

# The Health Costs of Firms

Vahid Moghani\*

April, 2026

## Abstract

How much do firms contribute to disparities in workers' healthcare expenditure? Using linked Dutch employer–employee administrative data, I exploit worker moves across firms to separate persistent worker differences from persistent firm differences. Moving the same worker to a firm one standard deviation higher in the firm-effect distribution raises expected annual healthcare expenditure by 17.8%. The implied top–bottom quintile gap is comparable to expenditure gradients by income, wealth, and education, and more than half of the dispersion remains within narrowly defined industrial codes. Supporting a health-risk interpretation, workers who move to higher-expenditure firms increase use of pain, anti-inflammatory, and muscle relaxant medications and face higher disability entry and long-run mortality. Because Dutch health insurance is largely financed outside the employment relationship, the results point to a fiscal externality: firms generate healthcare expenditure that they do not fully bear.

---

\*Erasmus School of Economics, Tinbergen Institute, and Erasmus Center for Health Economics Rotterdam. I thank Owen O'Donnell, Jan van Ours, Sonia Bhalotra, Tom Van Ourti, Hans van Kippersluis, Patrick Kline, Peter Hull, Jonathan Roth, Olivier Marie, Bettina Siflinger, Erik Schut, Lieke Beekers, Lisa Voois, Simon Rabaté and Shobhit Kulshreshtha for helpful comments and discussions. I gratefully acknowledge financial support from the Erasmus Initiative Smarter Choices for Better Health, which supported my position. I declare that there are no conflicts of interest.

## I Introduction

Firms affect workers' health through physical demands, psychosocial stressors, and management practices, and they also shape workers' use of healthcare.<sup>1</sup> Yet surprisingly little is known about whether otherwise similar workers incur systematically different healthcare expenditure depending on the firm that employs them. Such variation is important in its own right as a source of healthcare expenditure differences across workers, and it can also shed light on whether firms are a quantitatively relevant source of unequal workplace health risks. This question is particularly relevant in publicly financed healthcare systems, where part of the medical expenditure associated with workplace environments is borne outside the firm.<sup>2</sup>

Using linked employer–employee administrative data from the Netherlands, this paper delivers the first policy-interpretable estimates of how much annual healthcare expenditure varies across firms for otherwise similar workers. I use data on full-time male workers and leverage moves across firms to separate persistent worker differences from persistent firm differences, thereby recovering the distribution of firm effects in healthcare expenditure in a setting where basic health insurance is individual, universal, and portable across jobs rather than tied to employers, and where primary care is free at the point of use.<sup>3</sup> Cross-firm differences in expenditure are therefore less likely to reflect employer-driven differences in insurance coverage or cost sharing.

---

<sup>1</sup>Related work shows that firms affect worker health through workload and demand shocks, job loss, workplace safety, and employer-sponsored health investments (Hummels et al., 2025; Sullivan and Von Wachter, 2009; Johnson, 2020; Jones et al., 2019). Firms can also affect healthcare access through employer-sponsored insurance and provider networks (Tilipman, 2022; Ho and Lee, 2019; Bíró and Elek, 2025).

<sup>2</sup>The compensating-differentials literature uses wage–risk tradeoffs to price job-related health risk (Rosen, 1974; Marin and Psacharopoulos, 1982; Arnould and Nichols, 1983; Viscusi and Moore, 1987; Mas and Pallais, 2017; Sorkin, 2018) and shows that the cost of health-related benefits is largely shifted into wages (Gruber, 1994). It does not measure how realized healthcare expenditures vary across firms employing otherwise similar workers.

<sup>3</sup>Formally, I estimate a leave-one-out bias-corrected Abowd–Kramarz–Margolis (AKM) worker–firm two-way fixed effects model of log annual healthcare expenditure for full-time male workers from 2009–2016 (Abowd et al., 1999; Kline et al., 2020). I focus on men because women's job-to-job moves are closely tied to fertility decisions (Figure A7), which can violate the mobility assumptions behind the decomposition. The results for the female sample are reported in Appendix A8 and are similar to those for the male sample.

I document substantial variation across firms in healthcare expenditure among otherwise similar workers. Moving to a firm one standard deviation higher in the firm-effect distribution increases expected annual healthcare expenditure by 17.8%. In the study sample, the implied top–bottom quintile gap in firm effects is comparable in magnitude to gradients in healthcare expenditure by income, wealth, and education. More than 50% of this dispersion remains within narrowly defined industrial codes, suggesting that the variation is not simply between broad industries but also reflects differences across firms operating in similar lines of business.<sup>4</sup>

The main estimates are robust when restricting the sample to younger workers, for whom predictable health declines are less likely to drive job switches. Event-study evidence around job-to-job moves also shows no differential pre-trends in workers' healthcare expenditure prior to the move; together, these patterns are consistent with mobility not being driven by emerging health problems in aggregate. Moves to higher-expenditure firms are followed by increases in pain, anti-inflammatory, and muscle relaxant medications, with little corresponding change in categories more closely tied to screening or diagnosis, a pattern suggestive of workplace-related strain or injury. In the long-run, workers who move—including those displaced by firm closures—to higher-expenditure firms are more likely to enter disability and face higher mortality, suggesting that cross-firm variation in healthcare expenditure is linked to persistent differences in health risk. The firm gradient is also visible in GP expenditure and among high-spending workers who face near-zero marginal prices, suggesting that these differences are not solely driven by financial incentives.

This paper contributes to the literature on firm heterogeneity and inequality by adding healthcare expenditure and health to the set of worker outcomes shaped by firms. Existing work shows that firms pay different wage premia to otherwise similar workers and account for a substantial share of earnings inequality (Kline, 2024; Card et al., 2013, 2018, 2016), differ in amenities and working conditions (Sockin, 2022; Lamadon et al., 2022), and contribute to inequality through differential treatment by gender, origin, and other characteristics (Kline et al., 2022). The evidence on whether non-wage attributes offset or reinforce wage

---

<sup>4</sup>I use detailed 898 Dutch standard industrial classification (SBI) codes, which are comparable in spirit to SIC or NAICS industry classifications.

differences is mixed (Sorkin, 2022; Lamadon et al., 2022; Sorkin, 2018). I document that firms also generate substantial dispersion in healthcare expenditure with longer-run health implications, and that higher wage premia are associated with higher health-cost effects, consistent with compensating differentials.

The findings also connect the firm-heterogeneity literature to the health-inequality literature, where income, education, occupation, and place are documented as central sources of inequality (Cutler et al., 2006; Lleras-Muney et al., 2025; O'Donnell et al., 2015; Ravesteijn et al., 2018; Finkelstein et al., 2021, 2016). While earlier work identifies causal effects of specific workplace conditions on employee health (e.g., Hummels et al., 2025; Lee and Taylor, 2019), this paper shows that firm environments are a quantitatively important source of healthcare expenditure differences, with evidence that these differences are related to broader health risks.

Methodologically, this paper is the first to formally map the canonical mover-design event-study estimand into the variance components of the AKM decomposition (Finkelstein et al., 2016). Under standard AKM assumptions with homogeneous sorting, the coefficient converges in the large-cell limit to a standard-deviation share—the standard deviation of firm effects divided by the cross-firm standard deviation of average outcomes—and must therefore be squared to recover a variance share. This correction can materially change the inferred importance of firms or places: for example, mover-design coefficients of 0.5 in Finkelstein et al. (2016) and 0.3 in Ahammer et al. (2024) imply variance shares of 25% and 9%, respectively. While Lyubich (2025) notes that the mover estimand speaks to between-firm rather than within-firm dispersion, this paper derives that result formally.

The closest paper to this is Ahammer et al. (2024), contemporaneous and independent, who use Austrian linked data and a mover event-study design to document meaningful between-firm differences in healthcare expenditure. Their mover-design coefficient identifies a between-firm object: the standard deviation of firm effects relative to the cross-firm standard deviation of firm means. This paper instead estimates a leave-one-out bias-corrected AKM model and the associated variance decomposition, thereby recovering the dispersion of firm effects on individual annual healthcare expenditure, expressed as percentage differences for

otherwise similar workers across the firm-effect distribution (Kline, 2024). This yields an economically interpretable object, comparable to standard gradients by income, wealth, and education. I further document that these firm-induced expenditure differences predict medication use, disability, and mortality. Moreover, the Dutch setting strengthens interpretation because basic insurance is individual and portable across jobs, making cross-firm expenditure differences less likely to reflect employer-driven differences in coverage or cost sharing.

The results point to a fiscal externality in publicly financed healthcare systems such as the Dutch one. Because health insurance is financed outside the employment relationship, firms do not bear the marginal fiscal cost when their workers incur higher medical expenditure, which weakens incentives to reduce workplace health risks even under stringent safety regulation and experience rating in sickness and disability insurance. The fact that more than half of the dispersion remains within narrowly defined industrial codes suggests that this problem is not only about industry composition but also about firm-specific practices and organizational design. While wages may partly compensate workers for the disutility of the health risks, they do not reimburse the public financing system for the resulting medical costs. This creates an unpriced fiscal externality and points to the value of policies that better align firms' incentives with the healthcare expenditure they generate.

## II Background

This section summarizes the institutional features of the Dutch setting that matter for interpreting the estimates. Appendix A1 provides additional detail on the healthcare system, occupational-health obligations, sickness provisions, and the channels through which firms may affect employees' health and healthcare expenditures.

The Netherlands has a system of universal basic health insurance provided by private insurers under tight public regulation. Basic insurance is individual and portable across jobs: coverage does not change when workers change employers, and the statutory content of the basic package is the same across insurers. Cost sharing in the basic scheme is largely standardized through annual deductibles, but GP care is fully covered and exempt from the deductible. Access to care is

organized around the general practitioner, who serves as the first point of contact and gatekeeper for most non-urgent specialist and hospital treatment (Currie and Zwiers, 2025). For this paper, the key implication is that cross-firm differences in observed universal basic-package expenditure arise in a setting with limited employer-driven variation in insurance generosity, provider networks, or statutory cost sharing.

At the same time, Dutch employers face extensive obligations in occupational health and sickness management. Under the Working Conditions Act (*Arbowet*), firms must monitor workplace risks, contract occupational-health services or experts, provide access to a company doctor, and continue wage payments for up to two years during illness. These obligations are reinforced by experience rating in disability-related insurance. Yet day-to-day healthcare still runs primarily through the ordinary GP-based system, and occupational doctors provide work-related advice and reintegration support rather than medical treatment. The Dutch institutional setting therefore combines substantial employer responsibility for prevention and sickness management with limited employer control over the content of ordinary medical care.

Cross-firm differences in basic healthcare expenditure can arise through two broad channels: differences in workers' ability to seek care and differences in the health consequences of the workplace itself.<sup>5</sup> The Dutch institutional setting is particularly informative because several features narrow the scope for the first channel. Out-of-pocket payments for basic care are capped by the annual deductible, more than 90% of adults choose the minimum deductible, and the marginal price of additional basic care is zero for high spenders once the deductible is exhausted (Handel et al., 2020). Existing evidence also finds limited income effects on healthcare expenditure (Miller et al., 2024; Cesarini et al., 2016). Time barriers also appear modest. Collective agreements and labour law provide paid time off for sickness and medical visits, employers must continue wage payments during illness for up to two years, and reintegration is formally supervised by occupational physicians. Consistent with this, self-reported unmet medical need in the Netherlands is among the lowest in the EU and shows little income gradient

---

<sup>5</sup>Appendix A4 formalises this point in a Grossman-style framework (Grossman, 1972b), where individuals choose health investments given preferences, income, and time constraints, as well as a workplace-specific health-depreciation rate.

(e.g. [OECD and European Observatory on Health Systems and Policies, 2019, 2021](#)); lack of time is cited by less than 1% of respondents ([Eurostat, 2023](#)). Taken together, these features make the setting well suited to studying persistent expenditure differences across firms. In addition to the care-seeking channel, the workplace can directly affect employees' health. Relevant mechanisms include differences in safety and physical working conditions, psychosocial stress, environmental exposures, and health-related workplace norms ([Hamermesh, 1999](#); [Lee and Taylor, 2019](#); [Johnson, 2020](#); [Johansson et al., 2023](#); [Bhattacharya, 2014](#); [van Ours, 2019](#); [Hummels et al., 2025](#); [Blackburn et al., 2023](#); [Jolivet and Postel-Vinay, 2025](#); [Collaborators et al., 2018](#); [Loomis et al., 2018](#); [Tran et al., 2022](#); [Gan et al., 2011](#); [Royer et al., 2015](#); [Jones et al., 2019](#); [Simonsen and Skipper, 2025](#)).

### III Dataset Description, Sample Selection, and Descriptive Statistics

I use Dutch administrative employer–employee data covering all male employees living in the Netherlands.<sup>6</sup> The data contain information on employer–employee matches, demographics, income, contract start and end dates, and firms' Dutch standard industrial classification (SBI) codes.<sup>7</sup> I also observe workplace postal codes, individuals' residential locations over time, highest completed education, and annual total and component-level basic healthcare expenditure (general practitioner, medication, specialized mental healthcare, and hospital care).<sup>8</sup>

Employment information is available annually from 2009 to 2016, while health-care expenditure is observed annually from 2009 to 2019. I focus on male employees

---

<sup>6</sup>The corresponding results for female employees are reported in Appendix A8. For women, the patterns in Figure A7 indicate that firm changes are endogenous to childbearing decisions. Given the direct and indirect healthcare expenditure of childbearing ([Ahammer et al., 2023](#); [García-Gómez et al., 2026](#)), estimates based on firm moves should therefore be interpreted with caution for female workers.

<sup>7</sup>[Standaard Bedrijfsindeling \(SBI\)](#) is the Dutch standard industrial classification used by Statistics Netherlands and the Chamber of Commerce to assign each firm an industrial code. The hierarchy spans five levels, with up to five digits identifying increasingly specific industries (e.g. code 01131 “Growing of vegetables in open fields” versus code 01132 “Greenhouse growing of vegetables”). The analysis data contain 898 unique SBI codes.

<sup>8</sup>The baseline data do not contain occupation information. In Section V, I examine heterogeneity by education-proxied collar type and show that the estimated firm effects are largely shared across groups. Moreover, even within narrowly defined SBI industrial codes, I document substantial heterogeneity across firms, suggesting that the results are not driven solely by broad industry or occupation groups.

aged 25–65 who, in each year, work at least 11 months in jobs with a contract of at least 0.9 full-time equivalent.<sup>9,10</sup> When an individual is employed by multiple firms within the same year, I assign the worker to the firm where he spent the most months, since healthcare expenditure is observed annually.<sup>11</sup> These restrictions reduce reliance on functional-form assumptions relating work intensity to healthcare expenditure and simplify the assignment of workers to firms. A consequence is that workers who move late in the year remain assigned to their origin firm for that year.

For these full-time workers, I observe background characteristics such as birth year, migration background, and—for most individuals—education.<sup>12</sup> The data further include annual salary, employer identifier, employer sector, and firm size.

For each worker, I observe total and category-specific basic healthcare expenditure from 2009 to 2019. I observe all claims billed to the basic insurance package.<sup>13</sup> Total expenditure combines insurer reimbursements and individuals' deductible payments, which are not separately identified in the data. The categories are: (1) primary care (general practitioner services); (2) medication; (3) specialized mental healthcare outside primary care (psychology and psychiatry); and (4) hospital care, covering specialist treatments and inpatient admissions.<sup>14</sup> Employees in the defence sector (SBI code 8422) obtain healthcare through the Ministry of Defence, which operates a parallel system not captured in these data. I therefore exclude defence-sector firms and their employees, who account for about 1.5% of the sample.

The sample is restricted to firms connected through worker mobility, which is

---

<sup>9</sup>Full-time employment in the Netherlands ranges from 38 to 40 hours per week.

<sup>10</sup>Requiring at least 0.9 FTE excludes part-time workers, who constitute less than 20% of employed men between ages 25 and 65 ((CBS), 2025), and removes years in which employees shift into part-time work or leave employment. Sick employees continue to receive employer-paid wages for up to two years before entering sickness or disability schemes, so these periods remain recorded as employment, which is important for observing healthcare expenditure during sickness.

<sup>11</sup>The 0.9 FTE requirement effectively rules out simultaneous substantial contracts.

<sup>12</sup>Education information is complete for cohorts born after 1990. For older cohorts, education is observed only for those who interacted with administrative processes requiring educational documentation. Individuals with missing education are treated as a separate category. Results are robust to excluding these individuals.

<sup>13</sup>Care outside the basic package is therefore not observed. This is most relevant for short-term physiotherapy, which is often financed out of pocket or through supplementary insurance.

<sup>14</sup>Together, these categories represent about 90% of total basic healthcare expenditure; remaining expenditure includes categories such as physiotherapy and patient transport.

required for estimating the AKM model because it relies on an employer network linked by movers (Abowd et al., 1999). To address limited-mobility bias, I apply the leave-one-out correction proposed by Kline et al. (2020), which removes each employer–employee match iteratively during estimation. This procedure requires that the network remains connected after excluding any given match. As a result, the sample is further restricted to ensure connectivity under leave-one-out removal. Summary statistics for the main sample, the connected sample, and the movers sample appear in Table 1.

Table 1 shows that restricting the data to the connected set leaves the composition of the sample essentially unchanged. Relative to the initial full-time sample, workers in connected firms have very similar age, migration background, education, birth rates, and healthcare expenditure, with only small increases in average age (44.2 vs. 44.0), wages, firm size, and total healthcare expenditure. By contrast, within the connected set, movers differ systematically from stayers. Movers are on average about three years younger (42.6 vs. 45.6), somewhat better educated (higher shares with upper secondary, bachelor, and master degrees and fewer missing-education records), and earn slightly higher hourly wages. They are also more likely to experience a birth in a given year (5% vs. 3%) and have marginally fewer children on average, consistent with mobility being concentrated among younger workers. Movers tend to work in slightly smaller firms. Finally, movers have lower total healthcare expenditure, hospital expenditure, and specialized mental-health expenditure than stayers, consistent with their younger age and suggesting that movers are somewhat positively selected on health. The AKM framework assumes that firm effects are common across subgroups, so that effects identified from movers also apply to stayers and to workers with different observable characteristics; this restriction is not directly testable for stayers, who never switch firms. In Section V, I examine robustness of findings across subsamples defined by age, income, and baseline healthcare expenditure and show that the estimated firm effects and their dispersion are similar. In Section IV, I also provide additional empirical support for the common-firm-effect assumption.

For the main analysis, I model annual healthcare expenditure in logs. Under a multiplicative structure of health investments, such as Cobb–Douglas, taking

	<b>Initial Sample</b> (Full-time Employees) N = 18,705,218	<b>Connected Set</b> (Connected Firms) N = 16,976,071	<b>Stayers</b> (Connected Firms) N = 9,236,759	<b>Movers</b> (Connected Firms) N = 7,739,312
<b>Variable</b>	Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)
Age	43.99 (10.48)	44.24 (10.39)	45.60 (10.51)	42.62 (10.01)
Dutch background (dummy)	0.83 (0.38)	0.83 (0.37)	0.83 (0.37)	0.83 (0.37)
First-generation migrant (dummy)	0.10 (0.30)	0.10 (0.30)	0.10 (0.31)	0.09 (0.29)
Second-generation migrant (dummy)	0.07 (0.25)	0.07 (0.25)	0.06 (0.24)	0.07 (0.26)
Education: Elementary (dummy)	0.02 (0.16)	0.02 (0.15)	0.02 (0.15)	0.02 (0.15)
Education: Lower sec./MBO1 (dummy)	0.06 (0.24)	0.06 (0.24)	0.06 (0.24)	0.06 (0.24)
Education: Upper sec./MBO (dummy)	0.23 (0.42)	0.23 (0.42)	0.20 (0.40)	0.26 (0.44)
Education: Bachelor (HBO/WO) (dummy)	0.17 (0.37)	0.17 (0.38)	0.15 (0.36)	0.20 (0.40)
Education: Master/Doctorate (dummy)	0.11 (0.31)	0.11 (0.31)	0.10 (0.29)	0.12 (0.33)
Education: Missing (dummy)	0.41 (0.49)	0.41 (0.49)	0.46 (0.50)	0.34 (0.47)
Total # children	1.38 (1.23)	1.40 (1.23)	1.44 (1.24)	1.35 (1.21)
New child (dummy)	0.04 (0.20)	0.04 (0.20)	0.03 (0.18)	0.05 (0.22)
Gross payment of contract (€)	49,616 (46,120)	50,545 (46,821)	50,510 (47,553)	50,586 (45,932)
Mover (dummy)	0.44 (0.50)	0.46 (0.50)	0.00 (0.00)	1.00 (0.00)
Full-time factor (FTE)	0.99 (0.02)	0.99 (0.02)	1.00 (0.02)	0.99 (0.02)
Estimated hourly wage (€) <sup>a</sup>	24.32 (22.55)	24.76 (22.91)	24.48 (23.00)	25.09 (22.79)
Business size (employees)	692.46 (800.51)	741.76 (806.30)	763.30 (804.44)	715.98 (807.77)
Total healthcare expenditure (€)	1,143 (4,800)	1,149 (4,780)	1,260 (5,263)	1,015 (4,127)
Mental health (specialist) (€)	84.46 (1,434)	83.22 (1,389)	90.03 (1,573)	75.10 (1,130)
Hospital expenditure (€)	687.31 (4,040)	691.63 (4,036)	767.33 (4,451)	601.28 (3,474)
GP expenditure (€)	109.70 (69.00)	110.09 (68.58)	111.42 (70.71)	108.51 (65.92)
<b>Number of Observations</b>	18,705,218	16,976,071	9,236,759	7,739,312

<sup>a</sup> Hourly wage computed from gross contract payment assuming a full FTE works 52 weeks  $\times$  40 hours per week.

*Notes:* Universe of male employees aged 25–65 observed between 2009 and 2016. The sample is restricted to full-time jobs with FTE  $\geq$  0.9 and at least 11 months employed with the same main employer in a given year; if an individual is employed by multiple firms within the same year, the employer with the largest number of months is used. “Connected set” refers to employees in firms that belong to the leave-one-out connected component of the employer network, i.e. the set of firms that remain connected through worker moves when any single worker–firm match is removed. “Stayers” are worker–year observations for individuals who remain with the same employer throughout the 2009–2016 period; “Movers” are worker–year observations for individuals who change employer at least once over this period. Sectors are defined on the basis of the first two digits of the Dutch Standaard Bedrijfsindeling (SBI) classification, grouped into 17 broad sectors. Business size is measured as the midpoint of the categorical firm–size intervals reported in the data (interpreted as the approximate number of employees at the firm in the given year). Healthcare variables refer to annual expenditure under the basic health insurance package (insurer payments plus deductible). “Total healthcare expenditure” is the sum of all basic-package claims; “Mental health (specialist)” includes specialized mental healthcare outside primary care; “Hospital expenditure” cover specialist and inpatient hospital care; and “GP expenditure” refer to primary care delivered by general practitioners. Employees in the defence sector (SBI 8422), who obtain healthcare through a separate system, are excluded from all samples.

Table 1: Summary Statistics

logs yields additive individual and firm components.<sup>15</sup> This property is important

<sup>15</sup>A parsimonious Grossman-style health-capital model in Appendix A4 shows that, under multiplicative individual and firm components in health depreciation and care-seeking costs, log

for interpreting firm effects as “shared” across subgroups—such as blue- and white-collar workers, high- and low-income workers, or high- and low-spending individuals—because the same firm component on the log scale implies proportionally similar effects while preserving larger level effects for workers with higher individual expenditure.<sup>1617</sup>

A common concern with log transformations is the treatment of zero outcomes, since estimates can then depend on the measurement unit and on the mix of extensive- and intensive-margin responses (Chen and Roth, 2024; Mullahy and Norton, 2022). In my data, however, the share of observations with zero total healthcare expenditure is very small (Figure A1), so this issue is limited in practice. Using contemporaneous annual healthcare expenditure is also appropriate in this setting. Dutch employers must continue wage payments for up to two years during sickness, so periods of poor health remain attached to the firm where they arise. Moreover, Table A16 shows that firm effects on future healthcare expenditure are noticeably smaller than firm effects on contemporaneous expenditure.


Firms differ in the average basic healthcare expenditure of their workers, both within and across sectors (Figure 1). These raw differences, however, cannot be interpreted as firm effects. Workers and firms sort: lower-skilled and otherwise disadvantaged workers are more likely to be employed in low-wage, low-amenity firms (Card et al., 2013; Sockin, 2022), and health is strongly graded by socioeconomic status (O’Donnell et al., 2015; Danesh et al., 2024). As a result, firms employing workers with poorer underlying health or higher healthcare needs will mechanically display higher average healthcare expenditure even if the workplace itself has no effect on health. I therefore use within-worker changes in healthcare expenditure around firm moves to separate persistent worker components from firm components.

---


healthcare expenditure admits an additively separable decomposition into individual, firm, and age terms.

<sup>16</sup>Section V shows that firm effects and their variance are broadly similar across these subgroups, consistent with such a shared firm component.

<sup>17</sup>Additionally, the KSS correction is derived for OLS estimators in linear models; this motivates the use of a linear AKM specification in logs rather than nonlinear alternatives.



(a) Mean basic healthcare expenditure by industrial code (€, per worker-year)



(b) Distribution of firm-level mean basic healthcare expenditure (€, per worker-year)

Figure 1: Average annual basic healthcare expenditure sectoral and firm level variation

*Note:* Sample of male employees aged 25–65 in the connected set of firms used for the AKM estimation, observed between 2009 and 2016. Workers are restricted to full-time jobs with FTE  $\geq 0.9$  and at least 11 months with the same main employer in a given year; if an individual is employed by multiple firms within a year, the firm with the largest number of months is used. Healthcare expenditure refer to annual spending under the statutory basic health insurance package (insurer payments plus deductible) and exclude care covered only by supplementary insurance or in parallel systems (such as the defence sector, which is excluded). Sectors in panel (a) are defined from two-digit industrial codes grouped into 17 broad sectors. Panel (b) shows the cross-sectional distribution of firms’ mean annual basic healthcare expenditure per worker-year, where each firm’s mean is computed over all full-time male worker-year observations in the connected sample.

### III.I The mover design event study

A natural next step is therefore to examine healthcare expenditure among workers who move between firms with different average healthcare expenditure. This exercise is informative for two reasons. First, identification of firm effects from job moves requires that movers not exhibit systematic pre-trends in healthcare expenditure. Second, post-move changes in healthcare expenditure shed light on the dynamics through which the firm environment affects healthcare expenditure.

Mover-design event studies are motivated by the same additive worker-firm structure as the AKM model (Abowd et al., 1999; Finkelstein et al., 2016). Their usual interpretation is intuitive: if a worker moves to a firm whose coworkers have higher average outcomes and the worker’s own outcome rises proportionally,

firm-specific factors appear important; if the worker's own outcome does not change, they do not. This intuition is useful, but it is not yet precise enough for an AKM decomposition. The key point is that movers move, but they do not randomly reallocate across firms. Even within the mover sample, average coworker outcomes reflect both firm effects and worker sorting.

