Labor Market Competition and Job Loss*

Theodor Friederiszick† Axel Malmcrona‡ Nelson Sunuwar§

March 2026

Abstract

Mergers can harm workers both by weakening wage competition among their employers and by eliminating jobs through post-merger closures. The magnitude of both effects depends on how substitutable workers' jobs are, yet standard occupation and geography classifications do not capture this well. We develop a framework that measures job-level substitutability using employee diversion ratios: the probabilities of where workers would go if their job disappeared. We estimate a structural model of labor supply using Swedish administrative data covering more than 2.6 million workers and 615,000 jobs. Substitutability varies sharply across jobs, even within the same occupation and geography. We show that it is a powerful predictor of displacement costs: workers who lose jobs with many close substitutes recover earnings substantially faster, while those who lose jobs with few close substitutes experience larger and more persistent losses. The same framework delivers the objects needed for prospective merger screening: employer-level labor overlap for assessing wage effects on incumbent workers, and job-level diversion ratios for assessing harm from post-merger closures.

1 Introduction

Mergers can harm workers in two distinct ways. They can weaken wage competition for incumbent workers by reducing the set of employers that compete for their labor, and they can harm displaced workers when post-merger closures force them into worse alternatives. The magnitude of both effects depends on worker substitution patterns: how strongly the merging firms compete for the same workers, and where workers would go if their job disappeared. Standard approaches that define labor markets using occupation and

*Theodor Friederiszick and Nelson Sunuwar gratefully acknowledge funding from the University Research Priority Program (URPP) Equality of Opportunity. Theodor Friederiszick is especially grateful to his advisors Greg Crawford, David Dorn, Gautam Gowrisankaran, and Pietro Tebaldi for their guidance and support. We thank participants at the Columbia IO Colloquium and the NYU IO Colloquium, as well as Chris Conlon, Simon Jäger, Andreas Mueller, and Suresh Naidu for helpful comments and suggestions.

†University of Zurich. Email: theodor.friederiszick@econ.uzh.ch

‡University of Gothenburg. Email: axel.malmcrona@gu.se

§University of Zurich. Email: nelson.sunuwar@econ.uzh.ch

commuting zone do not capture these patterns well, because workers often substitute across both occupational and geographic boundaries.

We recover these substitution patterns from a structural model of labor supply in which heterogeneous workers choose among differentiated jobs, defined as occupation × establishment. The model recovers employee diversion ratios: for each worker-job pair, the probabilities of where the worker would go if that job disappeared. Rather than first defining a labor market and then asking whether two firms compete within it, we estimate substitution patterns across jobs and aggregate them to the job and employer level. This yields both a job-level measure of substitutability and an employer-level measure of labor overlap between firms.

We estimate the model using Swedish administrative data covering more than 2.6 million workers and approximately 615,000 jobs. Substitutability varies sharply across jobs, even within occupations and geographies. To validate the framework, we study plant closures and show that job-level substitutability is a powerful predictor of displacement costs. Workers displaced from jobs in the top quartile of substitutability experience an average earnings drop of 18 percent in the first year and nearly fully recover within seven years. Workers displaced from jobs in the bottom quartile suffer a 32 percent drop and still earn 15 percent less than baseline seven years later. Decomposing the index into its geographic and skill-based components shows that these channels matter at different horizons: proximity to alternative employers determines how quickly workers find new jobs, while the portability of their occupation and education determines whether those new jobs restore their prior earnings.

Recent merger challenges make clear that there is real demand for tools that measure how workers substitute across jobs. In challenging Kroger's proposed acquisition of Albertsons, the FTC alleged that the merger would weaken competition for union grocery labor, while the merging parties argued that workers could readily move to other retail jobs. In discussing the labor theory, the court noted that the evidence was "primarily anecdotal" and lacked economic modeling of how wages would change (United States District Court for the District of Oregon, 2024). Our framework delivers exactly these objects: employer-level labor overlap for screening potential wage effects on incumbent workers, and job-level diversion ratios for assessing harm from post-merger closures.

The paper connects to four strands of research. A first literature studies how mergers and acquisitions affect workers. Prager and Schmitt (2021) show that hospital mergers reduce wage growth for pharmacists and nurses but not for janitors. Arnold (2021) finds that ownership changes that increase local labor market concentration reduce worker earnings, with larger effects in already concentrated markets. Geurts and Van Biesebroeck

(2019) document that mergers motivated by market power lead to permanent employment reductions, whereas efficiency-motivated mergers can increase employment. Beuschlein et al. (2025) show that acquisitions reduce wages for workers who remain at the target firm. These papers establish that mergers can harm workers through multiple channels. Because they study realized mergers, however, they are informative about the effects of past transactions rather than prospective ones. Competition authorities also need ex ante tools that can assess where worker harm is likely to arise before a merger is consummated. We contribute to this literature by providing such a tool: job-level substitutability and employer-level labor overlap recovered from diversion ratios.

The paper also relates to work on job substitution, outside options, and job amenities. Using job-to-job flows, Sorkin (2018) shows that mobility patterns reveal substantial differences in non-wage job amenities. Caldwell and Danieli (2024) document that workers in the same job can nevertheless face heterogeneous outside options, with commuting as a key margin. Fan (2025) shows that substitutability declines with occupational distance. Relative to this literature, we recover a full directional diversion matrix at the occupation-by-establishment level. This lets us summarize substitutability with a simple index while retaining the pairwise substitution patterns needed for merger analysis.

A third strand asks how labor markets should be defined. Nimczik (2023) defines labor markets endogenously from worker flows. Fogel and Modenesi (2023) classify workers and jobs into markets based on comparative advantage, and show that these data-driven classifications outperform standard occupation and industry groupings in predicting wage responses to shocks. Jarosch et al. (2024) show that large employers can have market power because workers' outside options exclude the firm's own vacancies. These papers make clear that labor-market boundaries are difficult to draw cleanly. Following Farrell and Shapiro (2010) and Conlon and Mortimer (2021), we take a diversion-based approach instead. The advantage is that we do not need to define markets first. We estimate worker substitution directly and then aggregate it into objects that are relevant for wage competition, worker reallocation, and merger overlap.[1]

Finally, our validation exercise connects to the displacement literature. Jacobson et al. (1993) document large and persistent earnings losses after job displacement. Using Swedish administrative data, Eliason and Storrie (2006) show that these losses remain substantial even in a relatively flexible labor market, and Seim (2019) emphasizes the role of worker reallocation across firms and places after displacement. More recently, Athey et al. (2023) show that displacement effects are highly heterogeneous across workers, establishments,

[1]See also Fogel et al. (2024) on market-definition bias in studies of labor market power.

and markets. We contribute by providing a structural measure of job substitutability that helps explain this heterogeneity: workers suffer the largest and most persistent losses when the lost job had few close substitutes.