Consider the standard AKM structure in a simple two-period panel for an arbitrary outcome  $Y$ :

$$Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}, \quad (1)$$

where  $\alpha_i$  is a worker effect,  $\psi_{j(i,t)}$  is a firm effect, and  $\varepsilon_{it}$  is an idiosyncratic component. For movers, let

$$\Delta Y_i = Y_{i2} - Y_{i1}$$

denote the change in own outcomes, and let

$$\delta_i^{(-i)} = \bar{Y}_{d(i),2}^{(-i)} - \bar{Y}_{o(i),1}^{(-i)}$$

denote the leave-one-out difference in average coworker outcomes between the destination and origin firms. The first-difference mover-design slope is the population coefficient from regressing  $\Delta Y_i$  on  $\delta_i^{(-i)}$  among movers.

Appendix A3 shows that, under an AKM outcome structure, sufficiently large firms, and stable affine sorting—that is, a linear relationship up to a constant between worker and firm effects—this coefficient converges to a *standard-deviation share*: the standard deviation of firm effects divided by the cross-firm standard deviation of average firm outcomes among the firms visited by movers. Thus, the mover coefficient identifies a relative measure of how much firm effects contribute to cross-firm differences in average outcomes.

**Theorem 1.** Under the assumptions stated in Appendix A3, the mover-design coefficient converges to

$$\theta = \frac{\sigma_{\psi,M}}{SD_M(\bar{Y}_j)}, \quad (2)$$

where  $\sigma_{\psi,M}$  is the standard deviation of firm effects in the mover population, and  $SD_M(\bar{Y}_j)$  is the cross-firm standard deviation of average firm outcomes among the firms visited by movers.

**Proof.** Appendix A3 provides the full derivation, including the finite-firm expression with leave-one-out noise and the large-firm limit in (2).

The theorem clarifies the interpretation of the mover statistic. It is useful for assessing how much firm effects contribute to observed cross-firm dispersion in average outcomes, and the corresponding variance-share object is  $\theta^2$ . At the same time, the mover-design estimand does not by itself measure the contribution of firm effects to individual-level outcome variation.

This distinction matters for the interpretation of the results below. In this paper, I use the mover-design event study mainly to study pre-trends and timing: whether workers who move to higher-cost firms already exhibit higher own expenditure before the move, and whether expenditure adjusts at the time of the move. The main quantitative statement about how much firms matter for individual outcomes is based instead on the AKM model, which directly targets the relevant quadratic components.

Here, I focus on individuals who switch employer exactly once during 2009–2016 and work with age- and year-residualised log healthcare expenditure. Let  $\hat{h}_{it}$  denote the residual from a regression of  $h_{it}$  on age and calendar-year dummies. For each mover  $i$ , define

$$D_i \equiv \bar{\hat{h}}_{d(i)}^{(-i)} - \bar{\hat{h}}_{o(i)}^{(-i)},$$

where  $\bar{\hat{h}}_{o(i)}^{(-i)}$  and  $\bar{\hat{h}}_{d(i)}^{(-i)}$  are the leave-one-out mean residualised expenditure at the origin and destination firms. Thus,  $D_i$  measures how much higher or lower average residualised expenditure is at the destination firm relative to the origin.

Figure 2a relates the change in a mover’s own residualised expenditure between the pre- and post-move years to  $D_i$  nonparametrically.<sup>18</sup> Two features stand out. First, the relationship is close to linear: workers who move to firms with higher residualised coworker expenditure experience higher own expenditure after the move. Second, the fitted relationship has a positive intercept at  $D_i = 0$ : even movers between firms with similar average expenditure exhibit a modest increase in healthcare expenditure around the job change. This pattern is consistent with a common cost of moving—for example, due to job loss, or displacement—that

<sup>18</sup>Appendix Figure A6 reports a binscatter with the linear approximation and the distribution of  $D_i$  across movers. The figure focuses on the central range of  $D_i$ , where almost all moves lie; estimates are naturally less precise in the tails.

shifts expenditure for all movers, while the slope with respect to  $D_i$  captures how the firm environment shifts expenditure differentially across moves. In turn, the approximate linearity of the relationship is consistent with the log-additive specification motivated above.<sup>19</sup>

To study timing more systematically, I next estimate an event-study specification that follows one-time movers from four years before to four years after the switch. Using event times  $\mathcal{T} = \{-5, \dots, 5\}$  and taking  $t = -2$  as the reference year, I estimate

$$\hat{h}_{it} = \eta_i + \sum_{\tau \in \mathcal{T}, \tau \neq -2} \lambda_{\tau} \mathbf{1}\{t = \tau\} + \sum_{\tau \in \mathcal{T}, \tau \neq -2} \theta_{\tau} \mathbf{1}\{t = \tau\} D_i + u_{it}, \quad (3)$$

where  $\eta_i$  are individual fixed effects and  $\lambda_{-2} = \theta_{-2} = 0$ , so all coefficients are relative to event time  $t = -2$ .

Figure 2b plots the coefficients  $\theta_{\tau}$ , which trace how the mover design estimate unfolds over time. The event study delivers three main messages. First, the interaction coefficients  $\theta_{\tau}$  are close to zero and statistically flat in the years before the move, in line with no systematic pre-trends in healthcare expenditure that are correlated with the type of firm workers are about to join. This argues against selection into destination firms based on differential pre-move health shocks. Second, there is a sharp jump in  $\theta_{\tau}$  at the time of the move, and the new level is reached within the first few years. Workers who move to new firms quickly converge to the new expenditure level. Third, the coefficients show small (but existing) further dynamics beyond the first post-move year. This limited adjustment suggests that the main firm effect on healthcare expenditure is contemporaneous and that a time-invariant firm component is a good approximation for annual expenditure.

The nonparametric slope in panel (a) and the event-study evidence in panel (b) suggest that workers' healthcare expenditure shifts systematically when they move across firms, with effects that are approximately linear and realised quickly after

---

<sup>19</sup>By contrast, the origin–destination cell approach of Card et al. (2018) effectively imposes symmetry: the estimated effect of moving from firm  $a$  to firm  $b$  is constrained to be the mirror image of the effect of moving from  $b$  to  $a$ . In the setting of healthcare expenditure—where job loss and search themselves have health consequences—working at the individual level instead allows these common moving costs to show up as an intercept shift, while the slope with respect to  $D_i$  captures how the firm environment changes expenditure and provides a direct functional-form check. The four group split results are shown in Figure A8.

the move. Figure 4 suggests that moves to higher-expenditure firms are followed by increases in medication categories associated with pain, injury, or related physical strain. The figures also show little evidence of differential pre-trends or slow adjustment.<sup>20</sup> These patterns support the log-linear two-way fixed-effects specification used in the next section and suggest that the AKM firm effects capture contemporaneous differences in health-related expenditure that are plausibly linked to workplace environments, rather than artefacts of selective mobility.

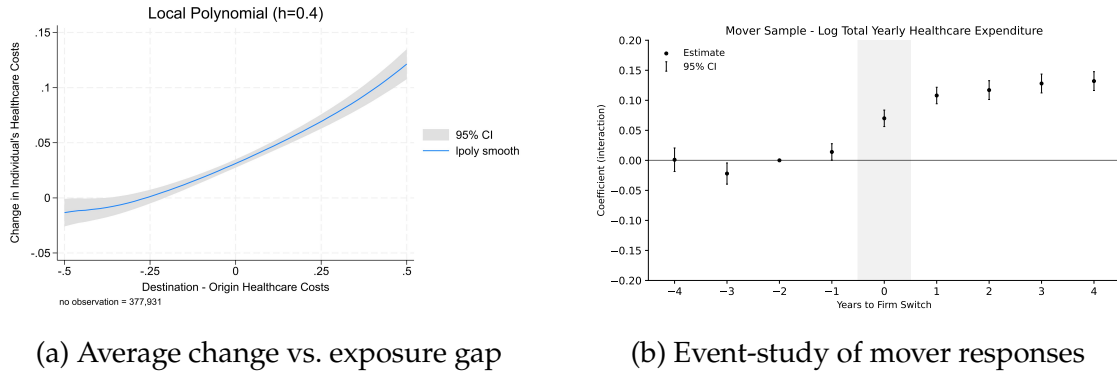


Figure 2: Workers' healthcare expenditure around moves between firms

*Note:* Sample of full-time male employees aged 25–65 who switch employer exactly once between 2009 and 2016 and belong to the connected set of firms. Healthcare outcomes are log total basic-package expenditures, residualised on age and calendar-year dummies. In panel (a), each point plots the average change in individual annual residualised expenditure between the years before and the years after the move against the “exposure gap”  $D_i$ , defined as the difference in leave- $i$ -out mean residualised expenditure between destination and origin firms; the line shows the (non-parametric) fitted relationship. In panel (b), the markers show coefficients from the mover event-study specification in equation (3), where the outcome is residualised log expenditure and the regressors are event-time indicators interacted with  $D_i$ , with event time measured relative to  $t = -2$  as the reference year.

## IV Method

To quantify how much firms contribute to variation in healthcare expenditure, I estimate a two-way fixed-effects model of log annual healthcare expenditure in

<sup>20</sup>The small gradual increase in healthcare expenditure over time suggests that some expenditure-inducing components may take longer to materialise.

the spirit of [Abowd et al. \(1999\)](#):

$$h_{it} = \alpha_i + \psi_{J(i,t)} + x'_{it}\beta + r_{it}, \quad (4)$$

where  $h_{it}$  is log basic-package healthcare expenditure (total paid by the insurer plus deductible paid by the individual) of worker  $i$  in year  $t$ ,  $J(i, t)$  indexes the firm employing  $i$  in  $t$ , and  $x_{it}$  collects time-varying observables (a 10-year split spline with a quadratic age profile by education group and calendar-year dummies). The term  $\alpha_i$  is a worker fixed effect capturing time-invariant determinants of expenditure such as genetic endowments, early-life health, persistent preferences for care, and long-run earnings capacity. The term  $\psi_{J(i,t)}$  is a firm fixed effect capturing the systematic impact of the workplace on employees' healthcare expenditure—for instance through workplace safety, physical and psychosocial risks, the organisation of work, and flexibility to seek care during work hours. The residual  $r_{it}$  captures remaining idiosyncratic shocks.

Annual basic-package healthcare expenditure is the main outcome because it is measured for every worker in every year and varies continuously rather than as a rare event. These features make it preferable, for the AKM analysis, to “hard” health outcomes such as mortality or disability, which are low-frequency tail events, observed at most once over the sample window and recorded as binary indicators. The outcome also maps naturally into a Grossman-style interpretation in which firms may affect expenditure through the workplace-specific health-depreciation rate as well as through other determinants of care use. [Appendix A4](#) formalises the log-linear specification in a simple Grossman-style life-cycle health-capital model. Under a multiplicative structure of health investments and care-use determinants, a log specification yields additive worker and firm components, so that a given difference in firm effects implies the same percentage difference in expenditure for workers with different baseline expenditure levels, while allowing larger level effects (in euros) for higher-spending workers.

#### *IV.I Identification and assumptions*

Equation (4) relies on several assumptions that I discuss here and assess empirically.

**Individual health trajectories.** The residual  $r_{it}$  captures, among other things, individual-specific health trajectories over the life cycle that are not fully absorbed by the flexible age and year controls in  $x_{it}$ . Following Card et al. (2013), I assume that any remaining worker-specific trend by age has zero mean and is not systematically related to firm type, conditional on the included age-by-education profiles and calendar-year dummies. This rules out unobserved individual characteristics that generate differential age patterns which are themselves correlated with the type of firm a worker joins. The rich age controls and the absence of pre-trends in the mover event studies (Figure 2) support this assumption: there is no evidence that workers who are about to move to high- versus low-cost firms already display diverging healthcare trajectories in the years before the move.

**Additive separability in logs.** The model assumes that worker and firm effects enter additively in logs, i.e. there are no systematic worker-firm interaction terms in  $r_{it}$ . To gauge the importance of such interactions, I estimate a saturated model with job-specific fixed effects, replacing  $\alpha_i + \psi_{J(i,t)}$  by a full set of worker×firm dummies, and compare the fit to the two-way fixed-effects model. As shown in Table 2, the  $R^2$  increases from 0.599 to 0.657 when moving from the AKM to the saturated specification, but the adjusted  $R^2$  rises only modestly, from 0.516 to 0.532, despite an increase of 56% in the number of parameters. Given that the variance of  $r_{it}$  is close to one, this implies that interactions explain only a small fraction of the residual variation.

Figure 3 plots the mean residuals  $\hat{r}_{it}$  from the AKM model by deciles of the estimated worker and firm effects. The averages are close to zero across almost all cells, with small deviations only in the cell combining the lowest worker and firm deciles, where truncation at zero and extensive-margin responses are most relevant. This pattern suggests that firm effects may differ somewhat for individuals at the very bottom of the expenditure distribution, but that additive separability in logs is a good approximation for the large majority of workers. Additionally, in Section V, I re-estimate firm effects separately for subsamples defined by education (blue- vs. white-collar), income, and baseline expenditure and show that the dispersion of firm effects and their correlation across subsamples are high, supporting the interpretation of a common firm component.

A further functional-form check comes from the mover design in Figure 2a. The non-parametric relationship between the change in an individual’s residualised log expenditure and the exposure gap  $D_i$  is approximately linear, with a common positive intercept at  $D_i = 0$ . The approximate linearity with a constant intercept is therefore consistent with the log-linear, additively separable structure in Equation (4).

Model	Two-Way Fixed Effects		Saturated with Job ID	
	$R^2$	Adjusted $R^2$	$R^2$	Adjusted $R^2$
	0.599	0.516	0.657	0.532
No. of Parameters	2,901,488		4,527,516	
# Person–Year Observations:	16,976,071			

*Note:* Sample of full-time male employees aged 25–65 in the connected set of firms, observed between 2009 and 2016 (# person–year observations = 16,976,071). The dependent variable is log annual basic-package healthcare expenditure (insurer payments plus deductible),  $h_{it} = \log(c_{it} + 1)$ . The two-way fixed-effects model (Eq. 4) includes worker and firm fixed effects and controls for age–by–education profiles and calendar-year dummies. The saturated model replaces worker and firm fixed effects with job-ID (worker×firm) fixed effects while keeping the same controls. The table reports  $R^2$ , adjusted  $R^2$ , and the total number of estimated parameters in each specification.

Table 2: Two-Way Fixed Effects vs. Saturated (Job-ID) Model

**Time-invariant worker and firm effects.** I treat  $\alpha_i$  and  $\psi_j$  as constant over the 2009–2016 period. Any transitory firm shocks or worker-specific changes in long-run health are absorbed into  $r_{it}$ . Given the relatively short panel and the limited dynamic adjustment after job moves shown in Figure 2b, this appears to be a reasonable approximation.<sup>21</sup>

**Selective mobility and pre-trends.** A central concern is that workers may move to firms with different health environments in anticipation of future health shocks, violating the implicit exogeneity condition that, conditional on  $\alpha_i$ , origin firm, and observables, the choice of destination is unrelated to unobserved changes in health. I provide several pieces of evidence against strong anticipatory mobility. First, the mover event study in Figure 2b, which interacts event-time indicators with the exposure gap  $D_i$  defined in Section III, shows that the coefficients  $\theta_\tau$  are flat and

<sup>21</sup>Allowing fully time-varying firm-specific health effects would require a longer panel and richer mobility than the data permit.

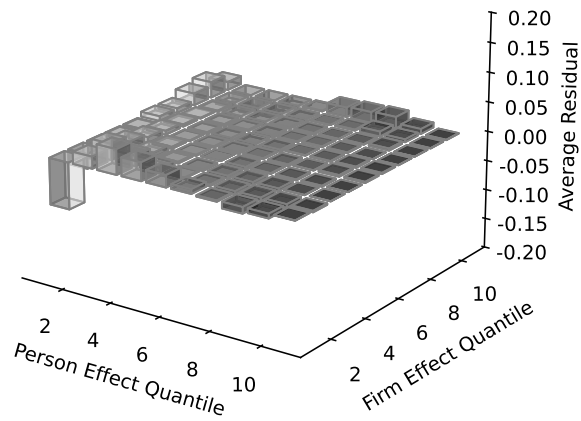


Figure 3: Mean residuals of the two-way fixed-effects model (Eq. 4) by decile of firm and individual fixed effects.

*Note:* Sample of full-time male employees aged 25–65 in the connected set of firms, observed between 2009 and 2016. The figure plots mean residuals from the two-way fixed-effects regression in Eq. 4, where the dependent variable is log annual basic-package healthcare expenditure,  $h_{it} = \log(c_{it} + 1)$ , and the controls are worker and firm fixed effects, age-by-education profiles, and calendar-year dummies. Residuals are averaged within 10×10 cells defined by deciles of the estimated worker and firm fixed effects (decile 1 = lowest, decile 10 = highest). The variance of the residuals is approximately 1, so the cell means shown in the figure represent small deviations relative to the overall idiosyncratic dispersion.

close to zero in the years before the move, while expenditure differentials open up sharply at the time of the move and then remain roughly constant. Second, Figure A8 reports pre- and post-move expenditure paths for workers who switch once between firms in the highest and lowest expenditure quartiles. Again, there is no evidence of diverging pre-move trends by move type. Finally, Table A7 shows that excluding workers above age 45, for whom future health declines are more common and predictable, leaves the estimated dispersion of firm effects essentially unchanged. This age split is consistent with life-cycle health-capital models in which individuals are close to full health throughout most of their prime working years and more rapid health deterioration sets in mainly at older ages (Galama and Van Kippersluis, 2019). Moreover, existing evidence for the Netherlands documents substantial worker and job dynamics, with job-to-job flows an important component of labour reallocation (e.g. Adema et al., 2025), making it plausible that many employer switches reflect career progress or job matching rather than responses to anticipated health problems. Taken together, these patterns suggest that selective mobility based on unobserved, destination-specific health trends is limited.

**Idiosyncratic shocks.** The remaining component of  $r_{it}$  reflects idiosyncratic, time-varying shocks to healthcare use that are not captured by the fixed effects or observed covariates. I assume that these shocks have zero mean and are not systematically related to firm type or the decision to move, conditional on  $\alpha_i$ ,  $\psi_j$ , and  $x_{it}$ . The large variance of  $r_{it}$  and the small average residuals across worker–firm cells in Figure 3 support this view that most of the residual variation reflects noise rather than systematic departures from the model structure.

#### IV.II *Object of interest and inequality interpretation*

The main object of interest is the dispersion of firm effects,

$$\sigma_\psi^2 \equiv \text{Var}(\psi_j), \quad \sigma_\psi = \text{sd}(\psi_j).$$

A move from firm  $A$  to firm  $B$  changes a worker's expected log expenditure by  $\Delta\psi = \psi_B - \psi_A$ , holding  $\alpha_i$  and  $x_{it}$  fixed, so that

$$\frac{\mathbb{E}[c_{it}^{(B)}]}{\mathbb{E}[c_{it}^{(A)}]} = \exp(\Delta\psi).$$

In particular, moving a worker from a firm to another firm one standard deviation higher in the firm-effect distribution changes expected expenditure by a factor  $\exp(\sigma_\psi)$ . I therefore report  $\sigma_\psi$  and  $\exp(\sigma_\psi) - 1$  as the primary measures of how much firms can shift individual healthcare expenditure.

One could instead normalise the variance of firm effects,  $\sigma_\psi^2$ , by the total variance  $\text{Var}(h_{it})$  and report variance shares (e.g. [Card et al., 2018](#)), or normalise the standard deviation of firm effects  $\sigma_\psi$  by the cross-firm standard deviation of firm means, as in the mover-design literature (e.g. [Finkelstein et al., 2016](#)) (see Appendix A3). As discussed by [Kline \(2024\)](#) and in Appendix A3, such relative measures can be difficult to compare across settings with different noise levels. Both the numerator and the denominator vary across time and contexts and, in the healthcare setting, the denominator is mechanically inflated by large idiosyncratic shocks at the individual-year level. A given variance share may therefore look small even when moving across the firm-effect distribution implies large percentage changes in individual expenditure. By contrast,  $\sigma_\psi$  and the implied factor  $\exp(\sigma_\psi)$  directly answer how much an individual's expected healthcare expenditure changes when their workplace changes.

This choice aligns closely with how health disparities is typically measured in the literature. For wages or incomes, univariate inequality indices (such as the variance or the Gini coefficient) have a natural welfare interpretation because the variable is monetary, cardinal, and directly comparable across individuals. For health expenditure, by contrast, such univariate indices are used much less frequently, in part because a large share of the dispersion reflects idiosyncratic biological risk and random health shocks<sup>22</sup>, and there is no clear normative benchmark

---

<sup>22</sup>In classic earnings data, regressions of log earnings on basic human-capital and labor-supply variables (schooling, experience) explain about 30% of the cross-sectional variance ( $R^2 \approx 0.3$ ; see [Mincer \(1974\)](#)), whereas analogous models for annual healthcare expenditure using only age, gender, region, and source of income explain on the order of 6% of the variance ( $R^2 \approx 0.06$ ; see [Van Kleef et al. \(2013\)](#)).

for what constitutes “acceptable” or unavoidable dispersion.<sup>23</sup> Instead, empirical work typically studies socioeconomic-related health disparities using bivariate measures based on the joint distribution of health and another variable such as income, education, or socio-economic status—for instance concentration indices, slope indices of disparities, or health gaps across income, education, or regional groups (e.g. [Kakwani et al., 1997](#); [Wagstaff, 2002](#); [Wagstaff and Doorslaer, 2004](#); [O’Donnell et al., 2015](#); [Finkelstein et al., 2016](#)). These measures answer questions of the form “by how much is health worse for individuals in disadvantaged positions?” rather than “what fraction of the total variance is explained by a given factor?”. In this sense,  $\sigma_\psi$  provides a firm-level analogue to these bivariate gradient measures, with firm type replacing income or education and healthcare expenditure replacing the health outcome.

#### IV.III *Limited mobility and leave-out variance correction*

Equation (4) can be estimated by OLS. However, variance components computed from the plug-in fixed-effect estimates,  $\{\hat{\alpha}_i, \hat{\psi}_j\}$ , are biased when worker mobility is limited. In employer networks where only few workers move between subsets of firms, each fixed effect is estimated with substantial noise even if it is unbiased individually. When these noisy estimates enter quadratically, as in  $\widehat{\sigma_\psi^2}$ , the variance of the estimation error inflates the quadratic form and leads to systematic bias ([Andrews et al., 2008](#); [Jochmans and Weidner, 2019](#); [Kline et al., 2020](#); [Bonhomme et al., 2023](#)). Intuitively, the plug-in variance of firm effects is biased upward because it adds true dispersion and noise dispersion. This problem is often referred to as *limited mobility bias*.<sup>24</sup>

To address this, I use the leave-out variance estimator proposed by [Kline et al. \(2020\)](#) (KSS). Their approach provides unbiased estimators for variance and covariance components of interest in linear models without imposing homoskedasticity or additional structural restrictions. Formally, any variance or covariance term can

<sup>23</sup>See, for example, [Wagstaff et al. \(1991\)](#), [Kakwani et al. \(1997\)](#), [Wagstaff and Doorslaer \(2004\)](#), [Erreygers and Van Ourti \(2011\)](#), and [Schlotheuber and Hosseinpoor \(2022\)](#), who distinguish overall from socioeconomic-related health disparities and show that empirical work and monitoring practice typically focus on the latter using concentration-type and related bivariate indices.

<sup>24</sup>The same logic applies to other quadratic terms such as  $Cov(\alpha_i, \psi_j)$ ; these are corrected using the same leave-out procedure.

be written as a quadratic form  $\theta = \nu' A \nu$ , where  $\nu$  stacks all model parameters and  $A$  is a known symmetric matrix. The naive plug-in estimator  $\hat{\theta}_{PI} = \hat{\nu}' A \hat{\nu}$  has a bias that can be expressed as a weighted sum of observation-specific error variances. KSS show how to estimate these error variances in a leave-one-out fashion and subtract the implied bias, yielding a bias-corrected estimator  $\hat{\theta}_{KSS}$ .

In practice, I implement the KSS correction at the worker–firm match level. Following [Kline et al. \(2020\)](#), I first partial out the observed covariates  $x_{it}$  by estimating Equation (4) and working with the residualised outcome  $h_{it} - x'_{it} \hat{\beta}$ . I then compute bias-corrected estimates of  $\sigma_\psi^2$  using the leave-one-out formulas of [Kline et al. \(2020\)](#), treating an entire worker–firm match as the leave-out unit to accommodate serial correlation within matches. Because recomputing the fixed effects separately for each leave-out observation is infeasible in a large panel, I follow KSS and use their analytic leave-out formulas and random-projection approximation based on [Achlioptas \(2003\)](#) to obtain computationally tractable estimates.<sup>25</sup> Appendix A5 provides the quadratic-form representation of the variance components and details of the implementation; see [Kline et al. \(2020\)](#) for a full derivation of the leave-out estimator. As a robustness check, I also re-estimate the variance components on a more strongly connected subset of firms; as expected, the plug-in bias is smaller but not negligible, and the KSS correction still reduces the estimated dispersion of firm effects (Table A15). These strongly connected firms are larger and more selected than the average employer and are therefore not representative of the full economy.

## V Results

Firms differ substantially in the healthcare expenditure they induce among otherwise similar workers. Table 3 reports KSS-corrected variance and covariance components for log annual healthcare expenditure  $h_{it}$  in Equation (4). The variance of firm effects is  $\text{var}(\psi_j) = 0.027$ , implying  $\sigma_\psi \approx 0.164$  log points. Moving a worker to a firm one standard deviation higher in the firm-effect distribution increases expected annual basic healthcare expenditure by  $\exp(\sigma_\psi) - 1 \approx 17.8\%$ .

---

<sup>25</sup>All AKM and leave-out computations are implemented in R using the package `LeaveOutKSS` developed by the author, which adapts the original Matlab routines of [Kline et al. \(2020\)](#) available as [GitHub: LeaveOutKSS-R](#).

The worker–firm correlation is close to zero ( $\rho(\alpha_i, \psi_j) = -0.018$ ), showing limited sorting on time-invariant expenditure.

To benchmark magnitudes, I compare the workplace gradient to the well-known gradients by wealth, income, and education that are standard in the health-disparities literature.<sup>26</sup> Within the sample, the change in (age- and year-adjusted) log healthcare expenditure from the bottom to the top quintile of the wealth distribution is about 0.26; for income, the corresponding difference is 0.29; and comparing workers with elementary or lower education to those with a master’s or doctorate yields a difference of 0.40 (Table A2). These gaps correspond to approximately 29%, 33%, and 49% higher annual basic healthcare expenditure, respectively. Using the estimated distribution of firm effects, the implied difference in expected expenditure between (same) workers employed at firms in the bottom and top quintile of the firm-effect distribution is about 0.46 log points, i.e. roughly 59% higher annual basic healthcare expenditure.<sup>27</sup> In this sample, the implied workplace gradient in expenditure is therefore comparable in magnitude to the observed gradients by wealth, income, and education.<sup>28</sup>

For completeness, Table 3 also reports variance shares: firm effects account for only about 1% of the total variance in  $h_{it}$  at the individual-year level. As discussed in Section IV, this variance share is not the most informative summary statistic,

---

<sup>26</sup>Highest completed education is used to define education; household wealth and total personal income are measured in 2009, before the main analysis period, to reduce concerns about reverse causality. These SES gradients are descriptive correlations rather than causal effects, and I report them only as a familiar benchmark for the magnitude of the firm-induced expenditure gradient.

<sup>27</sup>The “top–bottom quintile gap” is the difference in the mean firm effect between the top and bottom quintiles of the worker–year weighted latent distribution of  $\psi_j$  (Table A19). To translate the leave-one-out bias-corrected variance estimate of Kline, Saggio, and Sølvssten (KSS; Kline et al., 2020),  $\widehat{\sigma}_\psi^2$ , into this quintile-mean gap, a convenient benchmark is a normal approximation for  $\psi_j$ , which implies a gap of about 0.46 log points. Appendix A6 shows that correcting explicitly for sampling error in the AKM estimates  $\psi_j$  via deconvolution yields a very similar gap under a parametric benchmark with homoskedastic Gaussian estimation error, while a split-sample replicated-AKM deconvolution based on Kotlarski’s lemma recovers the latent distribution nonparametrically and allows for heterogeneous precision across firms; this yields a larger gap of 0.67 log points (Kotlarski, 1967; Anarat et al., 2025).

<sup>28</sup>The sample is positively selected in terms of socioeconomic characteristics: it consists of full-time male workers with relatively high income and wealth by construction (Figure A9), which likely underestimates the SES gradients in the Netherlands. In the same institutional context, Loef et al. (2021) show that healthcare expenditure is between 50% and 150% higher among individuals in the lowest education or income group than among those in the highest group across age bands. This suggests that, in the full population, SES-related gradients are larger than the 30–39% gaps observed in my sample.

because the denominator is dominated by large idiosyncratic individual-level shocks—such as cancers, accidents, and chronic conditions—that no employer can plausibly control and that are at least partly inherently random.

Statistic	KSS Estimates
Log Health Exp. Var.: $\text{var}(h_i)$	2.617 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.027 (1.03 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.516 (57.92 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	−0.004
Correlation $\rho(\alpha_i, \psi_j)$	−0.018
Explained share: workers & firms	0.587
No. Movers	1,225,799
No. Firms	132,820
No. of Observations	16,976,071
Average Health Exp. : $\text{mean}(h_i)$	5.608

*Notes:* Components are estimated via the leave-one-match-out Kline–Saggio–Sølvsten (KSS) correction (Kline et al., 2020) using random projections (50 simulations). Outcome is log healthcare expenditure per calendar year; sample spans full-time employees, 2009–2016. Shares in parentheses are relative to  $\text{var}(h_i)$ .

Table 3: Variance Decomposition of Log Healthcare Expenditure

### V.I Firm Fixed Effects and Health

Several results suggest that the estimated firm effects are more consistent with workplace-linked differences in health risk than with insurance design or short-run utilization incentives. Adding annual gross contract pay leaves both the dispersion and the ranking of firm effects essentially unchanged (Table A4). Dispersion remains sizeable among workers with high baseline expenditures, who typically exceed the deductible and face a zero marginal price for additional basic care, and firm rankings are strongly aligned across high- and low-baseline groups (Table A5). The same alignment holds when splitting by hourly wages (Table A6). Category-specific estimates point in the same direction. GP spending, which is fully covered and deductible-exempt, exhibits a firm-effect standard deviation close to that for total expenditure (Table A9). The patterns suggest that variation in firm effects is more consistent with differences in underlying health risk than with income-related differences in healthcare use.

To shed light on the underlying risk channels, I examine changes in medication use around firm moves using Equation (3). I consider both an indicator for any medication use and indicators for three subgroups: pain, anti-inflammatory, and muscle relaxant medications; cardiometabolic medications; and cancer-related medications. Consistent with the rise in healthcare expenditures for workers moving to higher-expenditure firms (Figure 2), Figure 4 shows an increase in overall medication use following such moves. This increase is not accompanied by changes in cardiometabolic or cancer-related medications, suggesting that differential screening or referral for previously undiagnosed conditions is unlikely to explain the firm-specific expenditure effects. Instead, the increase is concentrated in pain, anti-inflammatory, and muscle relaxant medications, a pattern more consistent with workplace strain, injury, or work-related stress than with broader changes in diagnosis intensity.

To further study the link between the estimated firm effects to health outcomes, I study two long-run measures observed in 2019. The first is an indicator for death by 2019, which captures cumulative severe health deterioration. The second is an indicator for disability receipt in 2019,<sup>29</sup> which captures persistent work-limiting health conditions. I relate these outcomes to firm effects on healthcare expenditure, conditioning on baseline demographics, pre-period expenditure, and origin-firm characteristics.

A practical concern is that  $\hat{\psi}_j$  is estimated from the same healthcare histories that also predict later disability and mortality, so single-sample estimation can mechanically link measurement error in  $\hat{\psi}_j$  to individual health shocks. I address this using the standard sample-splitting approach in the variance-components literature (Kline, 2025). I randomly split workers into two subsamples, estimate the AKM model separately in each,<sup>30</sup> and keep firms observed in both. Each worker is then assigned the (demeaned) firm effect estimated in the other subsample, which breaks the mechanical link between own shocks and the firm measure. Table 4 shows that workers with greater exposure to high-expenditure firms are more likely to be on disability and more likely to have died by 2019, conditional on migration background, postal-code fixed effects, and year-of-birth  $\times$  education

---

<sup>29</sup>Disability status is a truncated outcome: by construction, only individuals who are alive and not on other competing benefits or social welfare programmes can be on disability.

<sup>30</sup>See Table A18 for the KSS variance estimate within each subsample.

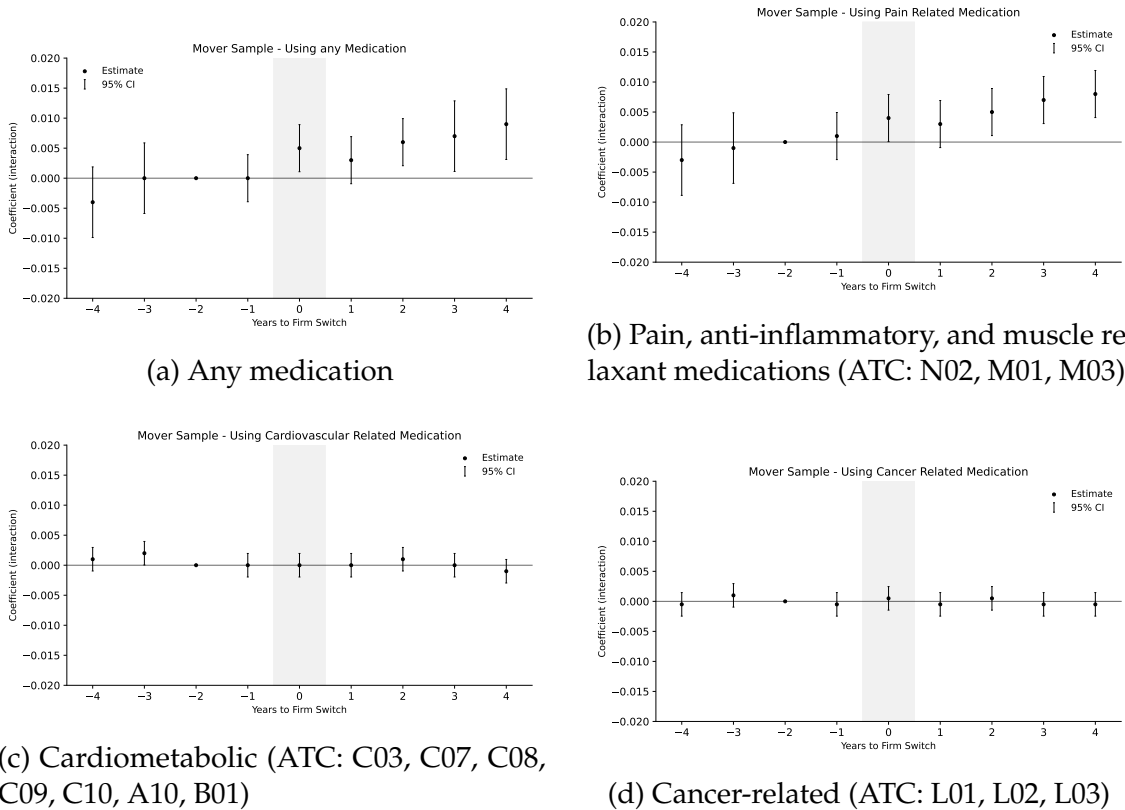


Figure 4: Medication use around moves between firms

*Note:* Sample of full-time male employees aged 25–65 who switch employer exactly once between 2009 and 2016 and belong to the connected set of firms. Each panel reports coefficients from the mover event-study specification in equation (3), where the outcome is an indicator for using (having any reimbursed dispensing of) the medication group in a calendar year, residualised on age and calendar-year dummies. Regressors are event-time indicators interacted with the exposure gap  $D_i$ , defined as the difference in leave- $i$ -out mean residualised total healthcare expenditure between destination and origin firms; event time is measured relative to  $t = -2$  (reference year). Markers show point estimates and vertical bars show 95% confidence intervals. Medication groups are defined by ATC prefixes: pain, anti-inflammatory, and muscle relaxant medications (N02 analgesics, M01 anti-inflammatory/antirheumatic products, M03 muscle relaxants; consistent with musculoskeletal strain and acute injury-related treatment), cardiometabolic medications (C03 diuretics, C07 beta blockers, C08 calcium channel blockers, C09 RAAS agents, C10 lipid modifiers, A10 diabetes drugs, B01 antithrombotics), and cancer-related medications (L01 antineoplastic agents, L02 endocrine therapy, L03 immunostimulants).

fixed effects. The relationship strengthens when I also control for the worker fixed-effect estimate from the AKM model, suggesting that the relationship is not driven only by time-invariant differences in baseline health.

I next use firm closures as an additional source of exogenous mobility. When

	Mortality in 2019		Disability in 2019	
	(1)	(2)	(3)	(4)
Average leave-out firm exposure	0.00111*** (0.00020)	0.00212*** (0.00020)	0.00377*** (0.00036)	0.00779*** (0.00035)
Worker FE control	No	Yes	No	Yes
Migration background controls	Yes	Yes	Yes	Yes
YOB × education FEs	Yes	Yes	Yes	Yes
Postal code FEs	Yes	Yes	Yes	Yes
<i>N</i>	2,485,666	2,485,666	2,432,508	2,432,508
<i>R</i> <sup>2</sup>	0.0417	0.0483	0.0445	0.0754

*Notes:* Each column reports a linear probability model estimated. The regressor is the worker-level panel average of the leave-out (sample-split) firm effect on log healthcare expenditure, computed over the years the worker is observed between 2009–2016. All specifications include migration background indicators, postal code fixed effects, and year-of-birth × education-category fixed effects. Columns (2) and (4) additionally control for the worker fixed effect estimate. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ .

Table 4: Average Exposure to High-Expenditure Firms and Subsequent Mortality and Disability (2019)

a firm exits, all workers separate in the same year and move to destination firms with different estimated health-cost effects. I focus on workers observed from the start of the panel and followed continuously, allowing me to trace healthcare trajectories around displacement. Following the displacement literature (e.g., [Halla et al. \(2020\)](#)), I define closure-induced job loss as a transition in which the employer identifier disappears from the panel and take the closure year as the last year the firm is observed. Figure [A11](#) shows that, after reemployment, healthcare expenditure shifts toward the average level of the destination firm.

Table 5 reports regressions of disability and mortality in 2019 on the leave-out estimated destination firm effect, controlling for baseline demographics, pre-closure average healthcare expenditure, and the origin firm’s health-cost effect.<sup>31,32</sup> The magnitudes are economically meaningful. Using the baseline AKM dispersion ( $\sigma_\psi = 0.164$ ), reemployment at a destination firm one standard deviation higher

<sup>31</sup>See Table [A10](#) for summary statistics of the closure sample.

<sup>32</sup>To reduce noise in the destination measure, I restrict the sample to destination firms above the 25th percentile of the firm-size distribution (at least 75 FTE workers). Origin and destination firm effects are highly correlated ( $\rho \approx 0.64$ ), so specifications with origin-firm fixed effects rely on limited within-origin dispersion and are correspondingly less precise and less robust. Results with origin fixed effects are similar but less precisely estimated.

in the health-cost firm-effect distribution is associated with a 0.22 percentage point higher probability of disability in 2019 and a 0.08 percentage point higher probability of death, roughly 6% of the respective sample means.<sup>33</sup> Figure 5 provides a complementary visualization: within broad origin bins, (residualized) mortality and disability are higher for workers reemployed at high-health-cost destination firms than for those reemployed at low-health-cost destinations.<sup>34,35</sup>

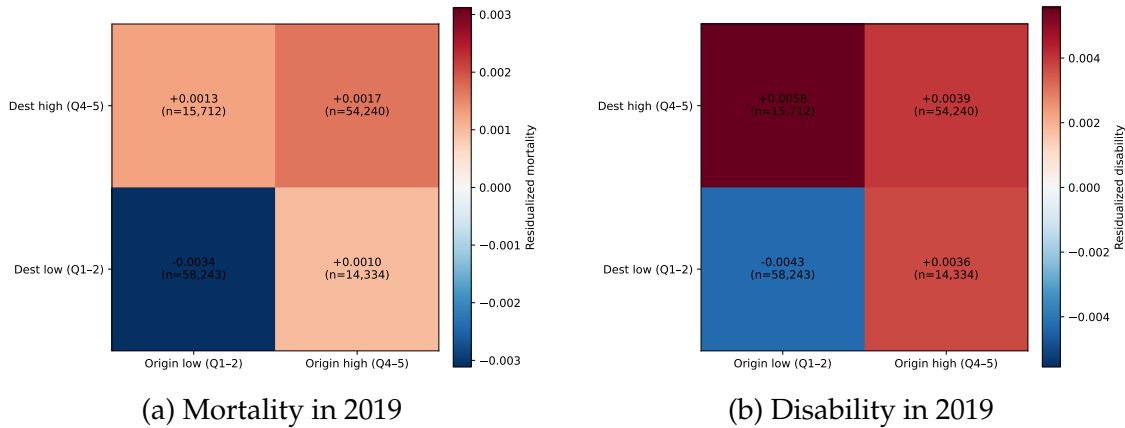


Figure 5: Long-run outcomes by origin and destination firm-effect bins after firm closures

*Note:* Sample of displaced male workers following firm closures, restricting destination firms to those above the 25th percentile of firm size (at least 75 employees). Each panel plots binned means of residualized outcomes, where residuals are obtained from the same control set used in Table 5 (baseline demographics, pre-move average healthcare expenditure, and pre-move postal-code fixed effects). The horizontal and vertical axes group workers by the origin and destination firm health-cost effect, respectively, into coarse bins (low vs. high). Cell annotations report the number of observations used to compute each binned mean. Appendix Figure A12 shows the same relationships using a more granular 5×5 (quintile-by-quintile) binning.

To characterise what the firm effects capture within healthcare expenditure

<sup>33</sup>These effect sizes use the preferred estimates in column (2) of Table 5 and the dependent-variable means reported in the table.

<sup>34</sup>Figure A12 shows a similar relationship using a finer 5×5 (quintile-by-quintile) grouping.

<sup>35</sup>As a placebo, I re-estimate the preferred specification using average pre-move healthcare expenditure as the outcome and the same covariate set as in Table 5. The destination-firm coefficient is not positive; if anything, it turns negative once I condition on the origin firm’s health-cost effect. Because origin and destination firm effects are highly correlated in the closure sample, placebo estimates that condition on origin-firm fixed effects should be interpreted cautiously. In any case, the lack of a positive relationship between destination fixed-effect and pre-move healthcare expenditures reduces concerns that the documented relationship reflects pre-existing differences in health (Table A11).

	(1)	(2)	(3)
<i>Outcome measured in 2019</i>			
<b>Panel A: Disability in 2019</b>			
Destination firm effect (leave-out)	0.0078*** (0.00133)	0.0137*** (0.00192)	0.0086*** (0.00292)
Mean of dep. var.		0.03929	
<i>N</i>	206,305	206,305	204,012
<i>R</i> <sup>2</sup>	0.066	0.066	0.116
<b>Panel B: Mortality in 2019</b>			
Destination firm effect (leave-out)	0.0033*** (0.00078)	0.0047*** (0.00113)	0.0034* (0.00173)
Mean of dep. var.		0.01313	
<i>N</i>	209,235	209,235	206,945
<i>R</i> <sup>2</sup>	0.048	0.048	0.083
Baseline controls & pre-move postal-code FEs	Yes	Yes	Yes
Origin firm-effect level (linear control)	No	Yes	No
Origin-firm fixed effects	No	No	Yes

*Notes:* Sample of displaced male workers following firm closures, restricting destination firms to those above the 25th percentile of firm size (at least 75 employees). The key regressor is the leave-out (sample-split) destination-firm effect on log healthcare expenditure. Column (2) additionally controls for the origin firm-effect level; column (3) includes origin-firm fixed effects. Baseline controls include pre-move average healthcare expenditure and demographics (as in the text) and pre-move postal-code fixed effects. Standard errors in parentheses. Significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Table 5: Firm Closures: Destination Firm Health-Cost Effects and Long-Run Outcomes (Continuous Measure)

and related health outcomes, I decompose expenditure by care category and relate firm exposure to cause-specific mortality. Dispersion in firm effects is present across care domains, and it is largest for hospital and specialist somatic care, smaller for GP services and medication—suggesting that firm effects are present for both severe and more routine care categories—and modest for specialised mental healthcare (Table A9). In the mortality data, exposure to high-expenditure firms is most strongly associated with subsequent cancer deaths, whereas links with injury-related<sup>36</sup>, cardiovascular, and other causes are smaller and estimated less precisely (Table A12)<sup>37</sup>. These patterns, along with the increase in the use of pain, anti-inflammatory, and muscle relaxant medications (Figure 4), suggest that workplace-related differences in healthcare expenditure are multidimensional and appear most closely linked to somatic and longer-run physical risks.

## *V.II Heterogeneity Across Sectors, Firms, and Workers*

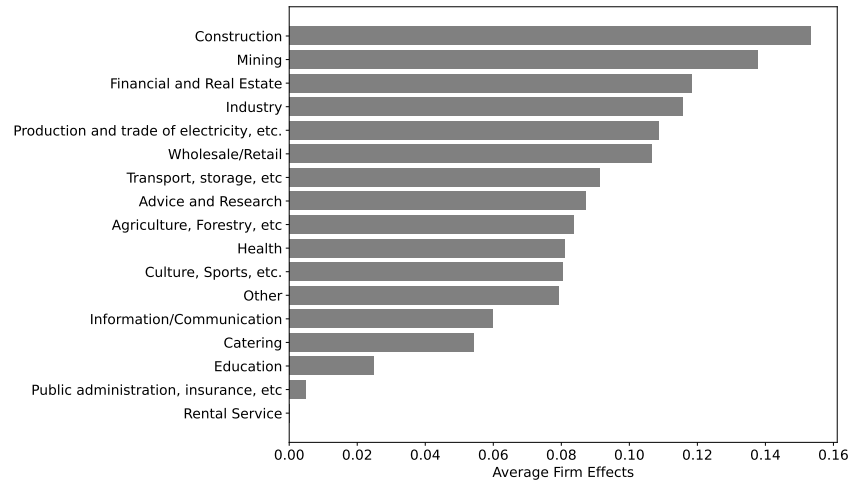
Because production technologies differ sharply across sectors, gaps between, for example, construction and finance in terms of induced health expenditure, mostly reflect what firms produce and leave limited scope for policy short of changing the activity mix. The more relevant question is how much dispersion remains among firms that ostensibly do the same thing. Figure 6 plots observation-weighted sector means of the workplace fixed effects. As expected, traditionally high-risk sectors such as mining, construction, and manufacturing are among the worst performers, and the financial sector—where work is less physically demanding but characterised by high pressure—also stands out with large adverse effects on employee healthcare expenditure, consistent with both physical and psychosocial channels. Using 17 broad sector classifications, sector differences account for only about one sixth of the dispersion in firm health effects. When I replace these broad groups with detailed, highly granular industrial codes, industrial-code

---

<sup>36</sup>Workplace mortality due to injuries is rare; in the Netherlands, annual fatal occupational accidents typically ranged between 23 and 45 cases in 2014–2020 [Centraal Bureau voor de Statistiek \(CBS\) \(2022\)](#).

<sup>37</sup>Cause-specific mortality is rare in the sample, so estimates for several causes are naturally imprecise. Cancer deaths are the most common cause-specific category in these data, whereas accidental deaths and other cause-specific deaths are much less frequent.

differences explain just under half of the dispersion,<sup>38</sup> so more than half of the dispersion remains within narrowly defined industrial-code cells. Thus, while some activities are clearly more hazardous than others, a large share of workplace-induced healthcare expenditure reflects differences between firms that do very similar work, pointing to firm-specific practices—and hence to scope for firm-level policy interventions—as a central driver.



*Notes:* The figure plots observation-weighted sector means of the estimated workplace fixed effects  $\psi_j$  from Equation 4. Higher bars indicate larger average impacts on employee healthcare expenditure. Sectors are ordered by magnitude. Because each sector aggregates many firms, sampling error in sector means is expected to be small. Sectors are defined using firms’ two-digit SBI industrial codes.<sup>a</sup> As the firm effects are identified only up to a constant, I normalise the average effect in the Rental Services sector to zero; all other sector means should be interpreted as differences relative to this baseline.

<sup>a</sup>Overzicht Standaard Bedrijfsindeling (SBI) codes are Dutch industrial codes used to classify firms by their main industry.

Figure 6: Average Sectoral Fixed Effects on total Healthcare Expenditure

To characterise which firms drive the estimated health-cost effects, I link the AKM firm effects to observable firm characteristics from the business registry

<sup>38</sup>The sector and industrial-code shares are computed in the subsample of firms that can be linked to the business registry, which removes about 12% of firms (around 1.4% of person-year observations). In this subsample, the overall variance of firm effects is smaller ( $\sigma_\psi \approx 0.11$  instead of 0.16; see Table A13), so the same sector differences translate into somewhat larger variance shares than in the full connected sample. Because some SBI industrial-code cells contain relatively few firms and workers, the between-code component may also be slightly upward biased; the reported industrial-code shares should therefore be interpreted as upper bounds, and the within-code component as a conservative estimate of within-firm heterogeneity.

and to wage premia from an AKM model of earnings.<sup>39,40</sup> Figure 7 reports these relationships as simple bivariate regressions of the firm health effect on each characteristic (one at a time). Firms with larger workforces, higher female and migrant shares, and older workforces have lower health-cost effects, while firms with more overtime have higher health-cost effects. Firms that employ an in-house company doctor are also associated with lower health-cost effects. Finally, the health-cost ranking is positively related to wage premia: high-expenditure firms pay higher wage premia. These patterns are consistent with compensating differentials—wages partially offset workplace conditions that are costlier for health.

Subsample analyses show that firms' health-cost effects are largely common across workers, but with meaningful heterogeneity in magnitude. Re-estimating the AKM model separately by age, income, baseline expenditure, and collar type, the estimated firm effects line up closely across groups: firms that are high-expenditure for one group tend also to be high-expenditure for others, as reflected in the cross-sample slope regressions of firm effects reported in Appendix Tables A5–A8. At the same time, the dispersion of firm effects is systematically larger for younger, lower-income, low-spending, and blue-collar workers, consistent with these groups being more exposed to channels such as physical hazards, work intensity, or tighter budget and time constraints.<sup>41</sup>

### V.III Robustness Checks


A key identifying assumption is that workers do not change employers in anticipation of unobserved health deterioration. The mover event studies in Figures 2

---

<sup>39</sup>This follows the compensating-differentials tradition (e.g. Rosen, 1986), and connects to recent firm-level evidence on wage premia and amenities (e.g. Sorkin, 2018; Sockin, 2022; Lamadon et al., 2022).

<sup>40</sup>Firm wage premia are estimated using the same specification as in Equation (4), replacing log healthcare expenditure with log annual earnings and restricting to full-time workers; see Table A3. Hourly wages are approximated by dividing annual earnings by a full-time schedule. Standard errors use the leave-one-out procedure of Kline et al. (2020), which is robust to heteroskedasticity and to correlation induced by common estimation error in the firm effects.

<sup>41</sup>Blue- and white-collar status is proxied from fields of study using a text-based classification. See Appendix A7 for details. The mapping from majors to collar types uses a simple text-based classification implemented with a large language model, in line with recent evidence that such models match or exceed human accuracy in straightforward coding tasks (e.g. Gilardi et al., 2023; Chang et al., 2024).



*Notes:* The figure reports OLS coefficients from separate bivariate regressions where the dependent variable is the estimated firm fixed effect in (log) healthcare expenditures. Each coefficient estimates the univariate regression coefficient of the firm health-cost effect on the indicated firm characteristic. Points show point estimates; whiskers show 95% confidence intervals computed using KSS leave-one-out standard errors (Kline et al., 2020), which are robust to heteroskedasticity and to correlation induced by common estimation error in the firm effects. Significance codes: \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

Figure 7: Correlates of Firm Fixed Effects in Healthcare Expenditure

and [A8](#) show no systematic pre-trends in healthcare expenditure by the type of firm workers are about to join. To further reduce scope for health-driven mobility, [Table A7](#) re-estimates the AKM model and corresponding variance decomposition for workers aged  $\leq 45$ , a life-cycle phase in which health and healthcare use are typically stable ([Galama and Van Kippersluis, 2019](#)). The KSS-corrected dispersion of firm effects is close to the baseline estimate in [Table 3](#), suggesting that the main results are not driven by late-career, health-motivated job changes.

A separate concern is that estimated firm effects proxy geography or access rather than workplace conditions. In the Netherlands, job changes rarely coincide with home moves ( $< 18\%$ ; [Kronenberg and Carree, 2010](#)), but I nonetheless add residential four-digit postcode fixed effects and, in a further specification, control for workplace-area healthcare access at PC4 (GP density, distance to pharmacies, and nearby hospitals). [Table A14](#) shows that the estimated firm-effect dispersion is essentially unchanged, indicating that residential sorting and local access do not drive the findings.

Because variance components from the AKM decomposition are sensitive to limited-mobility bias, the baseline estimates use the leave-one-match-out correction of [Kline et al. \(2020\)](#). [Table A15](#) shows that restricting to a highly connected subset of firms (at least 50 worker moves) reduces, but does not eliminate, plug-in bias and comes at substantial sample loss, while the KSS-corrected dispersion remains stable. [Table A18](#) provides an additional validation by estimating firm effects in two random worker splits, and reporting a cross-split covariance (Panel B) that is very close to the full-sample KSS estimate of  $\text{Var}(\psi_j)$  in [Table 3](#), consistent with the baseline dispersion reflecting signal rather than estimation noise.