The rest of the paper proceeds as follows. Section 2 develops the model and defines the diversion-based measures of job substitutability and labor overlap. Section 3 describes the Swedish administrative data and the estimation strategy. The remaining sections present the estimated substitution patterns and validate them using plant closures.

2 Model

2.1 Employee Diversion

Our central object is the **job removal diversion ratio**. In product-market merger analysis, diversion ratios are used to identify the closest competitors: if product j becomes unavailable, where do its consumers go instead? Products that capture a large share of that lost demand are close substitutes, so merging them is more likely to weaken competition. We study the labor-market counterpart. If job j disappears, where do its workers go?

In our setting, a job is an occupation-establishment pair. Many workers may hold the same job. For a pair of jobs (j, k), the diversion ratio measures the fraction of workers at job j who would move to job k if job j ceased to be available:

$$D_{jk} = \frac{\#\text{ workers moving from } j \text{ to } k \text{ if } j \text{ is removed}}{\#\text{ workers at } j} \tag{1}$$

This answers the question: *If job j disappears, where do its workers go?*[2]

This object is directional. A job may be an important fallback for workers leaving another job even if the reverse is not true. The full diversion matrix therefore summarizes worker substitution patterns without requiring us to impose occupation, industry, or geographic market boundaries in advance.

Observed worker flows need not recover this object (Qiu et al., 2024). Raw transitions combine workers who move because a job disappears with workers who move voluntarily for higher wages, better amenities, or personal reasons. These different motives imply different substitution patterns, so realized flows need not equal the counterfactual reallocations after job loss. We therefore estimate diversion ratios from a structural model of labor supply.

[2] As in product-market applications, diversion is a partial-equilibrium substitution object. We ask where workers from a lost job would go holding fixed the attractiveness of the remaining options. This abstracts from capacity constraints, vacancy creation, and equilibrium wage responses, but it isolates the substitution patterns that determine competitive pressure across jobs.

2.2 Labor Supply

We model labor supply in period t as a discrete choice problem in which workers choose among differentiated jobs and unemployment, taking their previous-period job $j_{i,t-1}$ as given. Each worker may choose any job in the labor market that is feasible given the worker's education.[3] We do not impose geographic market boundaries or require observed vacancies; geography enters utility only through distance-related disutility. Consider a labor market with N workers and J jobs. We suppress the time subscript on current-period objects for clarity.

Worker i's indirect utility from job j is

$$u_{ij} = \delta_j + \mu_{ij} + \varepsilon_{ij} \tag{2}$$

where δ_j is mean utility common across workers, μ_{ij} captures systematic heterogeneity, and ε_{ij} is an idiosyncratic Type I extreme value shock. The mean utility δ_j captures wages as well as observed and unobserved non-wage amenities. We write the heterogeneity term as

$$\mu_{ij} = g_d(d_{ij}^{\text{commute}}, X_i) + g_f(d_{ij}^{\text{family}}, X_i) + g_s(X_i)\mathbf{1}\{j \neq j_{i,t-1}\} + g_e(E_{ij}, X_i) + g_o(O_{ij}, X_i) \tag{3}$$

where X_i collects worker demographics, d_{ij}^{commute} is commuting distance, d_{ij}^{family} is the distance between the job and the worker's family location, E_{ij} contains indicators for over-qualification, partial subject match, and no subject match, and O_{ij} is the similarity between the worker's previous occupation and the occupation associated with job j. For workers whose previous-period state is unemployment, we set this term to zero for all jobs. The functions $g_d(\cdot)$ and $g_f(\cdot)$ are flexible distance splines, and all components are allowed to vary with demographics such as gender, age, education, family status, and foreign-born status. The outside option is unemployment, and its mean utility is normalized to zero.

Each worker's choice set $\mathcal{J}_i \subseteq \mathcal{J}$ is restricted to jobs that are feasible given the worker's education, as described in Section 3.2. Under the Type I extreme value assumption on ε_{ij}, the choice probability takes the logit form:

$$s_{ij} = \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{k \in \mathcal{J}_i} \exp(\delta_k + \mu_{ik})} \tag{4}$$

where the outside option represents unemployment.

[3]The set of jobs consists of all pairs observed in the period, that is, all occupation-establishment combinations with at least one worker employed.

Aggregating over workers, labor supply to job j is

$$L_j = N \times s_j = N \times \int_i s_{ij}\, dF(i) \tag{5}$$

where s_j is the market share of job j.

2.3 Diversion Ratios in the Model

The labor supply model gives us a direct way to compute the diversion ratio in (1). The idea is a simple counterfactual: remove job j from the choice set, recompute each worker's predicted employment shares, and measure how much each alternative k expands. Formally,

$$D_{jk} = \frac{1}{s_j}\left(s_k(\mathcal{J} \setminus j) - s_k(\mathcal{J})\right) \tag{6}$$

where $s_k(\mathcal{J})$ is job k's share when all jobs are available and $s_k(\mathcal{J} \setminus j)$ is its share when job j is removed. Diversion from j to k is therefore the increase in job k's employment share after removing job j, normalized by the size of the lost job. Jobs that absorb a large share of the displaced workers are close substitutes; jobs that absorb little are not.