Finally, the baseline outcome uses concurrent annual expenditure, which is the natural margin for workplace influences in this setting: workers remain tied to the employer during sickness absence for up to two years, and in the Grossman framework contemporaneous shifts in the health-depreciation rate map into contemporaneous expenditure flows. Consistent with this, [Table A16](#) shows that firm-effect dispersion attenuates with the horizon when future expenditure is used as the dependent variable, reinforcing that the baseline estimates primarily capture contemporaneous workplace-induced differences in healthcare expenditure.

## VI Discussion and Conclusion

The findings show that firms are a quantitatively important source of differences in workers' healthcare expenditure. In the Dutch setting, where basic health insurance is universal, individual, and portable across jobs, these differences are unlikely to be driven by employer-specific insurance coverage or cost sharing. Moving the same worker to a firm higher in the firm-effect distribution raises expected healthcare expenditure by 17.8%, and exposure to high-expenditure firms predicts higher subsequent disability and mortality. The evidence therefore suggests that firms affect health-related costs.

The estimated firm effect should be interpreted as a composite measure of workplace-related health costs. It may reflect physical risks, psychosocial strain, work pace, task composition, safety practices, managerial quality, peer norms, and prevention or reintegration policies. Because job moves change several of these margins at once, and because detailed occupations, tasks, and workplace practices are not observed, the data do not allow a decomposition into separate mechanisms. The estimates instead capture the net effect of being employed at a particular firm on contemporaneous healthcare expenditure.

Several pieces of evidence point to workplace-related health risk as an important component of this net effect. High-effect firms are concentrated in sectors with plausibly demanding physical or psychosocial environments, and firm effects are correlated with observable markers of demanding work environments, such as overtime. Workers who move to higher-expenditure firms increase use of pain, anti-inflammatory, and muscle-relaxant medications, while comparable changes are not visible for cardiometabolic or cancer-related medications. Firm effects also predict disability and mortality, including among workers displaced by firm closures. These patterns are difficult to reconcile with a pure utilisation or access interpretation.

Alternative explanations appear less central. The estimated dispersion is stable after controlling for pay, residential location, and local healthcare access. It remains sizeable for workers who face near-zero marginal prices for additional care and for GP expenditure, which is exempt from the deductible. These results suggest that prices, access, screening, and broad utilisation spillovers are unlikely to be the

main drivers. Narrower peer effects in care-seeking may still matter, especially for musculoskeletal complaints, and peer-driven health behaviour could also play a role. But such channels are unlikely to fully explain the sharp changes in contemporaneous healthcare expenditure, and the downstream increases in disability and mortality.

The findings also sharpen the interpretation of compensating differentials. Firms with higher health-cost effects tend to pay higher wage premia, consistent with workers being partly compensated for less favourable health environments. Yet even if wages adjust, the medical costs generated by these environments are largely financed outside the employment relationship. In a system such as the Dutch one, this creates a fiscal externality: firms may not fully internalise the healthcare expenditure associated with the risks they generate. This does not mean that firms face no health-related costs. Dutch employers already bear substantial responsibilities through sickness pay, occupational-health obligations, reintegration requirements, and experience rating in disability insurance. However, basic healthcare expenditure remains financed collectively and firms do not internalise directly the healthcare costs they induce.

Several caveats are important. The outcome is healthcare expenditure under the statutory basic package, not health itself and not the full social cost of work-related health risks. Care financed through supplementary insurance is not observed. Some part of the estimated firm dispersion may also reflect differences in screening, referral, or care-seeking behaviour rather than underlying health.

Two directions for future research are especially important. The first is to identify and quantify the mechanisms through which firms affect health: physical strain, work intensity, schedule control, managerial practices, psychosocial stress, or peer behaviour. The second is to study whether stronger cost internalisation can reduce the dispersion documented here. Broader experience rating or more employment-linked financing of healthcare could strengthen prevention incentives, but it could also create distortions, including selection against less healthy workers. Understanding this trade-off is central for policy design.

Overall, the results suggest that the firm should be treated as an important unit in the study of health disparities. If workplace environments generate part of the observed differences in healthcare expenditure and later health outcomes, then

policies focused only on individual behaviour or healthcare access will miss an important source of inequality.

## References

- Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67(2), 251–333.
- Achlioptas, D. (2003). Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of Computer and System Sciences* 66(4), 671–687.
- Adema, Y., L. Bettendorf, E. Van Bezooijen, D. Freeman, and B. Wache (2025). Dynamics, productivity, and innovation in the dutch economy.
- Ahammer, A., U. Glogowsky, M. Halla, and T. Hener (2023). The parenthood penalty in mental health: Evidence from austria and denmark.
- Ahammer, A., A. Packham, and J. Smith (2024). Firms and worker health. National Bureau of Economic Research, Working Paper No. w32011.
- Anarat, A., J. Krutmann, and H. Schwender (2025). A nonparametric statistical method for deconvolving densities in the analysis of proteomic data. *arXiv preprint arXiv:2506.01540*.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008). High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society Series A: Statistics in Society* 171(3), 673–697.
- Arnould, R. J. and L. M. Nichols (1983). Wage-risk premiums and workers' compensation: a refinement of estimates of compensating wage differential. *Journal of Political Economy* 91(2), 332–340.
- Bhattacharya, A. (2014). Costs of occupational musculoskeletal disorders (msds) in the united states. *International Journal of Industrial Ergonomics* 44(3), 448–454.
- Blackburn, B., T. Chan, E. Cherot, R. B. Freeman, X. Hu, E. Matt, and C. A. Rhodes (2023). Beyond burnout: from measuring to forecasting. Working Paper w30895, National Bureau of Economic Research.
- Bojilov, R. and A. Galichon (2016). Matching in closed-form: equilibrium, identification, and comparative statics. *Economic Theory* 61(4), 587–609.

- Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics* 41(2), 291–322.
- Bíró, A. and P. Elek (2025). Firm quality and health maintenance. *Journal of Health Economics* 103, 103045.
- Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics* 36(S1), S13–S70.
- Card, D., A. R. Cardoso, and P. Kline (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics* 131(2), 633–686.
- Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics* 128(3), 967–1015.
- (CBS), S. N. (2025, November). Werkzame beroepsbevolking; arbeidsduur (statline table 85275ned). Data set. Retrieved 25 November 2025 from <https://www.cbs.nl/nl-nl/cijfers/detail/85275NED>.
- Centraal Bureau voor de Statistiek (CBS) (2022). Dodelijke arbeidsongevallen, 2014–2020. Accessed: 2026-04-10.
- Cesarini, D., E. Lindqvist, R. Östling, and B. Wallace (2016). Wealth, health, and child development: Evidence from administrative data on swedish lottery players. *The Quarterly Journal of Economics* 131(2), 687–738.
- Chang, Y., X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, et al. (2024). A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology* 15(3), 1–45.
- Chen, J. and J. Roth (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics* 139(2), 891–936.
- Collaborators, G. . R. F. et al. (2018). Global, regional, and national comparative risk assessment of 84 behavioural, environmental and occupational, and metabolic risks or clusters of risks for 195 countries and territories, 1990-2017: a systematic

- analysis for the global burden of disease study 2017. *The Lancet* 392(10159), 1923–1994.
- Currie, J. and E. Zwiers (2025). Medication of postpartum depression and maternal outcomes: Evidence from geographic variation in dutch prescribing. *Journal of Human Resources* 60(4), 1093–1125.
- Cutler, D., A. Deaton, and A. Lleras-Muney (2006). The determinants of mortality. *Journal of Economic Perspectives* 20(3), 97–120.
- Danesh, K., J. T. Kolstad, W. D. Parker, and J. Spinnewijn (2024). The chronic disease index: Analyzing health inequalities over the lifecycle. Technical report, National Bureau of Economic Research.
- De Groot, N. and P. Koning (2016). Assessing the effects of disability insurance experience rating. the case of the netherlands. *Labour Economics* 41, 304–317.
- Erreygers, G. and T. Van Ourti (2011). Measuring socioeconomic inequality in health, health care and health financing by means of rank-dependent indices: a recipe for good practice. *Journal of health economics* 30(4), 685–694.
- Eurostat (2023). Unmet health care needs statistics. Statistics Explained, European Commission. Accessed 24 November 2025.
- 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(4), 1681–1726.
- Finkelstein, A., M. Gentzkow, and H. Williams (2021). Place-based drivers of mortality: Evidence from migration. *American Economic Review* 111(8), 2697–2735.
- Galama, T. J. and H. Van Kippersluis (2019). A theory of socio-economic disparities in health over the life cycle. *The Economic Journal* 129(617), 338–374.
- Galichon, A. and B. Salanié (2022). Cupid’s invisible hand: Social surplus and identification in matching models. *The Review of Economic Studies* 89(5), 2600–2629.

- Gan, W. Q., H. W. Davies, and P. A. Demers (2011). Exposure to occupational noise and cardiovascular disease in the united states: the national health and nutrition examination survey 1999–2004. *Occupational and environmental medicine* 68(3), 183–190.
- García-Gómez, P., A. C. Gielen, C. Riumallo Herl, and J. Rossi (2026). Health after birth: The persistent health penalty of becoming a parent. Draft manuscript.
- Gilardi, F., M. Alizadeh, and M. Kubli (2023). Chatgpt outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences* 120(30), e2305016120.
- Godøy, A. and I. Huitfeldt (2020). Regional variation in health care utilization and mortality. *Journal of Health Economics* 71, 102254.
- Grossman, M. (1972a). *The Demand for Health: A Theoretical and Empirical Investigation*. Columbia University Press.
- Grossman, M. (1972b). On the concept of health capital and the demand for health. *Journal of Political Economy* 80(2), 223–255.
- Gruber, J. (1994). The incidence of mandated maternity benefits. *The American economic review*, 622–641.
- Halla, M., J. Schmieder, and A. Weber (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics* 12(4), 253–287.
- Hamermesh, D. S. (1999). Changing inequality in markets for workplace amenities. *The Quarterly Journal of Economics* 114(4), 1085–1123.
- Handel, B. R., J. T. Kolstad, T. Minten, and J. Spinnewijn (2020). The social determinants of choice quality: Evidence from health insurance in the netherlands.
- Ho, K. and R. S. Lee (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review* 109(2), 473–522.
- Hummels, D., J. R. Munch, and C. Xiang (2025). No pain, no gain: Work demand, work effort, and worker health. *Review of Economics and Statistics* 107(4), 898–913.

- Jochmans, K. and M. Weidner (2019). Fixed-effect regressions on network data. *Econometrica* 87(5), 1543–1560.
- Johansson, N., S. C. de New, J. S. Kunz, D. Petrie, and M. Svensson (2023). Reductions in out-of-pocket prices and forward-looking moral hazard in health care demand. *Journal of Health Economics* 87, 102710.
- Johnson, M. S. (2020). Regulation by shaming: Deterrence effects of publicizing violations of workplace safety and health laws. *American economic review* 110(6), 1866–1904.
- Jolivet, G. and F. Postel-Vinay (2025). A structural analysis of mental health and labour market trajectories. *Review of Economic Studies* 92(3), 1920–1954.
- Jones, D., D. Molitor, and J. Reif (2019). What do workplace wellness programs do? evidence from the illinois workplace wellness study. *The Quarterly Journal of Economics* 134(4), 1747–1791.
- Kakwani, N., A. Wagstaff, and E. Van Doorslaer (1997). Socioeconomic inequalities in health: measurement, computation, and statistical inference. *Journal of econometrics* 77(1), 87–103.
- Kline, P. (2024). Firm wage effects. In *Handbook of Labor Economics*, Volume 5, pp. 115–181. Elsevier.
- Kline, P., E. K. Rose, and C. R. Walters (2022). Systemic discrimination among large u.s. employers. *The Quarterly Journal of Economics* 137(4), 1963–2036.
- Kline, P., R. Saggio, and M. Sølvssten (2020). Leave-out estimation of variance components. *Econometrica* 88(5), 1859–1898.
- Kline, P. M. (2025). Branching fixed effects: A proposal for communicating uncertainty. Technical report, National Bureau of Economic Research.
- Koning, P. (2009). Experience rating and the inflow into disability insurance. *De Economist* 157(3).
- Kotlarski, I. (1967). On characterizing the gamma and the normal distribution. *Pacific Journal of Mathematics* 20(1), 69–76.

- Kronenberg, K. and M. A. Carree (2010). Job and residential mobility in the netherlands: The influence of human capital, household composition and location. Technical Report MPRA Paper No. 25840, Munich Personal RePEc Archive (MPRA). Based on Statistics Netherlands (CBS) register microdata; full-time employees ( $\geq 0.8$  FTE), selected sectors, Sept. 2003–Sept. 2004. Reports counts used to compute that about 17.5% of job changes coincide with a residential move.
- Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review* 112(1), 169–212.
- Lee, J. M. and L. O. Taylor (2019). Randomized safety inspections and risk exposure on the job: Quasi-experimental estimates of the value of a statistical life. *American Economic Journal: Economic Policy* 11(4), 350–374.
- Lleras-Muney, A., H. Schwandt, and L. R. Wherry (2025). Poverty and health. *Annual Review of Economics* 17(1), 31–56.
- Loef, B., I. Meulman, G.-C. M. Herber, G. J. Kommer, M. A. Koopmanschap, A. E. Kunst, J. J. Polder, A. Wong, and E. Uiters (2021). Socioeconomic differences in healthcare expenditure and utilization in the netherlands. *BMC Health Services Research* 21(1), 643.
- Loomis, D., N. Guha, A. L. Hall, and K. Straif (2018). Identifying occupational carcinogens: an update from the iarc monographs. *Occupational and environmental medicine* 75(8), 593–603.
- Luan, W. P., R. Leal, J. S. Zhou, and J. S. Skinner (2025). Place-based variation in health care: Evidence from mandatory movers in the us military health system. Technical report, National Bureau of Economic Research.
- Lyubich, E. (2025). The role of people versus places in individual carbon emissions. *American Economic Review* 115(5), 1439–1484.
- Marin, A. and G. Psacharopoulos (1982). The reward for risk in the labor market: evidence from the united kingdom and a reconciliation with other studies. *Journal of Political Economy* 90(4), 827–853.

- Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review* 107(12), 3722–3759.
- Miller, S., E. Rhodes, A. W. Bartik, D. E. Broockman, P. K. Krause, and E. Vivalt (2024). Does income affect health? evidence from a randomized controlled trial of a guaranteed income. Technical Report w32711, National Bureau of Economic Research.
- Mincer, J. A. (1974). The human capital earnings function. In *Schooling, experience, and earnings*, pp. 83–96. NBER.
- Mullahy, J. and E. C. Norton (2022). Why transform y? a critical assessment of dependent-variable transformations in regression models for skewed and sometimes-zero outcomes. Working Paper w30735, National Bureau of Economic Research.
- Nederlandse Arbeidsinspectie (2019). Arbo in bedrijf 2018: Naleving van arboverplichtingen en blootstelling aan risico's. Technical report, Nederlandse Arbeidsinspectie, Den Haag. Rapport en tabellenboek.
- OECD and European Observatory on Health Systems and Policies (2019). Netherlands: Country health profile 2019. Technical report, OECD Publishing, Paris.
- OECD and European Observatory on Health Systems and Policies (2021). Netherlands: Country health profile 2021. Technical report, OECD Publishing, Paris.
- O'Donnell, O., E. Van Doorslaer, and T. Van Ourti (2015). Chapter 17 - health and inequality. In A. B. Atkinson and F. Bourguignon (Eds.), *Handbook of Income Distribution*, Volume 2, pp. 1419–1533. Elsevier.
- Prudon, R. (2025). Is delayed mental health treatment detrimental to employment? *Review of Economics and Statistics*, 1–46.
- Ravesteijn, B., H. v. Kippersluis, and E. v. Doorslaer (2018). The wear and tear on health: What is the role of occupation? *Health Economics* 27(2), e69–e86.
- Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy* 82(1), 34–55.

- Rosen, S. (1986). Chapter 12 the theory of equalizing differences. In *Handbook of Labor Economics*, Volume 1, pp. 641–692. Elsevier.
- Royer, H., M. Stehr, and J. Sydnor (2015). Incentives, commitments, and habit formation in exercise: evidence from a field experiment with workers at a fortune-500 company. *American Economic Journal: Applied Economics* 7(3), 51–84.
- Schlotheuber, A. and A. R. Hosseinpoor (2022). Summary measures of health inequality: a review of existing measures and their application. *International journal of environmental research and public health* 19(6), 3697.
- Shreekumar, A. and P. L. Vautrey (2022). Managing emotions: The effects of online mindfulness meditation on mental health and economic behavior. Technical report, MIT.
- Simonsen, M. and L. Skipper (2025). Healthy at work? evidence from a social experimental evaluation of a firm-based wellness program. *American Economic Journal: Economic Policy* 17(3), 471–495.
- Sockin, J. (2022). Show me the amenity: Are higher-paying firms better all around?
- Sorkin, I. (2018). Ranking firms using revealed preference. *The quarterly journal of economics* 133(3), 1331–1393.
- Sullivan, D. and T. Von Wachter (2009). Job displacement and mortality: An analysis using administrative data. *The Quarterly Journal of Economics* 124(3), 1265–1306.
- Teulings, C. N. (1995). The wage distribution in a model of the assignment of skills to jobs. *Journal of political Economy* 103(2), 280–315.
- Tilipman, N. (2022). Employer incentives and distortions in health insurance design: Implications for welfare and costs. *American Economic Review* 112(3), 998–1037.
- Törnberg, P. (2023). Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning.

- Tran, K. B., J. J. Lang, K. Compton, R. Xu, A. R. Acheson, H. J. Henrikson, J. M. Kocarnik, L. Penberthy, A. Aali, Q. Abbas, et al. (2022). The global burden of cancer attributable to risk factors, 2010–19: a systematic analysis for the global burden of disease study 2019. *The Lancet* 400(10352), 563–591.
- Van Kleef, R. C., R. C. Van Vliet, and W. P. Van de Ven (2013). Risk equalization in the netherlands: an empirical evaluation. *Expert Review of Pharmacoeconomics & Outcomes Research* 13(6), 829–839.
- van Ours, J. (2019). Health economics of the workplace: Workplace accidents and effects of job loss and retirement. In *Oxford Encyclopedia of Health Economics*. Oxford University Press.
- Vektis (2016). Verzekerden in beeld 2016. Technical report, Vektis.
- Viscusi, W. K. and M. J. Moore (1987). Workers' compensation: Wage effects, benefit inadequacies, and the value of health losses. *The Review of Economics and Statistics*, 249–261.
- Wagstaff, A. (2002). Inequality aversion, health inequalities and health achievement. *Journal of health economics* 21(4), 627–641.
- Wagstaff, A. and E. v. Doorslaer (2004). Overall versus socioeconomic health inequality: a measurement framework and two empirical illustrations. *Health economics* 13(3), 297–301.
- Wagstaff, A., P. Paci, and E. Van Doorslaer (1991). On the measurement of inequalities in health. *Social science & medicine* 33(5), 545–557.

# Appendices

## A1 Institutional Background Details

This appendix provides additional institutional detail for the background discussion in Section II.

### A1.I Dutch Healthcare and Insurance System

The Netherlands has a system of universal basic health insurance provided by private insurers under tight public regulation. All residents must purchase the basic package. Insurers must accept all applicants and charge community-rated premiums, while a national risk-equalization scheme compensates insurers for differences in enrollee risk. Financing is supplemented by tax revenue.<sup>42</sup> Basic insurance is individual and portable: coverage does not change when workers change jobs. Employers may offer collective contracts, sometimes with premium discounts of up to about 10%, but these contracts do not alter the statutory content or cost-sharing rules of the basic package (Handel et al., 2020).

Cost sharing in the basic scheme is largely standardized through annual deductibles. In 2015, all adults faced a compulsory deductible of €375 and could choose an additional voluntary deductible of up to €500. The deductible applied to most specialist and hospital care and to many prescription drugs,<sup>43</sup> but not to GP care, maternity care, or several preventive services. More than 90% of adults, including almost all individuals with high expected healthcare expenditure, chose the minimum deductible (Handel et al., 2020). Switching between insurers was also infrequent: only 6.8% of individuals changed insurer in 2015 (Handel et al., 2020).

Access to care is organized around the general practitioner. Individuals are required to register with a GP, typically near their residence (Currie and Zwiers, 2025). GPs provide the first point of contact for non-emergency care and act as gatekeepers for almost all non-urgent specialist and hospital treatment. GP consultations are fully covered and exempt from the deductible. Pharmacies are

---

<sup>42</sup>In 2015, the annual premium for a standard policy with the minimum deductible was around €1,200; low-income individuals received a healthcare benefit that offset part of the premium and deductible (Handel et al., 2020).

<sup>43</sup>Prescription-drug cost sharing in the Dutch basic scheme is not identical across all medications, since some drugs can also be subject to drug-specific co-payments or reimbursement limits in addition to the general deductible.

widespread and usually located near GP practices. This organization makes basic access to primary care fairly uniform across workers and largely independent of the employer.

The content of the basic package is set by the government and is the same across insurers. It covers medically necessary GP care, specialist and hospital services, most prescription drugs, maternity and neonatal care, emergency transport, mental health services,<sup>44</sup> and selected medical aids. Routine adult dental care, eyeglasses and contact lenses, and many elective or alternative treatments are excluded. Physiotherapy is reimbursed only for chronic conditions and only after the first 20 sessions, so most short-term physiotherapy is either paid out of pocket or financed through supplementary insurance.<sup>45</sup>

For this paper, the key implication is that basic insurance in the Netherlands is not attached to the employment relationship. Cross-firm differences in observed basic-package expenditure therefore arise in a setting with limited employer-driven variation in insurance generosity, provider networks, or statutory cost sharing.

#### *A1.II Occupational Health, Screening, and Sickness Provisions*

Dutch employers face extensive obligations in occupational health and sickness management. Under the Working Conditions Act (*Arbowet*), firms must contract a certified occupational health service or expert, maintain a written risk inventory and evaluation (*RI&E*) covering physical, chemical, ergonomic, and psychosocial risks, and implement a corresponding plan of action. Employers must also provide employees with access to a registered company doctor who supervises sickness absence and reintegration, and they must continue paying wages for up to two years during illness. When indicated by the *RI&E*, firms must offer a job-specific occupational health examination (*PAGO*), and they may also provide broader preventive medical examinations (*PMO*). Participation is voluntary except in a small set of legally defined high-risk or safety-critical jobs.

These obligations are reinforced by experience rating in disability-related

---

<sup>44</sup>Although workplaces have expanded mental health and well-being services in recent years (Shreekumar and Vautrey, 2022), in the setting of this study most mental healthcare is still provided through the formal basic health insurance system (Prudon, 2025).

<sup>45</sup>Vektis reports that 15.9% of insured persons in 2015 had no supplementary insurance, implying that about 84% held some form of supplementary insurance (Vektis, 2016, p. 23).

insurance. Employers pay an experience-rated contribution to the Return-to-Work Fund (Whk), which finances benefits under the Work and Income Act (WIA) and the Sickness Benefits Act (ZW). The premium depends on the firm's own inflow into sickness and disability benefits.<sup>46</sup> Dutch firms therefore face both legal duties and financial incentives to limit workplace-related health problems.

Compliance, however, is incomplete. Around 80% of firms report having a contract with an occupational health provider and roughly three quarters have a formal sickness-absence policy, but only about half maintain a written RI&E and fewer than 40% combine it with a concrete plan of action. When four core requirements are considered jointly—RI&E, an occupational health contract, in-company emergency response, and a designated prevention officer—only about one third of firms comply with all four ([Nederlandse Arbeidsinspectie, 2019](#)). Compliance is higher among large and high-risk firms than among small employers.

Preventive use of occupational health services is even narrower than formal coverage. Although most firms have access to a company doctor through an external provider, only about 7% report that employees used an open consultation hour in the previous year. Screening is similarly limited: PAGO/PMO is explicitly included in contracts for roughly one quarter of firms, but only about 4% report that an examination actually took place during the previous 12 months ([Nederlandse Arbeidsinspectie, 2019](#)). These programmes are concentrated in large firms and in sectors with well-identified exposures such as noise, chemicals, and heavy physical work. For most workers, day-to-day healthcare still runs through the ordinary GP-based system. Occupational doctors mainly provide work-related advice and reintegration support rather than medical treatment.

In short, the Dutch institutional setting combines substantial employer responsibility for prevention and sickness management with limited employer control over the content of ordinary medical care. That combination is useful for separating workplace health effects from insurance design.

---

<sup>46</sup>Studies find that linking premiums to firm-specific disability inflow reduces disability insurance (DI) entries among firms whose premiums actually adjust ([Koning, 2009](#)), and that removing experience rating raises DI inflow and lowers outflow for small employers ([De Groot and Koning, 2016](#)).

### *A1.III Workplace and Basic Healthcare Expenditure*

Cross-firm differences in basic healthcare expenditure can arise through two broad channels: differences in workers' ability to seek care and differences in the health consequences of the workplace itself.<sup>47</sup> The Dutch institutional setting is particularly informative because several features narrow the scope for the first channel. Out-of-pocket payments for basic care are capped by the annual deductible, more than 90% of adults choose the minimum deductible, and the marginal price of additional basic care is zero for high spenders once the deductible is exhausted. Existing evidence also finds limited income effects on healthcare expenditure (Miller et al., 2024; Cesarini et al., 2016). Time barriers also seem modest. Collective agreements and labour law provide paid time off for sickness and medical visits, employers must continue wage payments during illness for up to two years, and reintegration is formally supervised by occupational physicians. Consistent with this, self-reported unmet medical need in the Netherlands is among the lowest in the EU and shows little income gradient (e.g. OECD and European Observatory on Health Systems and Policies, 2019, 2021); lack of time is cited by less than 1% of respondents (Eurostat, 2023). Taken together, these features make the setting well suited to studying the persistent expenditure differences across firms.

Workplace health conditions provide a natural set of mechanisms through which firms can generate persistent expenditure differences. One set of mechanisms runs through safety and physical working conditions: injury risks and ergonomic strain vary across firms and sectors and respond to regulation and enforcement (Hamermesh, 1999; Lee and Taylor, 2019; Johnson, 2020; Johansson et al., 2023). Despite substantial progress, occupational injuries and work-related musculoskeletal disorders remain major health hazards (Bhattacharya, 2014; van Ours, 2019). A second set concerns psychosocial stress: high workloads, long hours, job insecurity, and poor management practices are linked to stress, depression, burnout, and related somatic conditions (Hummels et al., 2025; Blackburn et al., 2023; Jolivet and Postel-Vinay, 2025). Third, workplaces differ in environmental

---

<sup>47</sup>Appendix A4 formalises this point in a Grossman-style framework (Grossman, 1972b), where individuals choose health investments given preferences, income and time constraints, and a workplace-specific health-depreciation rate.

exposures such as noise, dust, chemicals, heat, and air pollution, which are established risk factors for respiratory, cardiovascular, renal, and cancer outcomes (e.g. Collaborators et al., 2018; Loomis et al., 2018; Tran et al., 2022; Gan et al., 2011). Finally, firms may shape daily health behaviours through food provision, exercise incentives, and workplace norms around smoking, alcohol use, and physical activity (Royer et al., 2015; Jones et al., 2019; Simonsen and Skipper, 2025).

## A2 Summary Statistics and Additional Results

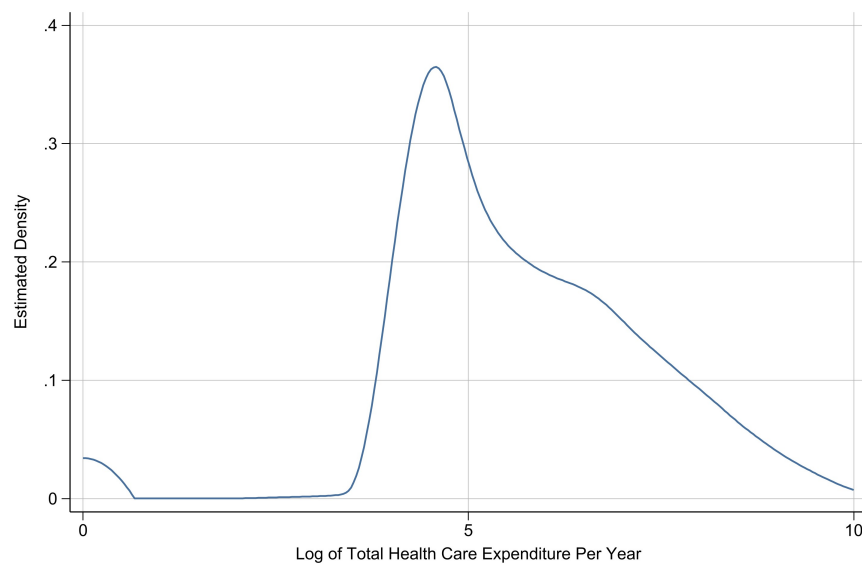


Figure A1: Log of total health expenditure per year for the connected set of employees (pooled male and female employees)

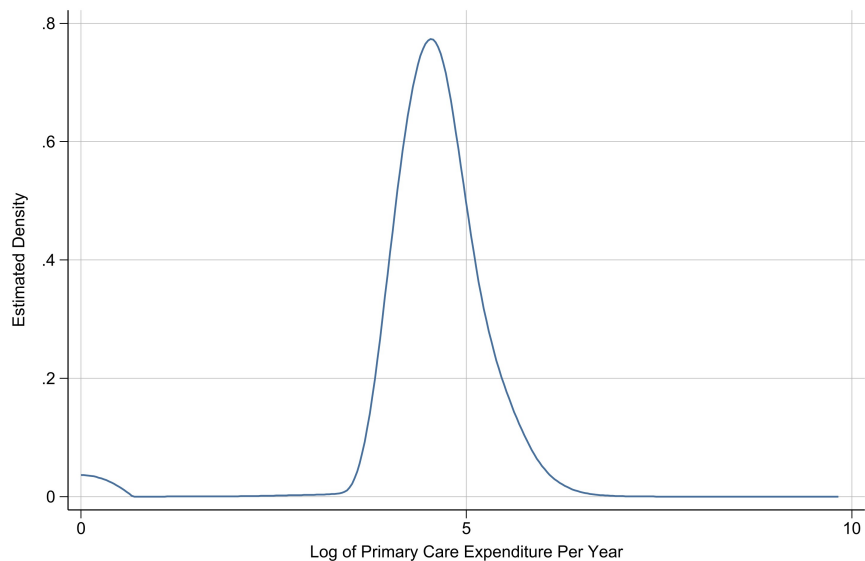


Figure A2: Log of primary health care expenditure per year for the connected set of employees (pooled male and female employees)

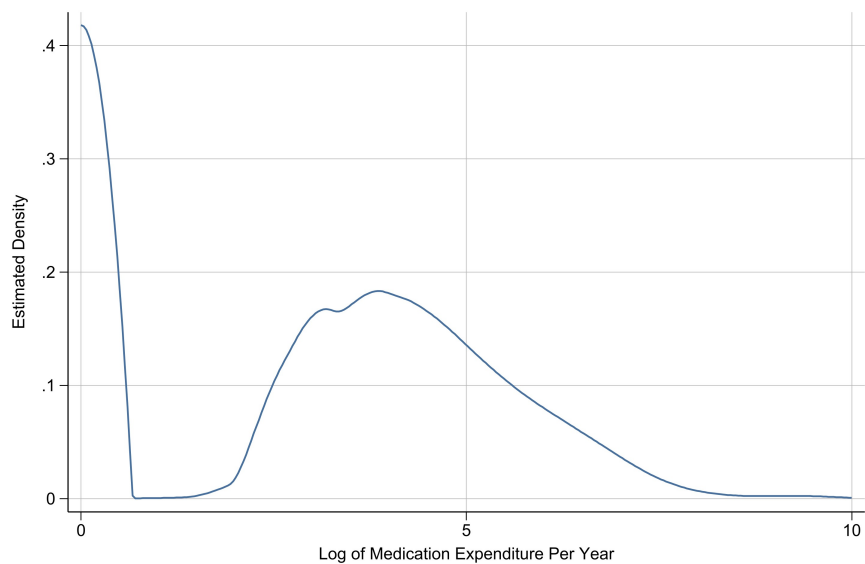


Figure A3: Log of medication costs per year for the connected set of employees (pooled male and female employees)

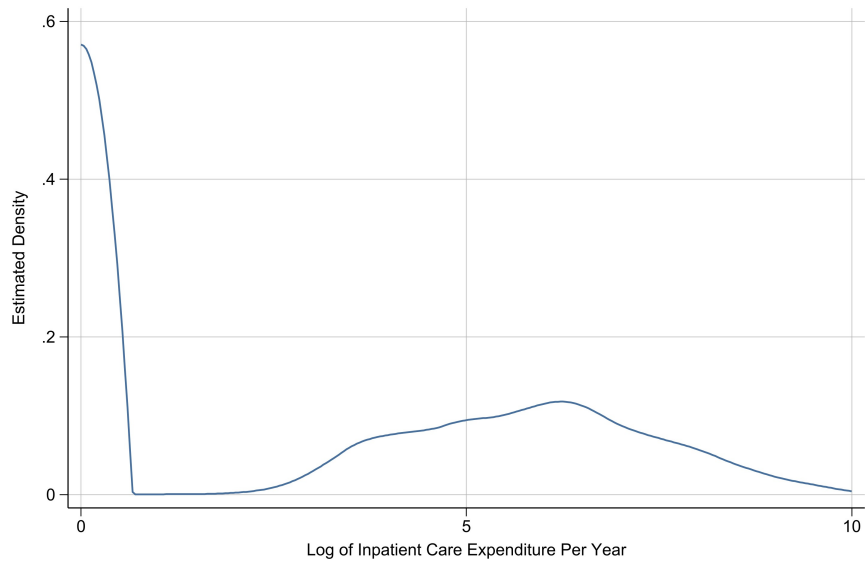


Figure A4: Log of inpatient care expenditure per year for the connected set of employees (pooled male and female employees)

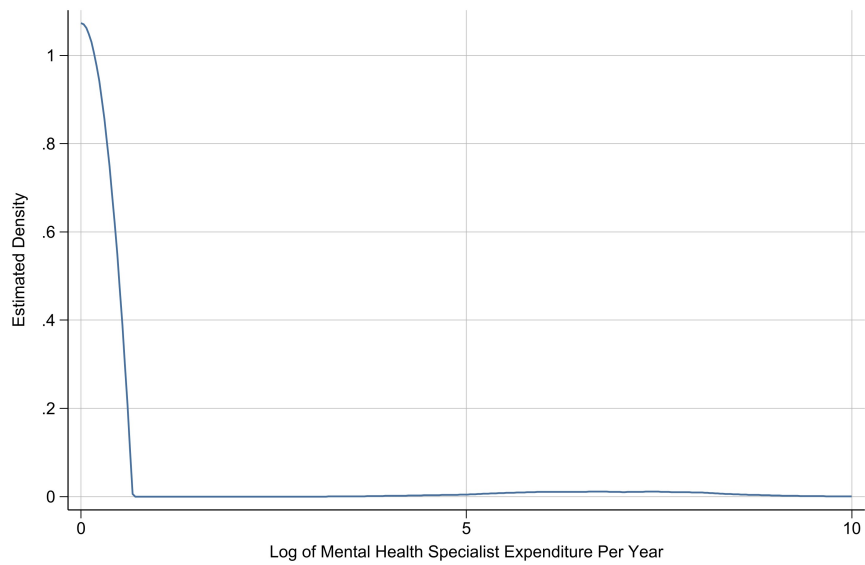




Figure A5: Log of mental health care expenditure (excluding primary care) per year for the connected set of employees (pooled male and female employees)




(a) Average Changes in Expenditure (binscatter)

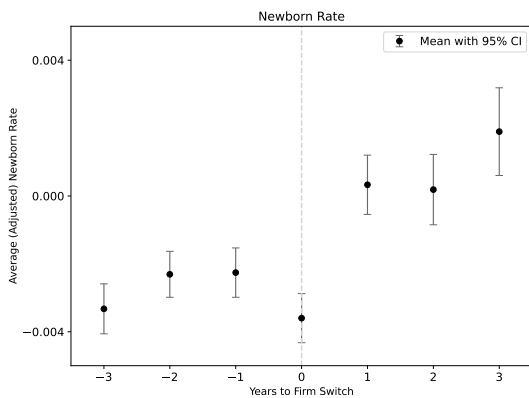


(b) Distribution of changes

Figure A6: Comparison of changes in workers' healthcare expenditure with changes in their colleagues' expenditure.



(a) Men



(b) Women

Note: The mean annual new birth rate is approximately 4% for both samples.

Figure A7: Average annual probability of having a new child around the time of a job switch (time = 0) for movers. Estimates are adjusted for age and education using a flexible age by education functional form using a linear regression.

Statistic	Plug-In Estimates (Biased)
Log Health Exp. Var.: $var(h_i)$	3.01 (Ref.)
Firm FE Var.: $var(\psi_j)$	0.33 (10.9%)
Employee FE Var.: $var(\alpha_i)$	1.94 (64.4%)
Firm and Employee FE's Cov.: $cov(\alpha_i, \psi_j)$	-0.04
No. Movers	1,232,499
No. Firms	132,899
No. of Observations	17,300,199

Note: This table displays the variance and covariance components from Equation (4), using the naive plug-in estimates of the parameters estimated with the OLS method from Equation (4).

Table A1: The variance decomposition results using the fixed effects estimated in the TWFE model (4): Male workers




Figure A8: Mean age-adjusted log healthcare expenditures of movers by coworker expenditure quartile.


Note: Average age-adjusted logarithmic healthcare expenditures for employees who changed firms between 2009 and 2016, remaining at the origin firm for at least two years and at the destination firm for at least two years. Firms are grouped into quartiles by the average coworker health expenditure within the firm.

	Low group mean $\hat{h}_{it}$	High group mean $\hat{h}_{it}$	Difference (low – high)	Implied % difference in spending <sup>a</sup>
Household wealth 2009 (bottom vs top quintile)	0.147	-0.111	0.258	29.4
Total income 2009 (bottom vs top quintile)	0.153	-0.133	0.286	33.1
Education (primary vs master/doctor)	0.257	-0.140	0.398	48.8

<sup>a</sup> The implied percentage difference is  $[\exp(\text{Difference}) - 1] \times 100$ .

Notes: Outcome is log annual basic-package healthcare spending residualised on age and calendar-year dummies,  $\hat{h}_{it}$ . Each cell reports the mean of  $\hat{h}_{it}$  for the indicated group, averaged over person-year observations for full-time male employees aged 25–65 in the main analysis sample (2009–2016). Household wealth and total personal income are measured in 2009 and divided into five quintiles; the corresponding rows compare the bottom and top quintiles of these distributions. Education is measured as highest completed level, and the row compares workers with only primary education (“Basisonderwijs”) to those with a master’s degree or doctorate.

Table A2: Gradients in age- and year-adjusted log healthcare spending by SES



(a) Household wealth quantile in 2009

(b) Total personal income quantile in 2009

Figure A9: Location of the analysis sample in the national wealth and income distributions

Note: Each panel plots the kernel density of the position of the full-time male main sample in the national distribution of the corresponding SES measure, expressed as percentiles (0–100) of the Dutch population in 2009. For every individual in the main analysis sample, I use the Statistics Netherlands reported rank for 2009. The sample consists of full-time male employees aged 25–65 in the 2009–2016 analysis period.




Figure A10: Distribution of estimated firm fixed effects

*Note:* Histogram of the demeaned estimated firm fixed effects  $\hat{\psi}_j$  from the AKM specification in Equation (4) for all (individual-year weighted) firms in the connected set used in the main analysis. The outcome is log annual basic-package healthcare spending,  $h_{it} = \log(c_{it} + 1)$ , and the model includes worker and firm fixed effects as well as age–education profiles and calendar-year dummies. The solid line shows the median.

Statistic	KSS Estimates
Log Wage Var.: $\text{var}(w_i)$	0.176 (Ref.)
Firm FE Var.: $\text{var}(\psi_j^w)$	0.017 (9.48 %)
Employee FE Var.: $\text{var}(\alpha_i^w)$	0.271 (154.12 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i^w, \psi_j^w)$	–0.001
Correlation $\rho(\alpha_i^w, \psi_j^w)$	–0.020
Explained share: workers & firms	1.621
No. Movers	1,225,799
No. Firms	132,820
No. of Observations	16,976,071
Average Log Wage: $\text{mean}(w_i)$	3.141

*Notes:* Components estimated via leave-one-match-out KSS correction (Kline et al., 2020) using random projections (50 simulations). The outcome is log hourly wage. Hourly wages are constructed from gross annual contract earnings, contractual full-time equivalent (FTE), and contract length, assuming that a full-time job corresponds to 52 weeks of 40 hours per week, as in the data construction in Section III. The sample consists of full-time male employees in the leave-one-out connected set of firms between 2009 and 2016. Shares in parentheses are relative to  $\text{var}(w_i)$ .

Table A3: Variance Decomposition of Log Hourly Wages

Panel A: KSS Variance Components	
Statistic	KSS Estimates
Log Health Exp. Var.: $\text{var}(h_i)$	2.617 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.027 (1.03 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.515 (57.89 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	−0.004
Correlation $\rho(\alpha_i, \psi_j)$	−0.018
Explained share: workers & firms	0.586
No. Movers	1,225,799
No. Firms	132,820
No. of Observations	16,976,071
Average Health Exp.: $\text{mean}(h_i)$	5.608
Panel B: Regression of Firm Effects with Pay Control on Baseline Firm Effects	
Dependent variable	$\hat{\psi}_j^{\text{HC,pay}}$
Coefficient on baseline firm FE (worker level) $\hat{\psi}_j$	1.000
No. of workers	16,976,071

*Notes:* Panel A reports components estimated via leave-one-match-out KSS correction (Kline et al., 2020) using random projections (50 simulations). The outcome is log annual basic-package healthcare expenditure per calendar year. The regression includes worker and firm fixed effects, age-by-education profiles, calendar-year dummies, and an additional control for workers’ annual gross contract payment (yearly labour income). Shares in parentheses are relative to  $\text{var}(h_i)$ . Panel B reports an OLS regression of the firm fixed effects from the healthcare model with the contract-payment control,  $\hat{\psi}_j^{\text{HC,pay}}$ , on the baseline healthcare firm fixed effects without that control,  $\hat{\psi}_j$ . The slope is close to one and the intercept is close to zero, indicating that adding the pay control leaves the pattern of estimated firm health effects almost unchanged.

Table A4: Variance Decomposition of Log Healthcare Expenditure (Controlling for Workers’ Annual Contract Payments)

	High baseline costs ( $\hat{\alpha}_i$ above median)	Low baseline costs ( $\hat{\alpha}_i$ below median)
<i>Panel A. KSS variance components</i>		
Log Health Exp. Var.: $\text{var}(h_i)$	2.101 (Ref.)	1.745 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.022 (1.03 %)	0.035 (1.98 %)
Employee FE Var.: $\text{var}(\alpha_i)$	0.550 (26.19 %)	0.563 (32.25 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	0.032	0.032
Correlation $\rho(\alpha_i, \psi_j)$	0.297	0.228
Explained share: workers & firms	0.303	0.379
No. Movers	551,201	606,960
No. Firms	80,230	89,384
No. of Observations	8,035,638	8,110,825
Mean log Health Exp.: $\text{mean}(h_i)$	6.447	4.775
<i>Panel B. Descriptives and cross-sample link</i>		
Mean annual basic-package costs (€) <sup>a</sup>	2,011	286.52
No. obs. used for mean costs <sup>a</sup>	8,488,038	8,488,038
Bias-corrected slope (worker level), firm FE (low baseline) on firm FE (high baseline) <sup>b</sup>		0.602
No. worker–year obs. in slope regression <sup>b</sup>		7,635,508

<sup>a</sup> Mean annual basic-package costs are computed directly from raw annual spending  $c_{it}$  (in euros) in the two samples, split by whether the estimated individual healthcare fixed effect  $\hat{\alpha}_i$  from the baseline AKM lies above or below its median (median split defined over all connected-set worker–year observations).

<sup>b</sup> Slope from an OLS regression of the firm fixed effects estimated in the low-baseline-costs subsample on the firm fixed effects estimated in the high-baseline-costs subsample, using worker–year observations in the overlap where both firm effects are defined. The reported slope is bias-corrected by rescaling the attenuated OLS coefficient using the ratio of KSS-corrected to plug-in firm-effect variances following [Kline et al. \(2020\)](#). Because KSS is implemented on the leave-one-out connected set *within each subsample*, the two subsamples need not form an exact partition of the baseline connected set, and some firms (and their worker–year observations) appear in only one subsample and therefore drop out of the overlap regression.

*Notes:* Panel A components are estimated via leave-one-match-out KSS correction ([Kline et al., 2020](#)) using random projections (50 simulations). The outcome is log annual basic-package healthcare expenditure per calendar year,  $h_{it} = \log(c_{it} + 1)$ . Shares in parentheses are relative to  $\text{var}(h_i)$  in the respective subsample.

Table A5: Variance Decomposition of Log Healthcare Expenditure by Baseline Individual Health-Cost Type

	High-wage workers <i>(above-median hourly wage)</i>	Low-wage workers <i>(below-median hourly wage)</i>
<i>Panel A. KSS variance components</i>		
Log Health Exp. Var.: $\text{var}(h_i)$	2.543 (Ref.)	2.706 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.022 (0.88 %)	0.039 (1.43 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.198 (47.08 %)	1.630 (60.22 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	−0.006	−0.005
Explained share: workers & firms	0.475	0.613
No. Movers	547,051	631,701
No. Firms	58,556	103,831
No. of Observations	8,061,912	8,278,053
Mean log Health Exp.: $\text{mean}(h_i)$	5.567	5.653
<i>Panel B. Descriptives and cross-sample link</i>		
Mean hourly wage (€) <sup>a</sup>	33.13	16.39
Mean annual contract payment (€) <sup>a</sup>	67,783	33,306
No. obs. used for wage descriptives <sup>a</sup>	8,488,041	8,488,030
Bias-corrected slope (worker level), firm FE (low wage) on firm FE (high wage) <sup>b</sup>		0.622
No. worker–year obs. in slope regression <sup>b</sup>		6,644,986

<sup>a</sup> Wage groups are defined using an estimated hourly wage (in euros) computed from annual contract payment and work intensity (as in Section III). The median hourly wage in the connected-set sample is 20.900,00; “high wage” corresponds to observations above this threshold. Panel B means are computed directly from the raw wage variables in the connected-set person–year sample and are descriptive.

<sup>b</sup> Slope from an OLS regression of the firm fixed effects estimated in the low-wage subsample on the firm fixed effects estimated in the high-wage subsample, using worker–year observations in the overlap where both firm effects are defined. The reported slope is bias-corrected by rescaling the attenuated OLS coefficient using the ratio of KSS-corrected to plug-in firm-effect variances following [Kline et al. \(2020\)](#). Because KSS is implemented on the leave-one-out connected set *within each wage subsample*, the two subsamples need not form an exact partition of the baseline connected set, and some firms (and their worker–year observations) appear in only one subsample and therefore drop out of the overlap regression.

*Notes:* Panel A components are estimated via leave-one-match-out KSS correction ([Kline et al., 2020](#)) using random projections (50 simulations). The outcome is log annual basic-package healthcare expenditure per calendar year,  $h_{it} = \log(c_{it} + 1)$ . Shares in parentheses are relative to  $\text{var}(h_i)$  in the respective subsample.

Table A6: Variance Decomposition of Log Healthcare Expenditure by Wage Group

	Young workers (age ≤ 45)	Old workers (age > 45)
<i>Panel A. KSS variance components</i>		
Log Health Exp. Var.: $\text{var}(h_i)$	2.390 (Ref.)	2.680 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.029 (1.23 %)	0.014 (0.54 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.369 (57.28 %)	1.649 (61.52 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	0.005	−0.001
Explained share: workers & firms	0.590	0.620
No. Movers	742,335	409,353
No. Firms	108,472	49,383
No. of Observations	8,653,958	7,154,491
Mean log Health Exp.: $\text{mean}(h_i)$	5.317	5.944
<i>Panel B. Cross-sample link</i>		
Bias-corrected slope (worker level), firm FE (young) on firm FE (old) <sup>a</sup>		0.846
No. worker–year obs. in slope regression <sup>a</sup>		7,236,057

<sup>a</sup> Slope from an OLS regression of the firm fixed effects estimated in the young subsample on the firm fixed effects estimated in the old subsample, using worker–year observations in the overlap where both firm effects are defined. The reported slope is bias-corrected by rescaling the attenuated OLS coefficient using the reliability ratio for the right-hand-side firm effect,  $\text{Var}_{\text{KSS}}(\psi)/\text{Var}_{\text{PI}}(\hat{\psi})$ , following [Kline et al. \(2020\)](#) (equivalently,  $\hat{\beta}_{\text{BC}} = \hat{\beta}_{\text{OLS}} \times \text{Var}_{\text{PI}}(\hat{\psi}_{\text{old}})/\text{Var}_{\text{KSS}}(\psi_{\text{old}})$ ). Because KSS is implemented on the leave-one-out connected set *within each age subsample*, the two subsamples need not form an exact partition of the baseline connected set; some firms appear in only one subsample and drop out of the overlap regression.

*Notes:* Panel A components are estimated via leave-one-match-out KSS correction ([Kline et al., 2020](#)) using random projections (50 simulations). The outcome is log annual basic-package healthcare expenditure per calendar year,  $h_{it} = \log(c_{it} + 1)$ . Shares in parentheses are relative to  $\text{var}(h_i)$  in the respective subsample.

Table A7: Variance Decomposition of Log Healthcare Expenditure by Age Group

	White-collar workers (education proxy)	Blue-collar workers (education proxy)
<i>Panel A. KSS variance components</i>		
Log Health Exp. Var.: $\text{var}(h_i)$	2.142 (Ref.)	2.286 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.020 (0.92 %)	0.032 (1.38 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.071 (50.01 %)	1.143 (49.99 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	–0.005	–0.012
Explained share: workers & firms	0.504	0.503
No. Movers	319,969	158,616
No. Firms	49,485	40,897
No. of Observations	3,719,229	1,903,635
Mean log Health Exp.: $\text{mean}(h_i)$	5.443	5.638
<i>Panel B. Cross-sample link</i>		
Bias-corrected slope (worker level), firm FE (white collar) on firm FE (blue collar) <sup>b</sup>		1.063
No. worker–year obs. in slope regression <sup>b</sup>		1,598,774

<sup>a</sup> Blue- and white-collar status is proxied using education (see Appendix A7 for the construction of the collar proxy used in the paper).

<sup>b</sup> Slope from an OLS regression of the firm fixed effects estimated in the white-collar subsample on the firm fixed effects estimated in the blue-collar subsample, using worker–year observations in the overlap where both firm effects are defined. The reported slope is bias-corrected by rescaling the attenuated OLS coefficient using the right-hand-side reliability ratio  $\text{Var}_{KSS}(\psi)/\text{Var}_{PI}(\hat{\psi})$  following Kline et al. (2020), i.e.  $\hat{\beta}_{BC} = \hat{\beta}_{OLS} \times \text{Var}_{PI}(\hat{\psi}_{blue})/\text{Var}_{KSS}(\psi_{blue})$ .



Notes: Panel A components are estimated via leave-one-match-out KSS correction (Kline et al., 2020) using random projections (50 simulations). The outcome is log annual basic-package healthcare expenditure per calendar year,  $h_{it} = \log(c_{it} + 1)$ . Shares in parentheses are relative to  $\text{var}(h_i)$  in the respective subsample.

Table A8: Variance Decomposition of Log Healthcare Expenditure by Collar Type (proxied by education)

Statistic	Primary care (GP)	Medication	Hospital (Inpatient + Specialist)	Mental health
Log Health Exp. Var.: $\text{var}(h_i)$	0.655	5.689	9.811	1.632
Firm FE Var.: $\text{var}(\psi_j)$	0.018	0.010	0.024	0.003
Employee FE Var.: $\text{var}(\alpha_i)$	0.521	3.209	2.911	0.371
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	–0.003	–0.001	–0.004	–0.002
Correlation $\rho(\alpha_i, \psi_j)$	–0.027	–0.006	–0.017	–0.061
Explained share: worker & firm	0.815	0.565	0.298	0.227
No. Movers	1,225,799	1,225,799	1,225,799	1,225,799
No. Firms	132,820	132,820	132,820	132,820
No. of Observations	16,976,071	16,976,071	16,976,071	16,976,071
Average Health Exp.: $\text{mean}(h_i)$	4.521	2.637	2.732	0.240

Notes: All columns report leave-one-match-out KSS estimates (Kline et al., 2020). Outcomes are the log of subcategory-specific healthcare expenditure: primary care (GP), medication, hospital (inpatient + specialist), and specialist mental health. The sample is the full connected set.

Table A9: Variance Decomposition of Log Healthcare Expenditure (KSS) by Cost Category



(a) Average change vs. destination exposure gap

(b) Event study around closure and reemployment

Figure A11: Workers' healthcare spending around firm closures and reemployment

*Notes.* Sample of full-time male employees aged 25–65 whose employer exits during the panel and who are observed in consecutive years from 2009 and around the closure. Healthcare outcomes are log total basic-package expenditures, residualised on age and calendar-year dummies. The closure year is the last year in which the origin employer identifier is observed. In panel (a), each point plots the average change in individual annual residualised spending from the pre-closure year to the post-reemployment year against the destination exposure gap, defined as the difference in leave-*i*-out mean residualised spending between the destination and origin employers. The line shows a non-parametric fit, where the intercept captures the common post-displacement increase and the slope captures the destination-firm component. In panel (b), the markers show coefficients from an event-study specification that interacts event-time indicators with the destination exposure gap, with event time measured relative to the closure year and  $t = -2$  as the reference year.