Under the logit choice probabilities in (4), this counterfactual share comparison can be written in closed form as

$$D_{jk} = \frac{1}{s_j}\int_i \frac{s_{ij}\, s_{ik}}{1 - s_{ij}}\, dF(i) \tag{7}$$

This expression integrates over all workers. Each worker's contribution depends on how attached they are to the lost job (s_{ij}), how attractive the alternative is (s_{ik}), and how concentrated their choices are ($1 - s_{ij}$). Workers with strong attachment to job j and a high probability of choosing k contribute the most to diversion from j to k.[4]

2.4 Aggregating Diversion Ratios

We ultimately want diversion-based overlap not only for individual jobs, but also for collections of jobs such as firms, occupations, or regions. For an origin set of jobs A and a destination set of jobs B, the aggregate diversion ratio is

$$D_{AB} = \frac{1}{\sum_{j \in A} s_j} \sum_{j \in A} \sum_{k \in B} s_j\, D_{jk} \tag{8}$$

[4]The expression follows from the individual-level diversion ratio $D_{jk,i} = s_{ik}/(1 - s_{ij})$, which is the conditional probability that worker i chooses job k when job j is unavailable (Conlon and Mortimer, 2021). Under the logit, this individual diversion ratio is the same whether job j is removed entirely or merely made less attractive.

This is the employment-weighted average of job-level diversion ratios from jobs in A to jobs in B, removing one job at a time. It measures the share of workers at A who would reallocate to B if their particular job disappeared.

When A and B are occupations or commuting zones, this formula delivers diversion-based overlap at those levels of aggregation. When A and B are the sets of jobs owned by two firms, it gives an employer-level measure of labor overlap: the fraction of one firm's workers who would move to the other firm if their job were eliminated. This is the object relevant for assessing how directly two employers compete for labor.

2.5 Job Substitutability Index

We summarize job-level substitutability with a concentration measure:

$$\text{Sub}_j = 1 - \sum_k (D_{jk})^2 \tag{9}$$

This index ranges from 0 to 1. When diversion is dispersed across many alternatives, Sub_j approaches 1. When diversion is concentrated in a few alternatives, Sub_j approaches 0. The substitutability index therefore provides a scalar summary of how easily workers at job j can be reallocated across the rest of the labor market.

3 Data

We use Swedish administrative registers from Statistics Sweden (SCB). The core data come from matched employer-employee records (JOBB and SYSSREG), which provide the universe of employment relationships in Sweden including earnings, occupation codes, industry classification, and establishment identifiers. We link these to individual records in LISA, which contains individual demographics including age, family composition, and residential location for all individuals aged 16–70, and importantly education level and subject. We also obtain workplace and home coordinates for all individuals, as well as family links, allowing us to measure both commuting costs and family-location ties.[5] Wage structure statistics (Lönestrukturstatistiken) provide full-time adjusted monthly wages for a majority of workers, and firm financial data (FEK) provide value added and other financial variables at the establishment level.

[5]For workers whose mother's location is missing, we impute it with the worker's own residence, so the family-distance measure coincides with commuting distance. Given the high quality of the Swedish register data, we interpret these cases as likely reflecting mothers who do not reside in Sweden.

3.1 Sample Definition

Our main estimation sample is a two-period panel of individuals aged 24–54 who are observed in both 2009 and 2010 and have non-missing education and demographic information. We exclude military personnel, politicians, police, and managers.[6] We also exclude individuals whose current-period or previous-period job is infeasible given their education, as discussed below.

A job is defined as an establishment-occupation combination observed in 2010, where occupations are classified at the 4-digit level using the SSYK96 classification.[7] The final estimation sample contains roughly 2.6 million individuals. Of these, about 2.20 million are employed in 2010 and distributed across approximately 615,000 jobs, while roughly 400,000 choose the outside option.

Two auxiliary data sources capture the structure of job substitutability: an education-occupation matrix from Statistics Sweden and occupation similarity scores from the Swedish Public Employment Service.

3.2 Education-Occupation Match

Statistics Sweden provides a classification matrix assessing the fit between education programs and occupations. The matrix combines normative assessments by education experts with statistical analysis of actual employment patterns. We map this to our data, yielding 93 education groups and 324 occupation codes.[8]

Each education-occupation pair is evaluated on two dimensions, illustrated in Appendix Figure A.1.[9] Level match captures whether the worker is under-qualified (education level below job requirements), correctly qualified, or over-qualified. Subject match captures how well the education's field of study aligns with the occupation's skill requirements: complete match (the education directly prepares for the occupation), partial match, or no match.

We use this matrix to restrict each worker's choice set to feasible occupations by excluding pairs where the worker is under-qualified.[10] This includes licensed occupations such as physicians, nurses, and teachers, that require a complete subject match. Hence,

[6] We exclude managers as this represents more a career progression than a job choice.

[7] There are 324 4-digit occupations in our matched data. A complete list can be found here: `https://www.scb.se/ssyk`

[8] A complete list of the education groups can be found in the Appendix or here:`https://www.scb.se/dokumentation/klassifikationer-och-standarder/svensk-utbildningsnomenklatur-sun/`

[9] See Tables A.1 and A.2 in the Appendix for the group definitions.

[10] Agarwal and Somaini (2025) develop the econometric framework for demand estimation under latent choice constraints. In our setting, the education-occupation matrix provides directly observed constraints, so we do not need to estimate the choice sets.

workers only choose among jobs they could plausibly obtain.[11] The screen is substantial but far from binding: the median worker retains 236 feasible 4-digit occupations, and the 90th percentile retains 289. Consider a dentist and a dental nurse. For the dentist occupation, only university-level dentistry education provides a complete subject match. Dental care education also provides the right subject but leaves the worker under-qualified.[12] For the dental nurse occupation, university-level dentistry education gives a complete subject match, but the worker will be overqualified. With dental care education a worker has the complete subject match and is correctly qualified for the job. This asymmetry captures an important labor market reality: dentists can move into dental nursing roles, but dental nurses cannot become dentists.

3.3 Occupation-Occupation Similarity

The Swedish Public Employment Service (Arbetsförmedlingen) provides occupation similarity scores derived from job advertisements. These scores are used in practice to guide unemployed workers toward alternative occupations. The underlying data consists of over 6 million job postings from 2016–2024, from which skills and competencies are extracted. Cosine similarity between occupation pairs is computed based on these skill profiles.[13] For our purposes, the key feature is that the measure is constructed from the text of job advertisements rather than from observed worker flows, so it captures task similarity without mechanically importing the mobility patterns we later seek to explain. The resulting scores range from 0 to 1, with higher values indicating greater skill overlap. We use the non-directional version of this measure and retain the full set of pairwise similarity scores. Directionality in feasible transitions is then introduced separately through the education-occupation feasibility restriction, which rules out moves barred by qualifications or licensing. The dentist and dental nurse example illustrates this separation: the task-based similarity between the two occupations can be high, but the education screen rules out moves into the licensed dentist occupation for workers without the required training.

Appendix Figure A.2 visualizes the similarity structure as a network where nodes represent occupations and edges connect each occupation to its ten most similar counterparts for readability. The layout positions similar occupations closer together, revealing skill-based clusters. Healthcare occupations form a tight cluster, as do IT and technology occupations.

[11] For workers with only compulsory education, who are not explicitly covered in the matrix, we assign feasibility to all occupations for which any education level is either qualified or over-qualified.

[12] The dentist occupation is licensed, so workers must hold university-level dentistry education to practice.

[13] The original data uses the SSYK 2012 classification; we transform it to SSYK96 to match our administrative records.

3.4 Summary Statistics

Table 1 presents summary statistics. About half of workers have a complete subject match between their education and occupation, while the other half work in fields unrelated to their training. Among employed observations, we require current and lagged jobs to be feasible given the worker's education, which removes roughly 300,000 employed observations. Hence, there are no under-qualified employed workers in the estimation sample. Interestingly, nearly all employed workers are correctly qualified for their current job, though many work in occupations not fitting their educational background. Panel C also shows that the estimation sample includes roughly 424,000 non-employed individuals, who enter the model as the outside option. The median commute is 6.8 kilometers and the median distance to the mother is 13.6 kilometers, consistent with non-negligible family-location ties. About one quarter of workers switch jobs year-to-year. Workers who switch jobs move, on average, to occupations with high skill overlap with their previous occupation.

Table 1: Summary Statistics

	Value
Panel A: Worker Characteristics	
Female	51.5
Foreign-born	18.2
Has child under 3	17.5
Number of children (0–17)	0.83
Lives in big city	20.0
Age 24–29	19.6
Age 30–39	30.6
Age 40–49	32.1
Age 50–54	17.7
Pre-secondary education (<9 years)	2.4
Pre-secondary education (9 years)	10.2
Secondary education (≤2 years)	49.3
Secondary education (3 years)	6.4
Post-secondary education (<3 years)	30.6
Post-secondary education (≥3 years)	1.2
Panel B: Matching and Mobility	
Share switching jobs	23.5
Median distance to job (km)	6.8
Median distance to mother (km)	13.6
Share under-qualified (by construction)	0.0
Share correctly qualified	96.5
Share over-qualified	3.5
Share with complete subject match	45.3
Share with partial subject match	4.5
Share with poor subject match	50.2
Mean occ. similarity (switchers)	0.60
Panel C: Sample Overview	
Total individuals	2,620,754
Employed individuals	2,196,707
Unemployed individuals	424,047
Jobs	615,110
Average job size	3.6

Notes: Sample of Swedish individuals aged 24–54 observed in both 2009 and 2010. Panel A reports demographic shares. Panel B reports matching quality among employed individuals based on the education-occupation and occupation-occupation matrices, as well as mobility patterns between 2009 and 2010. The mean occupation-similarity measure is computed among all job switchers, not only occupation switchers. Panel C reports post-restriction sample sizes. A job is defined as an establishment-occupation combination observed in 2010.

4 Estimation

4.1 Computational Challenge

Estimating a discrete choice model with more than 2.6 million workers and 615,000 job alternatives presents a computational challenge. The full choice set would require evaluating choice probabilities for $\approx 1.6 \times 10^{12}$ worker-job pairs.

We address this challenge through two complementary techniques: sampling of alternatives and exploiting the multinomial-Poisson equivalence.

4.2 Sampling of Alternatives

Following Train et al. (1987) and Frejinger et al. (2009), we estimate on a subset of alternatives for each worker. For each worker i, we draw $m = 20$ jobs uniformly at random from the worker's feasible choice set $\mathcal{J}_i$ (see Section 3.2), and then add the chosen alternative, the previous-period job, and the outside option. Duplicates are removed, so the sampled set contains 21 to 23 alternatives. If the previous-period job is not observed among current-period jobs, it cannot be added; in that case the state-dependence term equals one for all jobs except the outside option.

This reduces the computational burden while maintaining consistent parameter estimates through a sampling correction. Each alternative j in the sampled set K receives a weight

$$q_{K|j} = P(K \text{ selected} \mid j \text{ selected}) \propto \frac{1}{P(\text{sampling } j)} \tag{10}$$

Under uniform sampling, randomly drawn jobs receive weight $|\mathcal{J}_i|/m$. The chosen alternative, previous-period job, and outside option, which are always included when distinct and available, receive weight 1.

The corrected choice probability is

$$P_{ij} \mid K = \frac{\exp(\delta_j + \mu_{ij} + \ln(q_{K|j}))}{\sum_{k \in K} \exp(\delta_k + \mu_{ik} + \ln(q_{K|k}))} \tag{11}$$

This correction ensures that coefficient estimates are consistent with estimation on the full choice set.

4.3 Poisson Pseudo-Maximum Likelihood

Sampling of alternatives makes the stacked worker-alternative dataset manageable, but it does not solve the second computational bottleneck: estimation still requires a very large number of job fixed effects. In our application, a classical conditional logit implementation

with roughly 615,000 job effects is not computationally feasible. We therefore exploit the equivalence between the multinomial logit likelihood and the log-linear Poisson model (Baker, 1994), which allows us to use fast high-dimensional fixed-effects Poisson estimators.