<b>Closure Sample</b>	
<i>(Worker-year observations)</i>	
<i>N = 1,550,889</i>	
<b>Variable</b>	<b>Mean (SD)</b>
Age	46.03 (9.43)
Dutch background (dummy)	0.85 (0.36)
First-generation migrant (dummy)	0.08 (0.27)
Second-generation migrant (dummy)	0.07 (0.25)
Education: Elementary (dummy)	0.02 (0.13)
Education: Lower sec./MBO1 (dummy)	0.05 (0.21)
Education: Upper sec./MBO (dummy)	0.19 (0.40)
Education: Bachelor (HBO/WO) (dummy)	0.18 (0.38)
Education: Master/Doctorate (dummy)	0.11 (0.31)
Education: Missing (dummy)	0.46 (0.50)
Total # children	1.54 (1.20)
New child (dummy)	0.04 (0.19)
Gross payment of contract (€)	55,101 (40,796)
Mover (dummy)	1.00 (0.00)
Full-time factor (FTE)	1.00 (0.01)
Estimated hourly wage (€) <sup>a</sup>	26.84 (19.78)
Business size (employees)	1,012 (815.15)
Total health care costs (€)	1,159 (4,574)
Mental health (specialist) (€)	75.13 (1,188)
Hospital costs (€)	700.58 (3,867)
GP costs (€)	110.45 (65.88)
<b>Number of Observations</b>	<b>1,550,889</b>

<sup>a</sup> Hourly wage computed from gross contract payment assuming a full FTE works 52 weeks × 40 hours per week.

*Notes:* Firm-closure sample (worker-year observations). Variables are defined as in Table 1. All observations are movers by construction (Mover dummy = 1). Business size is observed for 1,544,519 worker-year observations (all other variables are observed for 1,550,889).

Table A10: Summary Statistics: Firm-Closure Sample

	(1)	(2)	(3)
	<i>Outcome measured pre-move</i>		
Destination firm health-cost effect (leave-out)	0.0068 (0.0092)	-0.3305*** (0.0132)	-0.3879*** (0.0201)
Mean of dep. var.	0.2054	0.2054	0.1960
<i>N</i>	209,235	209,235	206,945
<i>R</i> <sup>2</sup>	0.084	0.089	0.146
Baseline controls & pre-move postal-code FEs	Yes	Yes	Yes
Origin firm-effect level (control)	No	Yes	No
Origin-firm fixed effects	No	No	Yes

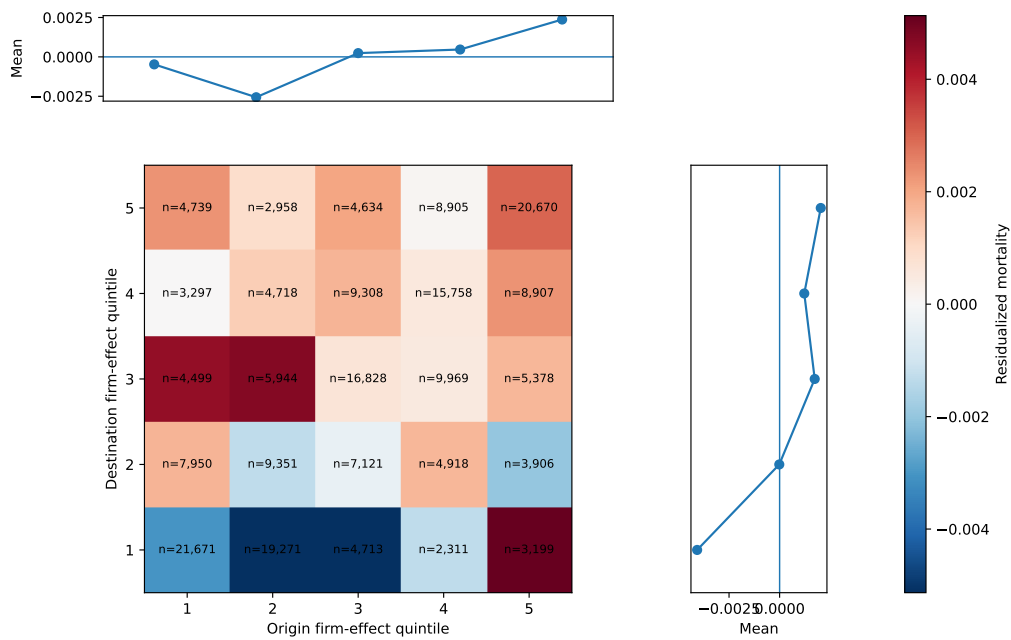
*Notes:* Sample of displaced male workers following firm closures, restricting destination firms to those above the 25th percentile of firm size (at least 75 employees). The outcome is average pre-move basic-package healthcare spending. Covariates are the same as in Table 5. Column (2) additionally controls for the origin firm-effect level; column (3) includes origin-firm fixed effects. Standard errors in parentheses. Significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Table A11: Firm Closures: Pre-Move Healthcare Spending

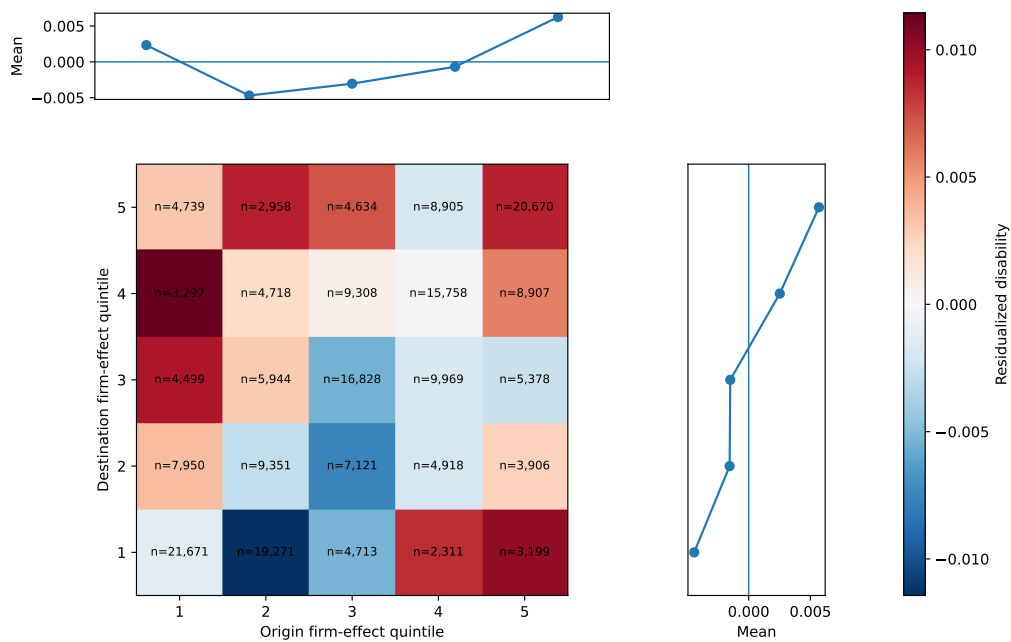
	(1)	(2)	(3)
<i>Outcome measured in 2019</i>			
<b>Panel A: Mortality (Accidental/Injuries)</b>			
Destination firm health-cost effect (leave-out)	0.000021 (0.000128)	0.000171 (0.000184)	-0.000170 (0.000282)
Mean of dep. var.		0.00034	
N	209,235	209,235	206,945
R <sup>2</sup>	0.020	0.020	0.072
<b>Panel B: Mortality (Cancers)</b>			
Destination firm health-cost effect (leave-out)	0.001772*** (0.000570)	0.002752*** (0.000820)	0.002729** (0.001267)
Mean of dep. var.		0.00687	
N	209,235	209,235	206,945
R <sup>2</sup>	0.037	0.037	0.069
<b>Panel C: Mortality (Cardiometabolic)</b>			
Destination firm health-cost effect (leave-out)	0.000341 (0.000366)	0.000774 (0.000527)	0.000570 (0.000815)
Mean of dep. var.		0.00278	
N	209,235	209,235	206,945
R <sup>2</sup>	0.024	0.024	0.057
<b>Panel D: Mortality (Chronic Respiratory)</b>			
Destination firm health-cost effect (leave-out)	0.000086 (0.000084)	0.000131 (0.000121)	0.000091 (0.000189)
Mean of dep. var.		0.00015	
N	209,235	209,235	206,945
R <sup>2</sup>	0.028	0.028	0.046
<b>Panel E: Mortality (Suicide/Substance)</b>			
Destination firm health-cost effect (leave-out)	0.000050 (0.000217)	-0.000458 (0.000312)	-0.000906* (0.000477)
Mean of dep. var.		0.00096	
N	209,235	209,235	206,945
R <sup>2</sup>	0.019	0.019	0.061
<b>Panel F: Mortality (Other)</b>			
Destination firm health-cost effect (leave-out)	0.000990*** (0.000313)	0.001377*** (0.000451)	0.001039 (0.000698)
Mean of dep. var.		0.00203	
N	209,235	209,235	206,945
R <sup>2</sup>	0.019	0.019	0.052
Baseline controls & pre-move postal-code FEs	Yes	Yes	Yes
Origin firm-effect level (control)	No	Yes	No
Origin-firm fixed effects	No	No	Yes

*Notes:* Sample of displaced male workers following firm closures, restricting destination firms to those above the 25th percentile of firm size (at least 75 employees). The regressor is the destination firm's leave-out health-cost effect. Column (2) additionally controls for the origin firm-effect level; column (3) includes origin-firm fixed effects. Baseline controls include demographics and pre-move postal-code fixed effects (as in Table 5). Cause-specific outcomes are constructed from the underlying cause-of-death ICD-10 codes recorded in the mortality register. Standard errors in parentheses. Significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Table A12: Firm Closures: Destination Firm Type and Cause-Specific Mortality in 2019



(a) Mortality in 2019



(b) Disability in 2019

Figure A12: Long-run outcomes by origin×destination firm-effect quintiles after firm closures

Note: Sample and residualization are the same as in Figure 5. The origin and destination firm health-cost effects are binned into quintiles (1 = lowest, 5 = highest). Cells show binned means of residualized outcomes, and annotations report the number of observations in each origin×destination cell.

Statistic	KSS Estimates
Log Health Exp. Var.: $\text{var}(h_i)$	2.609 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.013 (0.49 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.402 (53.74 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	0.006
Correlation $\rho(\alpha_i, \psi_j)$	0.045
Explained share: workers & firms	0.547
No. Movers	1,113,768
No. Firms	117,247
No. of Observations	16,737,318
Average Health Exp.: $\text{mean}(h_i)$	5.611

*Notes:* Components estimated via leave-one-match-out KSS correction (Kline et al., 2020) using random projections (50 simulations). Outcome is log healthcare expenditure per calendar year. The sample is restricted to full-time employees in firms with a non-missing sector code in the business registry, 2009–2016. Relative to the full connected sample in Table 3, the variance of firm effects is smaller in this restricted set, so firm effects account for a slightly smaller share of  $\text{var}(h_i)$ . Shares in parentheses are relative to  $\text{var}(h_i)$ .

Table A13: Variance Decomposition of Log Healthcare Expenditure: Registry-Sector Sample

Statistic	PC4 FE (Residence)	PC4 FE + Workplace Access
Log Health Exp. Var.: $\text{var}(h_i)$	2.393 (Ref.)	2.386 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.025 (1.03 %)	0.025 (1.03 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.360 (56.85 %)	1.368 (57.32 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	−0.006	−0.007
Correlation $\rho(\alpha_i, \psi_j)$	−0.031	−0.037
Explained share: worker, firm & edu–age	0.574	0.578
No. Movers	1,212,637	1,173,299
No. Firms	131,992	129,198
No. of Observations	16,806,677	16,263,162
Average Health Exp.: $\text{mean}(h_i)$	5.652	5.653

*Notes:* Components estimated via leave-one-match-out KSS correction (Kline et al., 2020) using random projections (50 simulations). Outcome is log healthcare expenditure per calendar year. The left column includes four-digit postcode (PC4) fixed effects for workers’ residential location; the right column additionally controls for PC4-level workplace-area healthcare access (number of GPs, distance to nearest pharmacy, number of nearby hospitals). Shares in parentheses are relative to  $\text{var}(h_i)$ .

Table A14: Variance Decomposition of Log Healthcare Expenditure (KSS, with Residential PC4 FEs and Workplace Access Controls)

Statistic	Plug-in (Full Sample)	Plug-in ( $\geq 50$ -Move Firms)	KSS ( $\geq 50$ -Move Firms)
Log Health Exp. Var.: $\text{var}(h_i)$	2.617 (Ref.)	2.675 (Ref.)	2.675 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.113 (4.33 %)	0.049 (1.84 %)	0.030 (1.13 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.839 (70.28 %)	1.790 (66.91 %)	1.520 (56.82 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	-0.077	-0.026	-0.008
Correlation $\rho(\alpha_i, \psi_j)$	-0.169	-0.086	-0.039
Explained share: worker, firm & edu–age	0.687	0.668	0.573
No. Movers	1,225,799	723,530	723,530
No. Firms	132,820	10,643	10,643
No. of Observations	16,976,071	11,295,814	11,295,814
Average Health Exp.: $\text{mean}(h_i)$	5.608	5.631	5.631

*Notes:* Plug-in columns report naive AKM quadratic components. The KSS column applies the leave-one-match-out bias correction of [Kline et al. \(2020\)](#). Outcome is log healthcare expenditure per calendar year. The  $\geq 50$ -move sample restricts to firms with at least 50 observed worker moves. Percent shares in parentheses are relative to  $\text{var}(h_i)$ .

Table A15: Variance Decomposition of Log Healthcare Expenditure: Plug-in vs. KSS, Full Sample vs.  $\geq 50$ -Move Firms

Statistic	1-year Ahead	2-year Ahead	3-year Ahead	4-year Ahead	5-year MA
Log Health Exp. Var.: $\text{var}(h_i)$	2.645 (Ref.)	2.723 (Ref.)	2.885 (Ref.)	3.044 (Ref.)	1.763 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.023	0.009	0.004	0.004	0.009
Employee FE Var.: $\text{var}(\alpha_i)$	1.490	1.544	1.662	1.723	1.610
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	-0.003	-0.001	-0.001	-0.003	0.000
Correlation $\rho(\alpha_i, \psi_j)$	-0.018	-0.005	-0.007	-0.039	-0.003
Explained share: worker, firm & edu–age	0.569	0.570	0.577	0.565	0.918
No. Movers	1,225,799	1,225,799	1,225,799	1,225,799	1,225,799
No. Firms	132,820	132,820	132,820	132,820	132,820
No. of Observations	16,976,071	16,976,071	16,976,071	16,976,071	16,976,071
Average Health Exp.: $\text{mean}(h_i)$	5.669	5.706	5.722	5.739	5.689

*Notes:* All columns report leave-one-match-out KSS estimates ([Kline et al., 2020](#)). Outcomes are the log of healthcare expenditure measured at different horizons relative to the baseline year: one-through four-year ahead, and a five-year moving average (current plus next four years). The sample is the full connected set.

Table A16: Variance Decomposition of Log Healthcare Expenditure (KSS) Across Horizons

### A3 AKM and the Mover Design

In the [Abowd et al. \(1999\)](#) (AKM) framework, outcomes are decomposed into individual and location (e.g., worker and firm) components. Using this model, one reports how much places matter—either as the variance share of location effects or, in a log-linear specification, as the expected change in outcomes when individuals move to locations whose estimated effects are one standard deviation higher ([Kline, 2024](#)).

The mover-design event study builds on the same underlying structure but adopts a different approach. Instead of estimating variance components directly, it exploits the correlation between changes in an individual’s outcome at the time of a move and the corresponding changes in the average outcomes of coworkers from the origin to the destination. This approach aims to capture how locations contribute to the observed cross-location dispersion in outcomes ([Finkelstein et al., 2016](#)).

To the best of my knowledge, despite extensive theoretical and empirical work on AKM models (e.g., see [Card et al. \(2018\)](#); [Kline et al. \(2020\)](#); [Andrews et al. \(2008\)](#)), the formal link between the mover-design estimand and the AKM *quadratic components* (i.e., variance components from the AKM decomposition) is not known. In applied work, mover-design coefficients are often interpreted as a variation share (e.g., see [Finkelstein et al. \(2021\)](#); [Godøy and Huitfeldt \(2020\)](#); [Luan et al. \(2025\)](#)), and sometimes used as a benchmark alongside AKM-based quadratic components ([Lyubich, 2025](#)).

I show that, when sorting patterns are homogeneous across individuals and firms and remain stable over time, and when locations are sufficiently large, the mover coefficient converges to a ratio of standard deviations: the standard deviation of location fixed effects divided by the cross-location standard deviation of average outcomes. Reading the mover coefficient itself as a “variance share” therefore overstates the role of locations; the relevant variance share is the square of this coefficient.

## Setup

In this appendix I focus on the *estimand* in the mover-design setting and its link to the AKM quadratic components, abstracting from sampling properties of estimators. Consider a balanced panel observed in two periods  $t \in \{1, 2\}$ , with individuals  $i$  located in  $J$  locations  $j \in \mathcal{J}$  each period. Let  $j(i, t)$  denote  $i$ 's location at  $t$ ,  $o(i) := j(i, 1)$  the origin, and  $d(i) := j(i, 2)$  the destination. Outcomes follow the AKM structure:

$$Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}, \quad (5)$$

where  $\alpha_i$  is a time-invariant individual effect,  $\psi_j$  a location effect, and  $\varepsilon_{it}$  an idiosyncratic error with  $\mathbb{E}[\varepsilon_{it} \mid \alpha_i, \psi_{j(i,t)}] = 0$ .

Define movers  $M := \{i : d(i) \neq o(i)\}$  and stayers  $S := \{i : d(i) = o(i)\}$ , and let  $w$  denote the (limit) share of movers in the population (identical across periods in a closed balanced panel). I allow the distributions of  $(\alpha_i, \psi_{j(i,t)})$  and their correlation to differ between  $M$  and  $S$ .

The standard first-difference mover regression is

$$\Delta Y_i = \theta \delta_i^{(-i)} + u_i, \quad (6)$$

with

$$\begin{aligned} \Delta Y_i &:= Y_{i2} - Y_{i1} = \underbrace{\psi_{d(i)} - \psi_{o(i)}}_{\Delta \psi_i} + \underbrace{\varepsilon_{i2} - \varepsilon_{i1}}_{\Delta \varepsilon_i}, \\ \delta_i^{(-i)} &:= \bar{Y}_{d(i),2}^{(-i)} - \bar{Y}_{o(i),1}^{(-i)}, \end{aligned}$$

where  $\bar{Y}_{j,t}^{(-i)}$  is the leave-one-out mean outcome in location  $j$  at time  $t$ . When (6) is run on movers,

$$\theta = \frac{\text{Cov}_M(\delta_i^{(-i)}, \Delta Y_i)}{\text{Var}_M(\delta_i^{(-i)})},$$

with  $\text{Cov}_M$  and  $\text{Var}_M$  computed in the mover subsample, while keeping the stayers in the mean calculations<sup>48</sup>.

<sup>48</sup>When the OLS is estimated solely on movers, the derivations in this appendix remain essentially unchanged with the constraint that the share of movers equals one.

**Assumptions.** For tractability I maintain:

*A1 (AKM outcome model).* Equation (5) holds and  $\varepsilon_{it}$  is mean independent of  $(\alpha_i, \psi_{j(i,t)})$ . This is the canonical two-way additive structure underlying mover-designs to unpack between-location variation into *location effects* and *sorting* (Finkelstein et al., 2016). Under the structure, conditional on  $(\alpha_i, \psi_{j(i,t)})$ , the observations  $(i, t)$  are i.i.d. across  $i$  and across  $t$ .

*A2 (Finite locations, no sorting on size).* Each location  $j$  has size  $n_{j,t} \equiv n_j$  (not necessarily equal across  $j$ ). Define  $c_j := 1/n_j$  and the mover-average  $c := \mathbb{E}_{i \in M, t} [c_{j(i,t)}]$ . Location sizes are independent of  $(\psi_j)$  and of move decisions. I use the standard finite-population approximation  $\frac{1}{n_j-1} \approx \frac{1}{n_j} = c_j$ . This assumption implies even in a very large economy, location sizes are not necessarily infinite, later I discuss the more restrictive case where the number of observations within locations are also very large.

The leave-one-out mean therefore mixes *signal* (systematic composition across locations) with *sampling noise* of order  $1/n_j$ . Ruling out sorting on size is convenient as it isolates the role of  $c$  as a pure precision term. Additionally, I assume, mover share  $w$  is constant across firms and across time (i.e.,  $w_{j,t} = w$  for all  $(j, t)$ )<sup>49</sup>.

*A3 (Affine relationship between individual and location effects, and stable sorting for movers across periods).* Here, *affine* means linear up to a constant. I assume that the assignment of individuals to locations is stable along two dimensions. First, sorting is stable *across time*: although movers change locations between periods, the strength of sorting—measured by the relationship between individual and location effects—remains constant before and after the move. This is not a strong restriction in near-equilibrium or short-horizon settings such as labor markets or location–health environments, where the mapping between types of individuals and types of locations typically evolves slowly and structural parameters are approximately constant.

Second, sorting is stable *across the distribution of location effects*. Within each group  $g \in \{M, S\}$  (movers, stayers), I impose an exact affine relationship between

---

<sup>49</sup>This is not essential in asymptotics, but it helps with tractability in small samples. The large sample results are robust to relaxing this assumption, with appropriate re-definition of  $\beta$  in (10).

individual and location effects:

$$\alpha_i = a_g + s_g \psi_{j(i,t)} + e_{i,t}, \quad s_g = \frac{\text{Cov}_g(\alpha, \psi)}{\text{Var}_g(\psi)} = \frac{\rho_g \sigma_{\alpha,g}}{\sigma_{\psi,g}}, \quad \mathbb{E}[e_{i,t} \mid \psi_{j(i,t)}, g] = 0, \quad (7)$$

with coefficients  $(a_g, s_g)$  constant over the support of  $\psi$  and across periods  $t \in \{1, 2\}$ . Movers and stayers may differ in the level and slope of this relationship,  $(a_M, s_M) \neq (a_S, s_S)$ , but for each group the sorting intensity does not vary with the level of location effects. Additionally, I assume any idiosyncratic assignment or transitory component is conditionally independent over time. So I rule out any correlation in the assignment residual once  $(\alpha_i, \psi_{j(i,t)})$  are conditioned on. Following (7), one can write  $\psi_{j(i,t)} = \tilde{\alpha}_g + \tilde{s}_g \alpha_i + \tilde{e}_{i,t}$  with  $\tilde{e}_{i,t}$  uncorrelated across time. The affine structure is an additional and relatively strong assumption relative to the canonical AKM framework, which remains agnostic about the functional form of sorting between individuals and locations. However, it is indispensable for interpretability: under (7), the correlation between coworkers' individual effects and the dependence between origin and destination locations can be expressed in closed form, allowing the mover-design estimand to be linked directly to the AKM variance components. Without this structure, the estimand would generally depend on higher-order and non-identifiable terms capturing nonlinear sorting patterns. At the same time, this restriction is plausible in equilibrium models and settings with stable sorting across different set of individuals and locations. In AKM-style variance decompositions, sorting is routinely summarized by a single covariance (equivalently, the linear-projection slope) between worker and location effects (Card et al., 2013; Kline et al., 2020; Kline, 2024)<sup>50</sup>.

Under A1-A3:

$$\gamma_M := \text{Corr}_M(\psi_{d(i)}, \psi_{o(i)}) \stackrel{\text{Under A3}}{=} \rho_M^2 \Rightarrow \text{Var}_M(\Delta\psi) = 2(1 - \rho_M^2) \sigma_{\psi,M}^2, \quad \sigma_{\psi,M}^2 := \text{Var}_M(\psi_{j(i,t)}). \quad (8)$$

Write  $\bar{Z}_{j,t}$  for the location-time mean of any  $Z_{it}$  computed over *all* individuals in

---

<sup>50</sup>Two familiar environments make an exact affine relationship a plausible theoretical benchmark. (i) In transferable-utility assignment models with quadratic mismatch, stable/equilibrium matching delivers a linear pairing of types (constant sorting intensity); (e.g. Teulings, 1995; Bojilov and Galichon, 2016; Galichon and Salanié, 2022). (ii) From a statistical point, if  $(\alpha, \psi)$  are jointly elliptically contoured with finite first moments, the conditional mean is exactly affine— $\mathbb{E}[\alpha \mid \psi] = a + s\psi$ —so normality is unnecessary.

$(j, t)$  (movers + stayers). The exact leave-one-out identity is

$$\bar{Z}_{j(i,t),t}^{(-i)} = \bar{Z}_{j(i,t),t} - \frac{1}{n_{j(i,t)} - 1} (Z_{it} - \bar{Z}_{j(i,t),t}) \approx \bar{Z}_{j(i,t),t} - c_{j(i,t)} (Z_{it} - \bar{Z}_{j(i,t),t}),$$

so when we apply it to  $Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}$  and difference across  $t = 2, 1$ , we obtain for movers:

$$\delta_i^{(-i)} = \Delta\psi_i + (1 + c_d) \bar{\alpha}_{d,2} - (1 + c_o) \bar{\alpha}_{o,1} + (1 + c_d) \bar{\varepsilon}_{d,2} - (1 + c_o) \bar{\varepsilon}_{o,1} - (c_d - c_o) \alpha_i - c_d \varepsilon_{i2} + c_o \varepsilon_{i1}, \quad (9)$$

where  $\bar{\alpha}_{j,t}$  and  $\bar{\varepsilon}_{j,t}$  are location-time means. Intuitively,  $\delta_i^{(-i)}$  equals the true change in location effects plus two mean-change terms (composition and errors) inflated by  $(1 + c)$  because the leave-one-out mean partially “puts back” individual  $i$ , minus the individual’s own error change scaled by  $c$ .

Using A3, and averaging over all individuals in  $(j, t)$  (movers+stayers) with mover share  $w$  (i.e.,  $w_{j,t} = w$  for all  $(j, t)$ ) and holding shares fixed across firms,

$$\bar{\alpha}_{j,t} = \beta \psi_j + \bar{\varepsilon}_{j,t}, \quad \text{Cov}_M(\psi_j, \bar{\varepsilon}_{j,t}) = 0. \quad (10)$$

with  $\beta := w s_M + (1 - w) s_S$ . The formula makes it clear that  $\beta$  is large when sorting is strong.