For each worker i and sampled alternative $j \in K_i$, let $y_{ij} = 1$ if j is chosen and zero otherwise. The conditional logit probabilities in the sampled set imply the Poisson mean

$$\mathbb{E}[y_{ij} \mid K_i] = \exp\left(\eta_i + \delta_j + x'_{ij}\beta + \ln q_{K_i|j}\right) \tag{12}$$

where η_i is a worker fixed effect absorbing the denominator of the conditional choice probability, δ_j is the job fixed effect, x_{ij} collects the worker-job characteristics in μ_{ij}, and $\ln q_{K_i|j}$ is the sampling correction. We estimate this equation using `ppmlhdfe` in Stata (Correia et al., 2020), absorbing worker and job fixed effects. The remaining coefficients identify preferences over distance, match quality, and state dependence, while the absorbed job fixed effects are recovered after estimation and used to construct predicted choice probabilities and diversion ratios.

4.4 Specification

Our main specification includes job fixed effects, commuting-distance splines over 0–5, 5–20, 20–50, and 50+ km, family-distance splines over 0–5, 5–50, and 50+ km, a switching indicator, occupation similarity, and indicators for over-qualification, partial subject match, and poor subject match. We then allow heterogeneity in skill match, short- and medium-distance commuting costs, and switching costs by worker demographics. The main interactions are with gender, low and high education, age and age squared, foreign-born status, big-city residence, the presence of a young child, and the number of children. Here, low education denotes high-school education, medium education denotes university education and is the omitted category, and high education denotes postgraduate education.

5 Results

5.1 Worker Preferences

Table 2 presents the main preference estimates from the worker choice model. Panel A reports the base coefficients. Distance enters through spline segments over 0–5, 5–20, 20–50, and 50+ km, while distance to mother enters through spline segments over 0–5, 5–50, and 50+ km. Panel B reports heterogeneity in skill match. Panels C and D report heterogeneity in the distance slopes for short and medium commutes. Panel E reports heterogeneity in switching costs.

Table 2: Estimation Results

Variable	Coeff.	S.E.
Panel A: Base Coefficients		
Distance (0–5 km)	-0.060^{**}	(0.026)
Distance (5–20 km)	-0.065^{***}	(0.00065)
Distance (20–50 km)	0.086^{***}	(0.0034)
Distance (50+ km)	-0.0049^{***}	(0.000039)
Dist. to mother (0–5 km)	-0.063^{***}	(0.0029)
Dist. to mother (5–50 km)	-0.0077^{***}	(0.00025)
Dist. to mother (50+ km)	-0.00068^{***}	(0.000029)
Switching cost	-1.433^{***}	(0.105)
Skill match	1.106^{***}	(0.014)
Educ. level too high	-0.690^{***}	(0.011)
Partial subj. match	-0.428^{***}	(0.011)
Poor subj. match	-0.834^{***}	(0.0071)
Panel B: Skill match × Demo.		
Female	-0.049^{***}	(0.0016)
Low educ.	0.032^{***}	(0.0025)
High educ.	-0.037^{***}	(0.0018)
Age	-0.0023^{***}	(0.00072)
Age^2	0.000079^{***}	(0.000009)
Panel C: Dist. (0–5 km) × Demo.		
Female	-0.012^{***}	(0.0027)
Age	0.030^{***}	(0.0013)
Age^2	-0.00039^{***}	(0.000017)
Low educ.	-0.119^{***}	(0.0043)
High educ.	-0.015^{***}	(0.0031)
Big city	-0.262^{***}	(0.0019)
Young child	-0.026^{***}	(0.0026)
Num. children	0.038^{***}	(0.00097)
Foreign born	-0.243^{***}	(0.0026)
Skill match	-0.035^{***}	(0.0011)
Panel D: Dist. (20–50 km) × Demo.		
Female	-0.015^{***}	(0.00038)
Age	-0.0040^{***}	(0.00019)
Age^2	0.000045^{***}	(0.000002)
Low educ.	0.0042^{***}	(0.00060)
High educ.	0.019^{***}	(0.00043)
Big city	0.031^{***}	(0.00049)
Young child	-0.0053^{***}	(0.00056)
Num. children	-0.0054^{***}	(0.00022)
Foreign born	0.014^{***}	(0.00052)
Skill match	-0.0034^{***}	(0.000068)
Panel E: Switching cost × Demo.		
Female	-0.102^{***}	(0.011)
Age	-0.099^{***}	(0.0052)
Age^2	0.00084^{***}	(0.000068)
Low educ.	-0.241^{***}	(0.018)
High educ.	0.280^{***}	(0.013)
Foreign born	-0.560^{***}	(0.012)
Distance (0–5 km)	-0.668^{***}	(0.0089)
Distance (20–50 km)	-0.135^{***}	(0.00062)
Constant	2.337^{***}	(0.042)

Notes: Poisson pseudo-maximum likelihood estimates with worker and job fixed effects. Panel A reports base preference coefficients. Panel B reports interactions of skill match with worker demographics. Panel C reports interactions of distance (0–5 km) with worker demographics. Panel D reports interactions of distance (20–50 km) with worker demographics. Panel E reports interactions of switching cost with worker demographics. Low education denotes high-school education, medium education denotes university education and is the omitted category, and high education denotes postgraduate education. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Commuting. The distance coefficients imply substantial commuting frictions, but in a clearly nonlinear way. Utility falls with distance in the first 20 km, while the positive coefficient on the 20–50 km spline indicates that the marginal penalty flattens over medium distances rather than continuing linearly. This does not imply that workers prefer medium commutes: it means only that the slope becomes less negative once commutes are already fairly long. Beyond 50 km, the remaining slope is close to zero. Distance to mother is negative in every segment, with the largest penalty at very short distances and much smaller effects further away, consistent with family proximity mattering most locally. The interaction patterns show that short-distance commuting costs are especially steep for workers in big cities and for foreign-born workers, and are also stronger for women, lower-educated workers, and workers with young children.

Switching cost. The large negative switching-cost coefficient implies substantial inertia in job choice: conditional on job fixed effects and the included worker-job characteristics, workers place considerable value on remaining in their current job rather than moving. The magnitude is economically large: in absolute value, the switching penalty exceeds the gain from moving from zero to perfect occupation similarity. This switching cost is stronger for women, lower-educated workers, and especially foreign-born workers, but weaker for highly educated workers. The negative interaction with the 0–5 km distance bin further implies that switching frictions remain especially strong even when alternative jobs are very nearby.

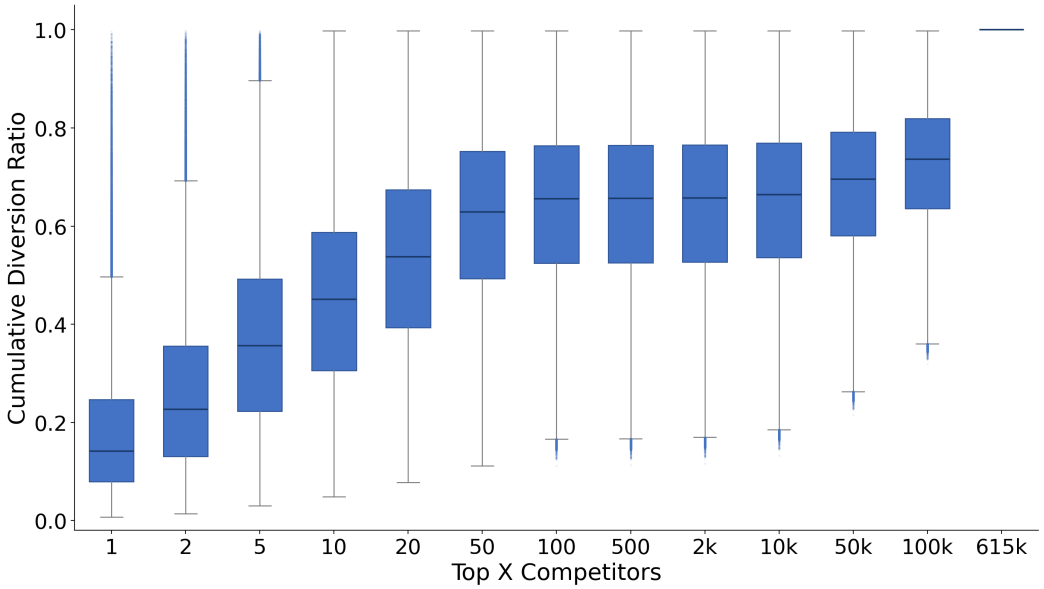
Education and skill fit. Workers strongly value occupational continuity: the coefficient on skill match is large and positive. In a similar vein, they dislike mismatch in formal training. Jobs for which workers are overqualified carry a large negative coefficient, and subject mismatch is also costly, especially when the mismatch is poor rather than partial. The poor-subject-match penalty is nearly as large as the overqualification penalty, which suggests that formal training mismatch is not a second-order margin. Panel B shows that skill match matters somewhat less for women and highly educated workers, and more for lower-educated workers. Taken together, the estimates suggest that workers trade off geography, switching costs, and match quality in systematic ways, rather than ranking jobs on any single dimension alone.

5.2 Diversion Concentration Across Alternatives

Figure 1 summarizes how much of a job's diversion is captured by its top outside alternatives. The x-axis expands approximately logarithmically, so moving to the right corresponds to a rapid increase in the number of outside jobs included.

The figure highlights two features of the estimated diversion patterns. First, the aver-

Figure 1: Concentration of Diversion in Top Alternatives



Notes: For each origin job j, the figure plots the cumulative diversion share to its top-X alternatives, $\sum_{k=1}^{X} D_{j,k^{(r)}}$, where $k^{(r)}$ denotes the r-th ranked alternative by diversion ratio. The x-axis expands approximately logarithmically in X. Higher values indicate that most displaced workers move to a small number of alternatives.

age job faces a large number of relevant competitors. The median job sends only about 15% of displaced workers to its top alternative, around 45% to its top 10 alternatives, and only about two thirds to its top 100 alternatives. Second, there is substantial heterogeneity across jobs. While many jobs have diffuse diversion, others have extremely concentrated outside options, with a small number of destinations capturing most of the diversion mass. These are precisely the cases in which antitrust authorities should be especially cautious: when a merger affects jobs with concentrated alternatives, workers have fewer close substitutes and the scope for post-merger wage pressure is greater.

The same heterogeneity should matter for job loss more generally. If a job disappears because an establishment closes, workers from jobs with diffuse diversion have many plausible reemployment paths and should be able to reallocate more easily. By contrast, workers from jobs with concentrated diversion must compete for a narrow set of alternatives, so congestion or mismatch in those destinations is more likely to translate into larger adjustment costs.

5.3 Geographic and Occupational Variation

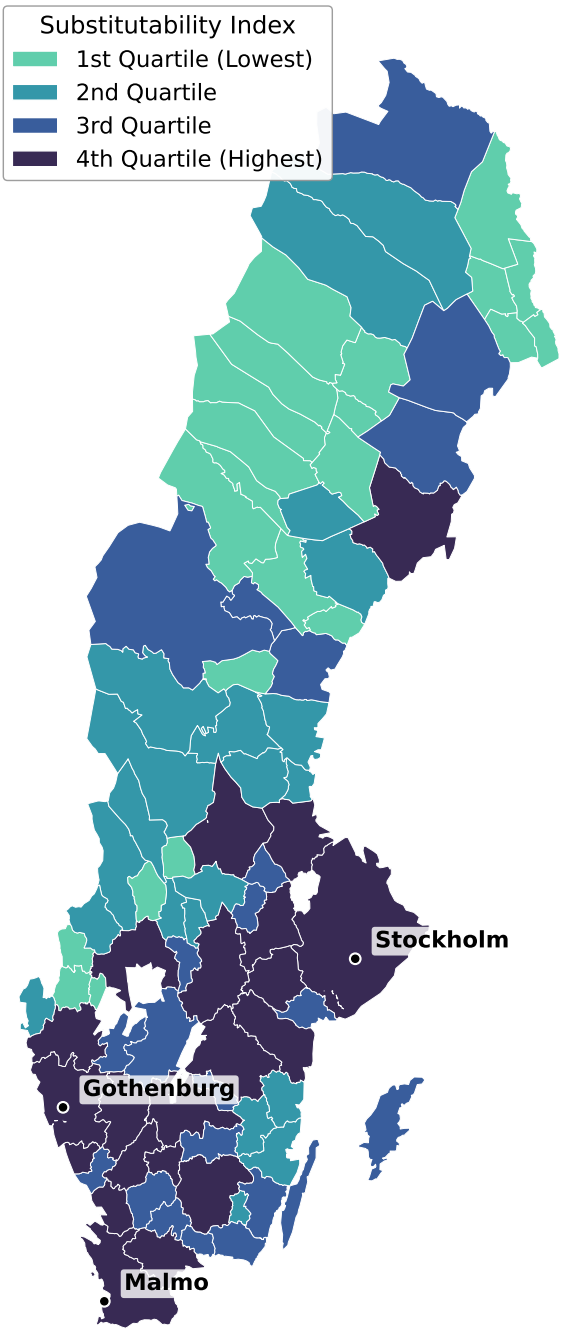
Figure 2 maps the average substitutability index across Sweden's 85 commuting zones. Urban labor markets, Stockholm, Gothenburg, and Malmo, exhibit high substitutability,