Starting from  $\bar{\alpha}_{j,t} = \beta \psi_j + \eta_{j,t}$ , for any mover  $i$ ,

$$\Delta \bar{\alpha}_i = \beta \Delta \psi_i + \Delta \eta_i, \quad \text{with} \quad \text{Cov}_M(\Delta \psi, \Delta \eta) = 0.$$

Using (8),

$$\text{Cov}_M(\Delta \psi, \Delta \bar{\alpha}) = \text{Cov}_M(\Delta \psi, \beta \Delta \psi) = \beta \text{Var}_M(\Delta \psi) = 2(1 - \rho_M^2) \sigma_{\psi, M}^2 \beta, \quad (11)$$

Under A2 (finite cells, fixed mover share, no sorting on size) and independence across individual error terms, write the individual projection  $\alpha_i = s_g \psi_{j(i,t)} + e_{i,t}$  with  $\sigma_{e,g}^2 = \sigma_{\alpha,g}^2 (1 - \rho_g^2)$  and  $\mathbb{E}[e_{i,t} | \psi_{j(i,t)}, g] = 0$ . Averaging over all individuals in  $(j, t)$  gives  $\eta_{j,t} = \bar{\varepsilon}_{j,t} = n_j^{-1} \sum_{i:j(i,t)=j} e_{i,t}$  and hence

$$\text{Var}_M(\eta_{j,t}) = c_j V_e, \quad V_e := w \sigma_{e,M}^2 + (1 - w) \sigma_{e,S}^2.$$

With  $\bar{\alpha}_{j,t} = \beta \psi_j + \eta_{j,t}$  and  $\text{Cov}_M(\psi_j, \eta_{j,t}) = 0$ , for any mover  $i$

$$\text{Var}_M(\Delta \bar{\alpha}_i) = \beta^2 \text{Var}_M(\psi_{d(i)} - \psi_{o(i)}) + \text{Var}_M(\eta_{d,2}) + \text{Var}_M(\eta_{o,1}).$$

By A3,  $\text{Var}_M(\psi_{j(i,t)}) = \sigma_{\psi,M}^2$  and  $\text{Corr}_M(\psi_d, \psi_o) = \rho_M^2$ , so  $\text{Var}_M(\psi_d - \psi_o) = 2(1 - \rho_M^2)\sigma_{\psi,M}^2$ . Independence across individual errors implies the two cell-mean residuals are uncorrelated across  $(d, 2)$  and  $(o, 1)$ , yielding

$$\text{Var}_M(\eta_{d,2}) + \text{Var}_M(\eta_{o,1}) = (c_d + c_o) V_e = 2c V_e.$$

Therefore

$$\text{Var}_M(\Delta \bar{\alpha}) = 2 \left[ c V_e + \beta^2 (1 - \rho_M^2) \sigma_{\psi,M}^2 \right]. \quad (12)$$

The decomposition shows that overall variation in  $\Delta \bar{\alpha}$  reflects both residual averaging noise from finite locations and systematic variation driven by firm effects and individual sorting. When sorting is strong in both groups, average composition aligns closely with firm effects, making  $\beta$  large, but movers tend to change between similar firms, so the correlation between origin and destination effects is high and across-move differences are small. When sorting is concentrated among movers, the composition–firm link strengthens but movers’ destinations also resemble their origins more closely, so the resulting changes in composition across firms remain limited. In contrast, when only stayers sort, movers switch between more distinct firms while  $\beta$  is still influenced by the sorting of stayers, generating larger differences across moves. As locations become large and sampling noise disappears, the remaining variation in  $\Delta \bar{\alpha}$  captures these systematic differences in firm effects and individual sorting patterns that drive the composition gap between origins and destinations.

*Error-mean components.* Because  $\bar{\varepsilon}_{j,t}$  averages  $n_j$  i.i.d. errors with variance  $\sigma_\varepsilon^2$ ,

$$\text{Var}_M(\Delta \bar{\varepsilon}) = 2c \sigma_\varepsilon^2, \quad \text{Var}_M(\Delta \varepsilon) = 2\sigma_\varepsilon^2, \quad \text{Cov}_M(\Delta \bar{\varepsilon}, \Delta \varepsilon) = 2c \sigma_\varepsilon^2.$$

These are purely mechanical finite- $n_j$  consequences of averaging and differencing.

*Mover Design Estimand*

Substituting (9), (11), and (12), and assuming that firm sizes at the time of the move are approximately equal<sup>51</sup>, as well as incorporating the noise components into  $\theta = \text{Cov}_M(\delta^{(-i)}, \Delta Y) / \text{Var}_M(\delta^{(-i)})$ , yields

$$\theta = \frac{2(1 - \rho_M^2) \sigma_{\psi, M}^2 [1 + (1 + c)\beta] + 2c^2 \sigma_\varepsilon^2}{2(1 - \rho_M^2) \sigma_{\psi, M}^2 [1 + (1 + c)\beta]^2 + 2(1 + c)^2 c V_e + 2(c + c^2 - c^3) \sigma_\varepsilon^2}, \quad (13)$$

where

$$\beta = w \frac{\rho_M \sigma_{\alpha, M}}{\sigma_{\psi, M}} + (1 - w) \frac{\rho_S \sigma_{\alpha, S}}{\sigma_{\psi, S}},$$

$$V_e = w \sigma_{e, M}^2 + (1 - w) \sigma_{e, S}^2, \quad \sigma_{e, g}^2 = \sigma_{\alpha, g}^2 (1 - \rho_g^2).$$

Both the numerator and denominator include a noise component that diminishes with increasing average firm size, as well as a systematic component that depends on the sorting behavior of both movers and stayers across locations. The key observation is that the numerator reflects two distinct sorting components. The first arises from the sorting of movers between their origin and destination locations. The second stems from the correlation between movers and their coworkers: even when leave-one-out averages are used, there remains a covariance between movers' and coworkers' location effects beyond the simple correlation between location effects and outcomes.

---

<sup>51</sup>If the firm sizes of the origin and destination differ, the denominator will include an additional noise term that vanishes as firm size increases.

### *The Case of Large Locations*

For simplicity, and following the existing literature, I assume that location sizes are sufficiently large such that noise components are negligible, i.e.,  $c \approx 0$ . Under this assumption, Equation 13 simplifies to:

$$\theta = \frac{1}{1 + \beta}. \quad (14)$$

Following Equation (10), it is straightforward to verify that the absolute value of the right-hand side of Equation (14), i.e.,  $|\frac{1}{1+\beta}|$ , exactly equals the ratio between the standard deviation of movers' location effects,  $\sigma_{\psi,M}$ , and the standard deviation of movers' average location outcomes in the cross-section they visit,  $SD(\bar{Y}) = \sigma_{\psi,M}(1 + \beta)$ .<sup>a</sup>

---

<sup>a</sup>Note that if individual sorting is negative,  $\beta$  may take negative values, and consequently, the mover design event-study estimate will fall outside the unit interval.

### *Discussion*

As shown in Equation 14, under A1-A3, the mover-design event-study estimate represents a standard-deviation share. By construction, it provides a relative comparison between the variation in location effects and the overall observed variation across locations. For applications aimed at quantifying the contribution of location to cross-location outcome differences (e.g., [Finkelstein et al. \(2016\)](#)), this estimate can serve as a useful comparative measure. However, its interpretation involves two important caveats.

First, the estimate is expressed as a standard-deviation share, which is less commonly employed as a measure of variation share in labor economics. In the AKM framework, the conventional measure of interest is the variance share, not the standard-deviation share. Therefore, researchers must take care to distinguish between the two. If the focus is on variance shares, the appropriate quantity is the square of the mover-design estimate.

Second, without information on the AKM quadratic components—namely, the variances of location and individual effects and the degree of sorting—the mover-design estimate is only a relative measure. When the goal is to assess the

absolute contribution of locations to outcomes<sup>52</sup>, or to quantify sorting between individuals and locations, the mover statistic alone is insufficient. The variance of location effects must be recovered by rescaling  $\theta$ ; and for any object involving  $\sigma_\alpha$  or  $\rho$ , recovery is not possible from the mover statistic alone.

Because the AKM framework underlies the mover-design event-study estimator, and since the informativeness of the mover design depends on large-location asymptotics with additional assumptions on the sorting structure (A3), it is generally preferable to estimate the AKM quadratic forms directly. Additionally, the mover design captures only movers' sorting behavior and provides limited information about stayers. However, when the share of movers is small relative to location sizes—making bias correction necessary in AKM estimates—or when the largest connected set in the data is unrepresentative, the mover-design event-study may deliver a more interpretable estimand than the AKM-based quadratic forms. This advantage holds only when the parameter of interest concerns the contribution of locations to cross-location variation, rather than their contribution to individual-level outcomes.

### *Simulations*

I simulate a two-period AKM economy under Assumptions A1–A3. In each replication, worker effects  $\alpha_i \sim \mathcal{N}(0, \sigma_\alpha^2)$  and location effects  $\psi_j \sim \mathcal{N}(0, \sigma_\psi^2)$  are drawn. Locations have (near-)equal sizes, and outcomes satisfy

$$Y_{it} = \alpha_i + \psi_{j(i,t)} + \varepsilon_{it}, \quad \varepsilon_{it} \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

Sorting uses rank-based matching on the latent score

$$\text{latent}_{it} = \rho \left( \frac{\sigma_\psi}{\sigma_\alpha} \right) \alpha_i + \sqrt{1 - \rho^2} e_{it}, \quad e_{it} \sim \mathcal{N}(0, \sigma_\psi^2),$$

so that  $\text{Corr}(\alpha_i, \text{latent}_{it}) = \rho$ . In period 1, workers are assigned by descending  $\text{latent}_{i1}$  to locations ordered by  $\psi_j$ . In period 2, a random set of movers fills newly opened vacancies, again by descending  $\text{latent}_{i2}$ . Leave-one-out firm-time means

---

<sup>52</sup>For instance, in a log-linear structure, the standard deviation of location effects is directly informative about the percentage change in the outcome level for the same individual moving across locations (see [Kline \(2024\)](#)).

$\bar{Y}_{jt}^{(-i)}$  are computed using both movers and stayers.

For movers, I estimate the mover-design regression with individual fixed effects and a time dummy,

$$Y_{it} = \eta_i + \lambda \mathbf{1}\{t = 2\} + \theta \left( \mathbf{1}\{t = 2\} \times [\bar{Y}_{d(i),2}^{(-i)} - \bar{Y}_{o(i),1}^{(-i)}] \right) + u_{it},$$

which is equivalent to regressing  $\Delta Y_i$  on  $r_i = \bar{Y}_{d(i),2}^{(-i)} - \bar{Y}_{o(i),1}^{(-i)}$  with  $\lambda = 0$  following the simple AKM structure.

Each replication yields the mover-design estimate and the realized share of the standard deviation of location effects relative to the cross-sectional standard deviation of average outcomes across locations:

$$\text{SD-share} = \frac{\text{SD}(\psi_{j(i,t)} \text{ for movers across } t = 1, 2)}{\text{SD}_{\text{mover}}(\bar{Y}_{jt})},$$

where  $\bar{Y}_{jt}$  are firm-time means using all workers, and  $\text{SD}_{\text{mover}}(\bar{Y}_{jt})$  is the mover-frequency-weighted cross-sectional SD over the  $(j, t)$  cells visited by movers. I also record realized  $\text{Var}(\alpha_i)$ ,  $\text{Var}(\psi_{j(i,t)})$ , and  $\text{Corr}(\alpha_i, \psi_{j(i,t)})$ .

I run 50 replications for each  $\rho \in \{-0.20, -0.05, 0, 0.05, 0.20\}$  with  $N = 500,000$  observations,  $K = 100$  firms,  $n_{\text{movers}} = 50,000$  movers, and  $(\sigma_\alpha, \sigma_\psi, \sigma_\varepsilon) = (1, 0.14, 0.001)$ .

In Table A17, the mover-design coefficient  $\hat{\theta}$  targets the SD share, and  $\hat{\theta}^2$  targets the variance share; the alignment is close across  $\rho$ . The main deviation appears under strong negative sorting: selection drives movers into  $(j, t)$  cells where location effects and sorting partly offset each other, compressing the across-location mean differences. This compression reduces the denominator in the SD share and makes the variation in leave-one-out means highly sensitive to idiosyncratic or small-sample noise. In that regime, squaring  $\hat{\theta}$  understates the true variance share—classic attenuation caused by a small, noisy denominator. With mild negative sorting the denominator is still compressed but not enough to induce a sign reversal, so  $\hat{\theta} > 1$  and  $\hat{\theta}^2$  exceeds one. As  $\rho$  moves toward and above zero, movers switch between more similar  $\psi$  ranks, the contribution of location to overall dispersion weakens, and both the SD share and  $\hat{\theta}$  decline together.

$\rho$	Mover design $\hat{\theta}$		SD share		Variance share		$\hat{\theta}^2$		Realized moments		
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	$\overline{Var(\alpha)}$	$\overline{Var(\psi)}$	$\overline{Corr(\alpha, \psi)}$
-0.200	-1.915	0.369	2.143	0.496	4.836	2.308	3.800	1.476	1.000	0.019	-0.198
-0.050	1.515	0.049	1.537	0.055	2.366	0.171	2.298	0.151	1.000	0.019	-0.050
0.000	0.990	0.011	0.995	0.011	0.990	0.023	0.980	0.023	1.000	0.020	0.000
0.050	0.732	0.015	0.734	0.014	0.539	0.021	0.535	0.021	1.000	0.020	0.050
0.200	0.406	0.017	0.407	0.017	0.166	0.014	0.165	0.014	1.000	0.019	0.199

Notes: Entries report Monte Carlo means and standard deviations over 50 replications for each  $\rho$ . “Mover design  $\hat{\theta}$ ” is the coefficient from the mover-design regression of  $\Delta Y_i$  on  $r_i = \bar{Y}_{d(i),2}^{(-i)} - \bar{Y}_{o(i),1}^{(-i)}$ , where  $\bar{Y}_{jt}^{(-i)}$  are leave-one-out firm-time means computed using movers and stayers.  $\hat{\theta}^2$  is squared within replication prior to averaging. “SD share” is  $SD(\psi_{j(i,t)})$  for movers, stacking  $t = 1, 2$  divided by  $SD_{\text{mover}}(\bar{Y}_{jt})$ , the mover-frequency-weighted cross-sectional SD of firm-time means across the  $(j, t)$  cells visited by movers;  $\bar{Y}_{jt}$  use all workers. “Variance share” is the square of the SD share, computed within replication prior to averaging. “Realized moments” are computed in the pooled two-period sample for assigned  $\psi_{j(i,t)}$  and  $\alpha_i$ :  $\overline{Var(\alpha)}$ ,  $\overline{Var(\psi)}$ , and  $\overline{Corr(\alpha, \psi)}$ .

Table A17: Simulation: Mover-design estimates and SD/variance shares

## A4 A Model

This appendix provides a simple Grossman-style health-capital model that micro-founds the log-linear two-way fixed-effects specification used in the main text. The aim is to show that, when health depreciation and the time cost of care-seeking have both individual and firm components, optimal medical spending can be written (in logs) as an additive function of an individual term, a firm term, and age. In this setting, the firm-specific component can be interpreted mainly as reflecting the workplace-specific health-depreciation rate, with a secondary role for differences in the time cost of accessing care.<sup>53</sup>

I extend the life-cycle health demand framework of [Grossman \(1972b,a\)](#). Individuals do not derive direct utility from health; instead, health affects utility indirectly through time lost to sickness and through the resources (time and money) devoted to healthcare. This structure is well suited to the Dutch context, where basic health insurance is universal and largely independent of the employer. Firms affect health-related spending primarily by changing the health environment (workplace risks, stress, exposures) and, to a lesser extent, the time needed to obtain care, rather than the price or generosity of basic coverage (see Section II).

The model introduces two firm-specific elements. First, a firm-specific health-depreciation component, capturing how the work environment accelerates or slows health deterioration through physical risks, psychosocial stress, and environmental exposures. Second, a firm-specific time cost of care-seeking, reflecting workplace flexibility, ease of taking time off for appointments, and the presence of occupational health services or screening. In the Dutch institutional setting, the first channel (health depreciation) is likely to be the dominant source of cross-firm differences in basic-package spending, while the second channel is unlikely to generate large or systematic differences (see Section II). I nevertheless keep the time-cost term in the model to make clear that, in principle, firms can affect both the health-depreciation environment and the convenience of care-seeking; in the empirical analysis, I interpret the estimated firm effects mainly through the depreciation channel.

---

<sup>53</sup>The model aims to explain how the workplace affects health and spending on healthcare, rather than why workers choose different firms.

## Setup

Consider an employee working in firm  $j$ , with intertemporal utility

$$\sum_{t=0}^T \beta^t U(Z_t), \quad (15)$$

where  $T$  is the length of life,<sup>54</sup>  $\beta$  is the subjective discount factor, and  $Z_t$  is a composite of consumption goods and leisure:

$$Z_t = Z(C_t, L_t), \quad (16)$$

with  $C_t$  denoting market consumption and  $L_t$  leisure at age  $t$ .

Following [Grossman \(1972b\)](#), poor health reduces the available time endowment (normalised to 1). The employee faces the time constraint

$$T_t^w + T_t^h + L_t = 1 - T_t^{loss}, \quad (17)$$

$$T_t^{loss} = T^{loss}(H_t), \quad (18)$$

where  $H_t$  is the health stock at age  $t$ ,  $T_t^{loss}$  the time lost due to sickness, and  $T^{loss}(\cdot)$  a decreasing function of  $H_t$ . The variables  $T_t^w$  and  $T_t^h$  denote time spent working and time spent on health-care activities (care-seeking) at age  $t$ , respectively.

The health stock evolves according to

$$H_{t+1} - H_t = I_t - \delta_{t,j} H_t, \quad (19)$$

$$I_t = I(\Theta_t), \quad (20)$$

$$T_t^h = \chi_j \Theta_t, \quad (21)$$

where  $I_t$  is health investment at age  $t$ ,  $\delta_{t,j}$  is an age- and firm-specific depreciation rate,  $\Theta_t$  is a composite of medical care goods, and  $\chi_j$  measures the time cost of obtaining a given amount of medical care. A higher  $\chi_j$  means more time must be spent on care-seeking per unit of  $\Theta_t$ ; this reflects lower flexibility, more rigid

---

<sup>54</sup>In line with [Grossman \(1972b,a\)](#), life continuation is predicated on maintaining a minimum health level, denoted as  $\bar{H}$ . Therefore, an individual's lifespan can be affected by their health and health investment decisions.

schedules, or less convenient access to care. In a firm with lower  $\chi_j$ , employees can obtain the same care with less time loss.

At age zero (which may be interpreted as entry into the labour market), the individual faces the full-income budget constraint

$$\sum_{t=0}^T \frac{f_t \Theta_t + p_t C_t}{(1+r)^t} = \sum_{t=0}^T \frac{w_{t,j} T_t^w}{(1+r)^t} + A_0, \quad (22)$$

where  $f_t$  and  $p_t$  are the market prices of medical care and consumption goods, respectively,  $w_{t,j} T_t^w$  is labour income at age  $t$ ,  $w_{t,j}$  is the wage rate in firm  $j$ ,  $A_0$  is initial wealth, and  $r$  is the interest rate. In the Dutch setting, basic-package prices and coverage are set nationally and do not vary across firms; I therefore treat  $f_t$  as homogeneous across firms, so that firm-specific differences in medical spending arise from differences in quantities  $\Theta_t$  rather than from price differences.

Normalising the discount factor so that  $\beta = \frac{1}{1+r}$ , the utility maximisation problem is equivalent to

$$\begin{aligned} & \max_{C_t, L_t, \Theta_t, T_t^h, T_t^w} \sum_{t=0}^T \frac{U(Z_t)}{(1+r)^t} \\ \text{s.t. } & \sum_{t=0}^T \frac{f_t \Theta_t + p_t C_t + w_{t,j} (T_t^h + L_t + T_t^{\text{loss}})}{(1+r)^t} = \underbrace{\sum_{t=0}^T \frac{w_{t,j}}{(1+r)^t} + A_0}_A, \end{aligned} \quad (23)$$

where the term  $w_{t,j} (T_t^h + L_t + T_t^{\text{loss}})$  collects forgone earnings from time spent in care-seeking, leisure, and sickness. The right-hand side  $A$  represents full lifetime potential income (the value of all time at work plus initial wealth).

#### A4.I Optimal Health Investment Level

Focusing on health investments, consider the first-order condition with respect to  $I_{\tau-1}$ . In equilibrium,

$$\frac{\pi_{\tau-1}}{(1+r)^{\tau-1}} = - \sum_{t=\tau}^T w_{t,j} \frac{\frac{\partial H_t}{\partial I_{\tau-1}} \frac{dT_t^{\text{loss}}}{dH_t}(\cdot)}{(1+r)^t}, \quad (24)$$

where

$$\pi_\tau = f_\tau \frac{\partial \Theta_\tau}{\partial I_\tau} + w_{\tau,j} \frac{\partial T_\tau^h}{\partial I_\tau}$$

is the marginal cost of health investment (the combined monetary and time cost at age  $\tau$ ). Equation (24) states that the present value of the cost of an additional unit of health investment at age  $\tau - 1$ ,  $\pi_{\tau-1}$ , equals the present value of the future gains from reduced time lost to sickness (and therefore higher effective labour supply and earnings).

Following the standard recursive manipulation in Grossman (1972b,a), this condition can be rewritten as

$$(1 + r)\pi_{\tau-1} = -w_{\tau,j} \frac{dT^{\text{loss}}}{dH_t}(\cdot) + (1 - \delta_{\tau,j})\pi_\tau, \quad (25)$$

or, equivalently,

$$\pi_{\tau-1} (r + \delta_{\tau,j} - \tilde{\pi}_\tau) = -w_{\tau,j} \frac{dT^{\text{loss}}}{dH_t}(\cdot), \quad (26)$$

where  $\tilde{\pi}_\tau = \frac{\pi_\tau - \pi_{\tau-1}}{\pi_{\tau-1}}$  is the gross growth rate of the marginal cost of health.<sup>55</sup>

#### A4.II Assumptions and a Closed-Form Demand Function

To obtain a closed-form expression for medical care demand, I introduce a set of standard simplifying assumptions.

**(i) Linear investment in health goods.** Assume health investments have constant returns to scale with respect to health goods:

$$I_\tau = I_0 \Theta_\tau, \quad (27)$$

---

<sup>55</sup>For (26) to hold in this simple form, I follow Grossman (1972b,a) and assume that cross-terms of the form  $\delta_{\tau,j} \tilde{\pi}_\tau$  are negligible. This keeps the focus on the main dependence on  $\delta_{\tau,j}$  and  $w_{\tau,j}$  without materially affecting the qualitative implications.

with  $I_0 > 0$  a constant. Introducing diminishing returns would be more realistic but would prevent a simple closed-form solution. Under (27),

$$\pi_\tau = \frac{1}{I_0} (f_\tau + \chi_j w_{\tau,j}). \quad (28)$$

In the Dutch context, the monetary price component  $f_\tau$  under the basic package is largely uniform across firms and, for high spenders, the marginal out-of-pocket price is often zero once the deductible is exhausted. For our purposes, any non-negligible but *firm-invariant*  $f_\tau$  only affects the intercept of log spending and does not contribute to firm-level variation. To simplify the algebra and focus on cross-firm differences, I therefore treat the time-cost term  $\chi_j w_{\tau,j}$  as the relevant driver of variation in  $\pi_\tau$  across firms and write, up to a proportionality constant,

$$\pi_\tau \propto \chi_j w_{\tau,j}. \quad (29)$$

This implies that the growth rate of  $\pi_\tau$  is approximately the growth rate of wages:

$$\tilde{\pi}_\tau \approx \tilde{w}_\tau. \quad (30)$$

**(ii) Sick time as a function of health.** Following the literature (e.g. Galama and Van Kippersluis, 2019), assume

$$T_\tau^{loss} = \phi H_\tau^{-\rho},$$

with  $\phi > 0$  and  $\rho > 0$ . Substituting into (26) yields

$$\delta_{\tau,j} \frac{1}{I_0} \chi_j \left( 1 + \frac{r - \tilde{w}_\tau}{\delta_{\tau,j}} \right) = \rho \phi H_\tau^{-(\rho+1)}. \quad (31)$$

**(iii) Individual and firm components in depreciation and wages.** Extend the functional forms commonly used in the life-cycle literature and assume

$$\delta_{\tau,j} = \delta_0 \delta_j e^{\kappa(\tau)},$$

where  $\delta_0$  and  $\delta_j$  are individual- and firm-specific factors determining the level of health depreciation, and  $\kappa(\tau)$  is an increasing function of age. By normalisation,

the population averages of  $\delta_0$  and  $\delta_j$  equal one, so that  $e^{\kappa(\tau)}$  is the average age profile of depreciation.

Similarly, following [Abowd et al. \(1999\)](#), assume wages factor as

$$w_{\tau,j} = w_0 w_j e^{l(\tau)},$$

where  $w_0$  is a common level term,  $w_j$  a firm-specific wage component, and  $l(\tau)$  an age profile. This implies  $\tilde{w}_\tau \approx l'(\tau)$ .

Taking logs and rearranging, we obtain

$$(1 + \rho) \log(H_\tau) \approx \log(\rho) + \log(\phi) + \log(I_0) - \log(\delta_0) - \log(\delta_j) - \log(\chi_j) - \kappa(\tau) - \log\left(1 + \frac{r - l'(\tau)}{\delta_{\tau,j}}\right). \quad (32)$$

Assuming either  $(r - l'(\tau)) \ll \delta_{\tau,j}$  or, more generally, that the last term in (32) can be well approximated by a function of age only,<sup>56</sup> we can write

$$\log(H_\tau) \approx \varepsilon \left[ \log(\rho) + \log(\phi) + \log(I_0) - \log(\delta_0) - \log(\delta_j) - \log(\chi_j) - \check{\kappa}(\tau) \right], \quad (33)$$

where  $\varepsilon = \frac{1}{\rho+1}$  and  $\check{\kappa}(\tau)$  is an age function that absorbs  $\kappa(\tau)$  and the term in brackets in (32).

Using (19) and (27), and approximating health around a local steady state so that  $I_\tau \approx \delta_{\tau,j} H_\tau$ , we obtain an expression for medical care goods:

$$\begin{aligned} \log(\Theta_\tau) \approx & \underbrace{\varepsilon \log(\rho) + \varepsilon \log(\phi)}_C + \underbrace{(\varepsilon - 1) \log(I_0) + (1 - \varepsilon) \log(\delta_0)}_{\alpha_0} \\ & + \underbrace{(1 - \varepsilon) \log(\delta_j) - \varepsilon \log(\chi_j)}_{\alpha_j} + \underbrace{\kappa(\tau) - \varepsilon \check{\kappa}(\tau)}_{\eta(\tau)} + \log\left(1 + \frac{\varepsilon \check{\kappa}'(\tau)}{\delta_{\tau,j}}\right). \end{aligned} \quad (34)$$

As before, assuming either  $\varepsilon \check{\kappa}'(\tau) \ll \delta_{\tau,j}$  or that the last term in (34) can be

<sup>56</sup>Formally, one can either neglect the last term in (32) or absorb it into an age-specific term, provided that the dependence on  $\delta_0$  and  $\delta_j$  is weak. This is a standard approximation in health-capital models and preserves the separability between age and individual/firm factors.

approximated by an age function,<sup>57</sup> we obtain the compact expression

$$\log(\Theta_\tau) \approx C + \alpha_0 + \alpha_j + \check{\eta}(\tau), \quad (35)$$

where  $\check{\eta}(\tau)$  is an age function and

$$\alpha_j = (1 - \varepsilon) \log(\delta_j) - \varepsilon \log(\chi_j).$$

In the Dutch basic-insurance system, the per-unit price of covered care is set nationally and does not vary across firms. Total basic-package spending at age  $t$  can therefore be written as  $c_t \propto f_t \Theta_t$ , with  $f_t$  common across firms. Taking logs and using (35), we obtain

$$\log c_t \approx \underbrace{\alpha_0 + C + \log f_t}_{\text{individual term}} + \underbrace{\alpha_j}_{\text{firm term}} + \underbrace{\check{\eta}(t)}_{\text{age term}}.$$

For a given worker  $i$ , let  $c_{it}$  denote basic-package healthcare costs at age  $t$ , and define  $h_{it} = \log(c_{it} + 1)$  as in the main text. Relabelling the individual component as  $\alpha_i$ , the firm component as  $\psi_j$ , and the age term as part of  $x'_{it}\beta$  yields exactly the log-linear two-way fixed-effects structure used in the main text:

$$h_{it} = \log(c_{it} + 1) \approx \alpha_i + \psi_j + x'_{it}\beta + r_{it},$$

with  $r_{it}$  collecting idiosyncratic shocks.