consistent with thick local markets offering many nearby alternatives. Rural commuting zones in northern Sweden, by contrast, have substantially lower average indices, reflecting thinner markets where the disappearance of a single employer can leave workers with few comparable options. The differences are significant: the average index in Stockholm (0.97) exceeds that in the smallest commuting zones (around 0.77) by 20 percentage points.

Figure 3 shows how substitutability varies across broad occupation groups, using the same grouping as the similarity network in Appendix Figure A.2. IT and technology, business and finance, and service occupations have the highest average substitutability, consistent with these workers having skills that transfer across many employers. Agriculture, healthcare, and science occupations have the lowest, reflecting that workers in these fields depend on a narrower set of employers: a nurse is typically tied to hospitals and care facilities, while a software developer can work at firms in many different industries. The interquartile ranges show that there is substantial heterogeneity within each group: even among healthcare occupations, some jobs have highly dispersed alternatives while others are concentrated.

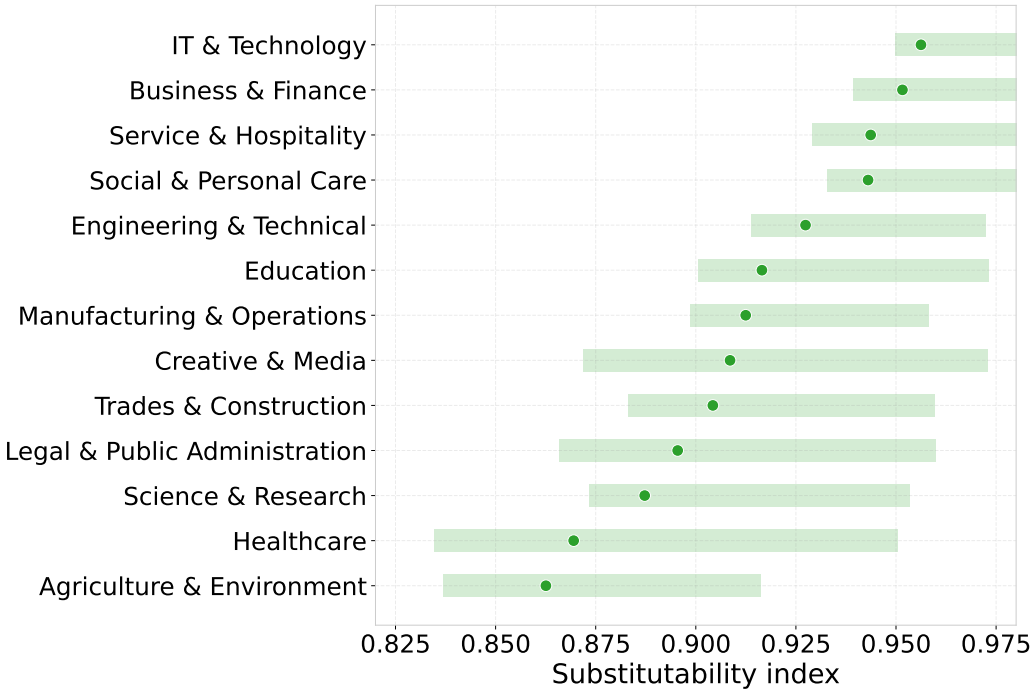
Taken together, these patterns show that the Swedish labor market is highly competitive for most jobs, but with meaningful variation across geography and occupations. The next section studies whether this variation predicts actual displacement costs using plant closures.

Figure 2: Substitutability Across Commuting Zones



Notes: Mean substitutability index by commuting zone (lokala arbetsmarknader, 2004 classification), grouped into quartiles. The index equals $\text{Sub}_j = 1 - \sum_k D_{jk}^2$, where D_{jk} is the diversion ratio from job j to alternative k. Averages are unweighted across jobs within each commuting zone. Darker shading indicates higher substitutability. Dots mark Sweden's three largest cities (English names). Municipality boundaries are dissolved to commuting zones using the SCB 2004 classification.

Figure 3: Substitutability by Occupation Group



Notes: Mean substitutability index (dots) and interquartile range (shaded bars) by broad occupation group. Occupation groups follow the SSYK96 classification used in Appendix Figure A.2. The index equals $\text{Sub}_j = 1 - \sum_k D_{jk}^2$. Averages are weighted by the number of jobs in each four-digit occupation within each group.

6 Do Diversion Ratios Predict Displacement Costs?

We validate our structural estimates using plant closures. When an establishment closes, its workers are forced to reveal their next-best alternatives, paralleling the use of second-choice data in product market demand estimation. This setting is also directly relevant for the merger application that motivates the paper: roughly 13 percent of establishments involved in mergers in Sweden subsequently close, affecting approximately 12,000 workers per year (Appendix Figure A.3). When a merger eliminates a job, the question of where displaced workers can go is exactly the one our diversion framework addresses.

Following Athey et al. (2023), we identify establishments that cease to exist between November 2010 and November 2011 and track the subsequent labor market outcomes of their workers. Our structural model is estimated without using this plant closure data; the comparisons below are therefore fully out of sample. We use the closure sample in two complementary ways. First, we compare realized reallocation patterns after closure to model-implied diversion ratios across commuting zones and occupations. Second, we ask whether our scalar substitutability index predicts which displaced workers suffer larger and more persistent earnings losses after their job disappears.