Two main observations follow from this framework. First, in a Grossman-type health-investment model with firm-specific health depreciation and time costs of care-seeking, the logarithm of healthcare expenditures can be decomposed into an individual-specific term, a firm-specific term, and an age profile. The individual term reflects the individual-specific depreciation factor and returns to healthcare investments (through  $\delta_0$  and  $I_0$ ). The firm-specific term  $\alpha_j$  captures a scaled firm-specific depreciation factor  $\delta_j$  and, in principle, the firm's flexibility for care-seeking through  $\chi_j$ . In the Dutch institutional context, where basic prices and coverage are uniform across firms and time-related barriers to care are limited,

---

<sup>57</sup>Here, I assume either the last term in equation (34) is negligible, or it can be absorbed into  $\eta(\tau)$ , leaving the dependence on  $\delta_0$ ,  $\delta_j$ , and  $\chi_j$  unchanged.

$\delta_j$  provides the most natural interpretation of the estimated firm effects  $\psi_j$  in the main text as workplace-specific health-depreciation rates.

Second, although health and healthcare expenditures result from multiplicative interactions between individual and firm factors (for example through products of  $\delta_0$ ,  $\delta_j$ , and  $\chi_j$ ), taking logarithms maps these multiplicative effects into an additively separable function of individual and firm components plus an age term. This log-additive structure is precisely what underlies the AKM specification in the main analysis and justifies interpreting the firm fixed effects  $\psi_j$  as a common proportional shift in healthcare spending across workers within a firm.

## A5 Leave-Out Variance Correction for the AKM Model

This appendix summarises how I implement the leave-out variance correction of [Kline et al. \(2020\)](#) (KSS) for the AKM specification in Equation (4). The goal is to obtain unbiased estimates of variance and covariance components of worker and firm effects in the presence of limited mobility.

### A5.I Quadratic-form representation and plug-in bias

Write the AKM model in stacked form as

$$h = Zv + \varepsilon,$$

where  $h$  collects log healthcare expenditures  $h_{it}$ ,  $Z$  is the design matrix containing worker and firm dummies and observed covariates, and  $v$  stacks all parameters,  $v = (\alpha_1, \dots, \psi_1, \dots, \beta')'$ . Let  $z'_o$  denote the  $o$ -th row of  $Z$ , and  $S_{zz} = \sum_{o=1}^n z_o z'_o$  the Gram matrix.

Many objects of interest in the AKM decomposition can be written as quadratic forms

$$\theta = v' A v,$$

for some known symmetric matrix  $A$ ; for example,  $Var(\psi_j)$  and  $Cov(\alpha_i, \psi_j)$  correspond to particular choices of  $A$ .<sup>58</sup> The natural plug-in estimator is

$$\hat{\theta}_{PI} = \hat{v}' A \hat{v},$$

where  $\hat{v}$  is the OLS estimator from Equation (4). As shown by [Kline et al. \(2020\)](#),  $\hat{\theta}_{PI}$  is generally biased when worker mobility is limited: even if each element of  $\hat{v}$  is unbiased, the quadratic form combines true dispersion and the dispersion of the estimation error.

Under weak conditions, KSS show that the bias of the plug-in estimator can be written as

$$\text{Bias}(\hat{\theta}_{PI}) = \mathbb{E}[\hat{\theta}_{PI}] - \theta = \sum_{o=1}^n z'_o S_{zz}^{-1} A S_{zz}^{-1} z_o \sigma_o^2,$$

<sup>58</sup>See [Kline et al. \(2020\)](#) for explicit  $A$  matrices for standard AKM variance components.

where  $\sigma_o^2 = \mathbb{E}[\varepsilon_o^2]$  is the observation-specific error variance. Limited mobility bias is thus a weighted sum of these variances, with weights given by the leverage-like terms  $z_o' S_{zz}^{-1} A S_{zz}^{-1} z_o$ .

#### A5.II Leave-out error-variance estimation

To obtain an unbiased estimator of  $\theta$ , KSS propose estimating each  $\sigma_o^2$  using a leave-one-out construction. Let  $\hat{v}_{-o}$  denote the OLS estimator from Equation (4) computed on the sample that excludes observation  $o$ . The leave-out prediction for  $h_o$  is  $z_o' \hat{v}_{-o}$  and the associated error-variance estimator is

$$\hat{\sigma}_o^2 = h_o (h_o - z_o' \hat{v}_{-o}). \quad (36)$$

By construction,  $h_o$  is independent of  $\hat{v}_{-o}$  (because  $h_o$  is not used in its estimation), and  $z_o' \hat{v}_{-o}$  is an unbiased predictor of  $h_o$  conditional on  $Z$ . It follows that  $\mathbb{E}[\hat{\sigma}_o^2 | Z] = \sigma_o^2$ , so (36) is an unbiased estimator of the observation-specific error variance; see [Kline et al. \(2020\)](#) for a formal derivation.

Given  $\hat{\sigma}_o^2$ , an unbiased estimator of  $\theta$  is then

$$\hat{\theta}_{KSS} = \hat{v}' A \hat{v} - \sum_{o=1}^n z_o' S_{zz}^{-1} A S_{zz}^{-1} z_o \hat{\sigma}_o^2, \quad (37)$$

which satisfies  $\mathbb{E}[\hat{\theta}_{KSS}] = \theta$  under the conditions in [Kline et al. \(2020\)](#). Intuitively,  $\hat{\theta}_{KSS}$  subtracts an estimate of the plug-in bias from the naive quadratic form.

In large panels, recomputing  $\hat{v}_{-o}$  separately for every  $o$  is infeasible. KSS show that for linear models the leave-one-out prediction error can be written as

$$h_o - z_o' \hat{v}_{-o} = \frac{h_o - z_o' \hat{v}}{1 - \ell_o},$$

where  $\ell_o = z_o' S_{zz}^{-1} z_o$  is the leverage of observation  $o$ . This identity implies that the leave-out prediction  $z_o' \hat{v}_{-o}$ , and therefore  $\hat{\sigma}_o^2$ , can be computed from the full-sample OLS quantities without repeated re-estimation.

### A5.III Implementation for the AKM decomposition

In practice, I implement the KSS correction as follows. First, I partial out the observed covariates  $x_{it}$  from the outcome by estimating Equation (4) and forming the residualised dependent variable  $h_{it} - x'_{it}\hat{\beta}$ . The dimensionality problem affects the worker and firm effects, not the coefficients on  $x_{it}$ , so  $\hat{\beta}$  is estimated precisely.

Second, I apply the leave-out formulas of [Kline et al. \(2020\)](#) to the worker–firm match level, treating an entire match (rather than a single person–year observation) as the leave-out unit to allow for serial correlation of errors within matches, as recommended by KSS. Using the quadratic-form representation, I construct  $\hat{\theta}_{KSS}$  in (37) for the variance of firm effects,  $Var(\psi_j)$ , and for the covariance between worker and firm effects,  $Cov(\alpha_i, \psi_j)$ .

Finally, because the AKM design matrix is very large, I use the random-projection approximation proposed by [Kline et al. \(2020\)](#), based on the database-friendly projections of [Achlioptas \(2003\)](#), to evaluate the quadratic forms in (37) efficiently. This approximation delivers estimates that are numerically very close to the exact leave-out quantities at a fraction of the computational cost; see [Kline et al. \(2020\)](#) for details and accuracy guarantees.

## A6 Deconvolving the Firm Effects

Throughout, I work with the worker–year weighted distribution of latent firm effects. In the worker-year panel, each observation  $i$  is employed at firm  $j(i, t)$  and therefore carries the corresponding firm effect. The object of interest is the distribution of  $\psi_{j(i,t)}$  on the worker–year scale.

### A6.I Parametric benchmark: homoscedastic Gaussian measurement error

The AKM estimation produces firm effects  $\hat{\psi}_j$ . In the worker-year panel, each observation  $i$  is employed at firm  $j(i, t)$  and therefore carries the corresponding estimated firm effect  $\hat{\psi}_{j(i,t)}$ . The object of interest is the distribution of latent firm effects at the worker-year level, i.e. the distribution of  $\psi_{j(i,t)}$ . I treat the worker-year series  $\hat{\psi}_{j(i,t)}$  as a noisy measurement of the latent firm effect:

$$\hat{\psi}_{j(\cdot)} = \psi_{j(\cdot)} + \varepsilon_{j(\cdot)}, \quad \varepsilon_j \overset{\text{approx}}{\sim} N(0, \sigma_\varepsilon^2),$$

where  $\sigma_\varepsilon^2$  is a single (average) estimation-error variance. These are a working approximation, and violated by construction. The estimation error is firm-specific, estimated through limited mobility in some firms, and shared across worker-years within a firm (so repeated values within firm are perfectly dependent through  $\varepsilon_j$ ). Later, in Appendix A6.II, I relax these structural assumptions by using a non-parametric estimation procedure leveraging a sample splitting approach.

I estimate  $\sigma_\varepsilon^2$  using a KSS-based variance decomposition. Let  $\widehat{\text{Var}}(\hat{\psi}_{j(\cdot)})$  denote the sample variance of  $\hat{\psi}_{j(\cdot)}$  across worker-years. Let  $\widehat{\text{Var}}^{\text{KSS}}(\psi)$  denote the KSS bias-corrected estimate of the variance of latent firm effects on the same worker-year weighted scale. I set

$$\hat{\sigma}_\varepsilon^2 = \widehat{\text{Var}}(\hat{\psi}_{j(\cdot)}) - \widehat{\text{Var}}^{\text{KSS}}(\psi)$$

I standardise the worker-year series by  $\hat{\sigma}_\varepsilon$ ,

$$z_i = \frac{\hat{\psi}_{j(\cdot)}}{\hat{\sigma}_\varepsilon},$$

so that  $z_i \mid \theta_{j(\cdot)}^* \sim N(\theta_{j(\cdot)}^*, 1)$  with  $\theta_j^* = \psi_j / \hat{\sigma}_\varepsilon$ . I then estimate the mixing distribution of  $\theta^*$  using `deconvolveR::deconv()` with `family = "Normal"` on a grid  $\{\tau_k\}_{k=1}^K$

spanning the support of  $z_i$ . The procedure fits a penalised g-model for the discrete masses over the grid, with spline degree `pDegree` and regularisation strength `c0`. In the implementation I set `pDegree = 5` and `c0 = 1`. From the estimated discrete mixing distribution  $\{(\tau_k, \hat{g}_k)\}$ , I form the CDF  $\hat{F}^*(t) = \sum_{k: \tau_k \leq t} \hat{g}_k$  and obtain the 20th, 40th, 60th, and 80th percentiles  $q_p^*$  of the latent distribution of  $\theta^*$ . Mapping back to the original scale yields worker-year weighted quintile cut-offs for  $\psi$ :

$$q_p = \hat{\sigma}_\varepsilon q_p^*, \quad p \in \{0.2, 0.4, 0.6, 0.8\}.$$

Finally, the mean latent firm effect within each quintile  $q \in \{1, \dots, 5\}$  is computed as the mean of  $\psi$  under the estimated mixing distribution truncated to that quintile:

$$\mathbb{E}[\psi \mid \psi \in q] \approx \hat{\sigma}_\varepsilon \frac{\sum_{k \in q} \tau_k \hat{g}_k}{\sum_{k \in q} \hat{g}_k}.$$

I report the difference between the top and bottom quintile means as a deconvolution-based firm gradient for the worker-year weighted latent distribution.

#### *A6.II Replicated AKM estimates and nonparametric deconvolution*

AKM firm effects are estimated with error, and the precision of  $\hat{\psi}_j$  varies systematically with firm size, mobility patterns, and network connectivity. Moreover, the estimation error is firm-specific: all worker-years in firm  $j$  inherit the same  $\hat{\psi}_j$ . This makes an i.i.d. measurement-error approximation for the worker-year series  $\hat{\psi}_{j(i,t)}$  inappropriate. Building on sample-splitting approaches that construct replicated fixed-effect estimates in two-way fixed effects models (Kline, 2025), I implement a split-sample replicated-AKM procedure that delivers two noisy but conditionally independent measurements of each firm's AKM effect. I then treat these replicated firm effects as a two-measurement deconvolution problem and use Kotlarski-type methods to nonparametrically recover the latent distribution of  $\psi_{j(i,t)}$  (Kotlarski, 1967; Anarat et al., 2025). This yields a noise-corrected estimate of the full firm-effect distribution, including tail objects such as quintile cut-offs and top–bottom quintile mean gaps, while allowing for heterogeneous precision across firms. To the best of my knowledge, this is the first application in the matched employer–employee AKM literature that uses replicated firm effects with

Kotlarski-type deconvolution to recover the full latent distribution of firm effects.

To obtain independent noise, I randomly partition workers into two mutually exclusive sets,  $S = 1$  and  $S = 2$ , and estimate the AKM model separately in each subsample (including all worker-years for workers in that subsample). I restrict attention to firms that belong to the intersection of two splits. To ensure comparability across splits, I impose the same normalization in each estimation by recentering firm effects so that  $\sum_i \hat{\psi}_{j(i,t)}^{(s)} = 0$  for  $s \in \{1, 2\}$ . This produces two firm-effect estimates for each firm in the connected set,

$$\hat{\psi}_j^{(1)} = \psi_j + u_j^{(1)}, \quad \hat{\psi}_j^{(2)} = \psi_j + u_j^{(2)},$$

where  $u_j^{(1)}$  and  $u_j^{(2)}$  are estimation errors induced by sampling variation in the two disjoint worker sets. Under random partitioning and the maintained AKM assumptions,  $u_j^{(1)} \perp u_j^{(2)} \mid \psi_j$  and  $\mathbb{E}[u_j^{(s)} \mid \psi_j] = 0$  for  $s \in \{1, 2\}$  (Kline, 2025). In the worker-year panel, the replicated series are  $\hat{\psi}_{j(i,t)}^{(s)}$  for  $s \in \{1, 2\}$ .

A useful implication is that the cross-split covariance identifies the variance of latent firm effects on the worker-year scale:

$$\text{Cov}(\hat{\psi}_{j(\cdot)}^{(1)}, \hat{\psi}_{j(\cdot)}^{(2)}) = \text{Var}(\psi_{j(\cdot)}).$$

I use this identity to test the robustness of the KSS bias-corrected variance estimates computed in the main sample. The covariance of the two firm effects is 0.026 which is very close to the KSS variance estimated in the main sample (Table 3). Moreover, the KSS variance estimated in each split sample is similar close to the total sample (Table A18).

The split-sample construction delivers a replicated-measurement system for the latent firm effect. Abstracting from the AKM context, suppose

$$Z_1 = Y + X_1, \quad Z_2 = Y + X_2,$$

with  $Y \perp (X_1, X_2)$  and  $X_1 \perp X_2$ . Independence implies the joint characteristic function factorises as

$$\varphi_{Z_1, Z_2}(t_1, t_2) = \varphi_Y(t_1 + t_2) \varphi_{X_1}(t_1) \varphi_{X_2}(t_2).$$

Statistic	Split 1 (KSS)	Split 2 (KSS)
<i>Panel A. AKM variance components estimated within split (KSS-corrected)</i>		
Log Health Exp. Var.: $\text{var}(h_i)$	2.632 (Ref.)	2.622 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.030 (1.13 %)	0.029 (1.11 %)
Employee FE Var.: $\text{var}(\alpha_i)$	1.598 (60.73 %)	1.486 (56.68 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	−0.007	−0.004
Correlation $\rho(\alpha_i, \psi_j)$	−0.033	−0.020
Explained share: workers & firms	0.613	0.575
No. Movers	576,056	576,410
No. Firms	84,852	85,160
No. of Observations	8,023,188	8,021,807
Average Health Exp. : $\text{mean}(h_i)$	5.614	5.613
<i>Panel B. Cross-split covariance of firm effects (common firms)</i>		
$\text{cov}(\hat{\psi}_j^{(1)}, \hat{\psi}_j^{(2)})$		0.026

*Notes:* Columns report leave-one-match-out KSS-corrected components (Kline et al., 2020) from AKM estimated separately in two random worker splits (all worker-years for workers in a split are included). Outcome is log healthcare expenditure per calendar year; sample spans full-time employees, 2009–2016. Shares in parentheses are relative to  $\text{var}(h_i)$  within the corresponding split. Panel B reports the covariance between firm fixed effects estimated in the two splits, computed on the set of firms observed in both splits. Under independent estimation error across splits, this covariance identifies the variance of the latent firm effect distribution.

Table A18: Variance Decomposition of Log Healthcare Expenditure: Sample Splits

Kotlarski’s lemma implies that  $\varphi_Y$  (and hence the distribution of  $Y$ ) is nonparametrically identified from the joint law of  $(Z_1, Z_2)$  under mild regularity conditions (in particular, non-vanishing characteristic functions in a neighbourhood of the origin) (Kotlarski, 1967). In my application,  $(Z_1, Z_2)$  corresponds to  $(\hat{\psi}_{j(\cdot)}^{(1)}, \hat{\psi}_{j(\cdot)}^{(2)})$  on the worker-year scale, and  $Y$  corresponds to  $\psi_{j(\cdot)}$ .

While identification is straightforward, estimation is numerically delicate. Deconvolution is an ill-posed inverse problem: at high frequencies  $|t|$ , characteristic functions can be small, so sampling error in  $\hat{\varphi}$  is amplified when forming ratios and performing Fourier inversion, which can generate oscillatory densities and unstable tail functionals. I therefore estimate the latent distribution of  $\psi_{j(\cdot)}$  using the NPDFD package (Anarat et al., 2025), which implements a Fourier-based deconvolution procedure for two-measurement settings with explicit regularisation. In the main specification I use `mode = "empirical"`, which forms characteristic functions using the empirical analogue.

Table A19 shows that the two parametric benchmarks deliver similar estimate of firm-effect inequality. The normal approximation calibrated to  $\widehat{Var}^{KSS}(\psi)$  implies a top–bottom quintile mean gap of 0.46, and the homoskedastic Gaussian-error deconvolution delivers 0.44. This tight agreement indicates that, in this setting, the KSS-corrected variance maps cleanly into tail inequality measures, and that replacing the latent distribution by a flexible deconvolution estimate under an average Gaussian error variance does not meaningfully change the implied quintile gap. By contrast, the split-sample replicated-AKM nonparametric deconvolution delivers a larger gap of 0.67, but it also implies a latent variance around 0.037, which exceeds the variance anchored by the KSS/covariance evidence. This variance inflation is an evidence that the nonparametric deconvolution is over-dispersing the latent distribution and mechanically pushing mass into the tails. I therefore interpret the 0.67 estimate as an upper bound on firm-effect inequality, with the 0.44–0.46 range providing a tightly disciplined baseline.

Statistic	Normal $\psi$ (KSS var.)	Gaussian error, homosk.	Split-sample, nonparam.
$\mathbb{E}[\psi   Q_5] - \mathbb{E}[\psi   Q_1]$	0.460	0.437	0.673

Notes: Each column reports  $\mathbb{E}[\psi | Q_5] - \mathbb{E}[\psi | Q_1]$ , where  $Q_1$  and  $Q_5$  are the bottom and top quintiles of the *latent* firm-effect distribution. Column 1 assumes  $\psi \sim N(0, \widehat{Var}^{KSS}(\psi))$  and computes the gap analytically. Column 2 assumes homoskedastic Gaussian estimation error in  $\hat{\psi}_j = \psi_j + \varepsilon_j$  with  $\varepsilon_j \sim N(0, \hat{\sigma}_\varepsilon^2)$  and estimates the latent distribution by deconvolution (no parametric restriction on  $\psi$ ). Column 3 uses split-sample replicated AKM estimates and Kotlarski-type nonparametric deconvolution.

Table A19: Firm-Effect Inequality: Top–Bottom Quintile Mean Gap across Corrections

## A7 LLM-based collar classification from education codes

Because the administrative records do not contain occupational titles, I proxy blue- versus white-collar status using educational specialization. The data include Dutch SOI education codes and, for each code, an education/program title in Dutch (OPLEIDINGSNAAM). I treat this education title as a concise summary of the type of work the individual is trained for, and I map it into a broad collar category.

I implement this mapping using an automated large-language-model (LLM) classification routine. For each education title, I query the OpenAI Chat Completions API (model `gpt-4o-mini`) with a fixed prompt that asks the model to assign the title to one of three categories: (1) *blue collar*, for programs that typically lead to manual, physical, trade, or craft occupations; (2) *white collar*, for programs that typically lead to office-based, professional, managerial, or administrative occupations; and (3) *borderline*, for ambiguous or mixed programs (for example, titles that plausibly combine technical and supervisory content or are too generic to map cleanly). The prompt restricts the model to return only the numeric label (1–3), with no accompanying text. I run the model with temperature zero and a short output constraint, so the classification is deterministic and machine-readable.

In the analysis, I collapse these labels into the collar proxy used in the subsample exercises. White-collar workers are those whose education titles are classified as category (2), and blue-collar workers are those classified as category (1). The borderline category (3) captures mixed or insufficiently specific education titles; they are excluded from collar subsample analysis.

This approach is motivated by the fact that the task is a coarse, low-dimensional coding problem: the goal is not to infer a specific occupation, but to distinguish whether a program is predominantly geared toward manual/trade work versus office/professional work. A growing literature finds that LLMs achieve human-comparable accuracy in such straightforward classification settings. [Gilardi et al. \(2023\)](#) show that GPT-4 outperforms human crowdworkers in text annotation. [Törnberg \(2023\)](#) documents expert-level reliability in political text coding. [Chang et al. \(2024\)](#) report similar evidence from large-scale evaluations in social science contexts. Since mapping education titles into broad collar categories is simpler than the nuanced political or attitudinal coding studied in these papers, the resulting

proxy is expected to be reliable for the purpose of stratifying the AKM estimates by worker type.

## A8 Female Sample Analysis

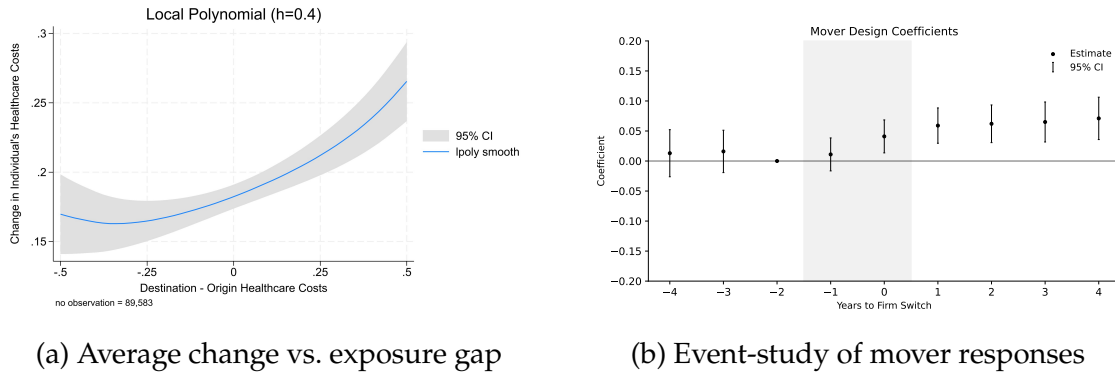


Figure A13: Workers' healthcare spending around moves between firms

*Note:* Sample of full-time female employees aged 25–65 who switch employer exactly once between 2009 and 2016 and belong to the connected set of firms. Healthcare outcomes are log total basic-package expenditures, residualised on age and calendar-year dummies. In panel (a), each point plots the average change in individual annuale residualised spending between the years before and the years after the move against the “exposure gap”  $D_i$ , defined as the difference in leave- $i$ -out mean residualised spending between destination and origin firms; the line shows the (non-parametric) fitted relationship. In panel (b), the markers show coefficients from the mover event-study specification in equation (3), where the outcome is residualised log spending and the regressors are event-time indicators interacted with  $D_i$ , with event time measured relative to  $t = -2$  as the reference year. Because workers are assigned to the firm where they spend most months, the year immediately before the move already contains some exposure to the destination firm.

Statistic	KSS Estimates
<i>Panel A. KSS variance components</i>	
Log Health Exp. Var.: $\text{var}(h_i)$	2.617 (Ref.)
Firm FE Var.: $\text{var}(\psi_j)$	0.020 (0.78 %)
Employee FE Var.: $\text{var}(\alpha_i)$	3.158 (120.70 %)
Firm–Employee Cov.: $\text{cov}(\alpha_i, \psi_j)$	−0.004
Correlation $\rho(\alpha_i, \psi_j)$	−0.017
Explained share: workers & firms	1.211
No. Movers	366,643
No. Firms	44,578
No. of Observations	4,784,634
Average log Health Exp.: $\text{mean}(h_i)$	6.114
<i>Panel B. Cross-sample link</i>	
Bias-corrected slope, firm FE (male) on firm FE (female) <sup>a</sup>	1.149
No. worker–year obs. in slope regression <sup>a</sup>	4,607,681

<sup>a</sup> Slope from an OLS regression of the firm fixed effects estimated in the male subsample on the firm fixed effects estimated in the female subsample, using worker–year observations in the overlap where both firm effects are defined. The reported slope is bias-corrected by rescaling the attenuated OLS coefficient using the reliability ratio for the right-hand-side firm effect,  $\text{Var}_{KSS}(\psi)/\text{Var}_{PI}(\hat{\psi})$ , following [Kline et al. \(2020\)](#) (equivalently,  $\hat{\beta}_{BC} = \hat{\beta}_{OLS} \times \text{Var}_{PI}(\hat{\psi}_{\text{female}})/\text{Var}_{KSS}(\psi_{\text{female}})$ ). In this application,  $\hat{\beta}_{OLS} = 0.151$ ,  $\text{Var}_{PI}(\hat{\psi}_{\text{female}}) = 0.155$ , and  $\text{Var}_{KSS}(\psi_{\text{female}}) = 0.020$ , yielding  $\hat{\beta}_{BC} = 1.149$ . Because KSS is implemented on the leave-one-out connected set within each sex-specific subsample, the male and female samples need not form an exact partition of the baseline connected set; some firms appear in only one subsample and therefore drop out of the overlap regression.

*Notes:* Panel A components are estimated via leave-one-match-out KSS correction ([Kline et al., 2020](#)) using 50 random projections. The outcome is log annual basic-package health-care expenditure per calendar year,  $h_{it} = \log(c_{it} + 1)$ . Shares in parentheses are relative to  $\text{var}(h_i)$  and need not sum to 100% when the firm–worker covariance is negative.

Table A20: Variance Decomposition of Log Healthcare Expenditure (Females)