6.1 Validation Design and Interpretation

Our structural diversion ratios are partial equilibrium objects: they measure where a worker would go if a single job disappeared, holding all other employment options fixed. Observed transitions after plant closures, by contrast, reflect general equilibrium outcomes in which many workers are displaced simultaneously and compete for the same alternatives. Not all displaced workers can obtain their most preferred alternative if capacity is limited, and congestion may push some workers toward less preferred options. Despite this distinction, the substitutability index should still predict relative outcomes across workers: those with more dispersed alternatives face a lower probability that all of their next-best options are congested, and should therefore fare better on average even in a general equilibrium setting.

6.2 Diversion Across Commuting Zones

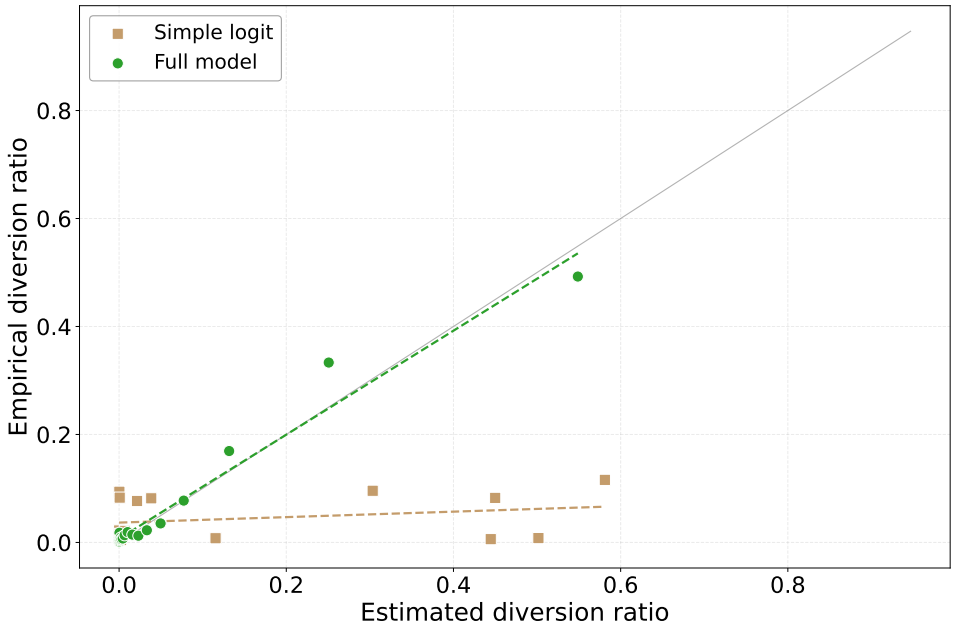
We first compare model-implied diversion ratios to empirical transition patterns from the plant closure sample, aggregated across reemployment destinations at the commuting-zone and occupation levels and excluding transitions into unemployment.

Appendix Figure A.6 displays diversion matrices across commuting zones for plant closures and our model. Both matrices share a strong diagonal, indicating that displaced workers typically re-match locally. The full model also reproduces the main off-diagonal

patterns in the data, whereas the simple logit benchmark in Appendix Figure A.8 is much flatter.

Figure 4 summarizes the same comparison in binned form. The simple logit benchmark is essentially off the mark, but this is not surprising: when workers choose jobs only through mean utility, diversion is mechanically proportional to destination job shares. As a result, the simple logit model sends too much displaced labor toward large destinations and produces a relationship that is far too flat. The full model, by contrast, is much closer to the 45-degree line and captures the high-diversion bins that matter most for local labor market substitution.

Figure 4: Model Fit: Diversion Across Commuting Zones



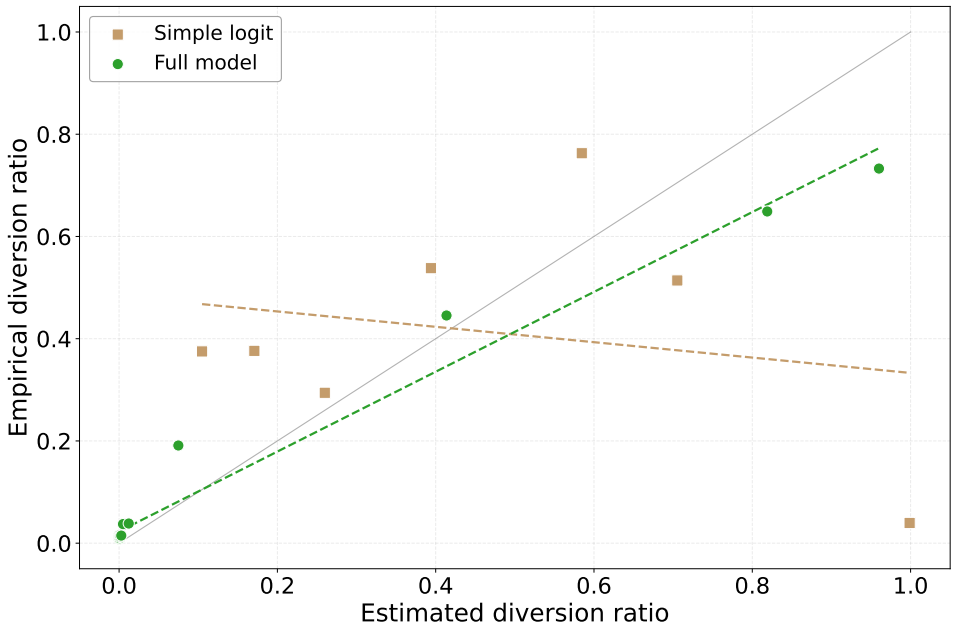
Notes: Each marker represents a bin of origin-destination commuting zone pairs, with the x-axis showing model-estimated diversion ratios and the y-axis showing empirical reallocation shares observed after plant closures. The empirical shares are computed from the plant closure sample described in Section 6.1, excluding transitions into unemployment. The full model uses the estimated preference parameters from Table 2. The simple logit benchmark uses only job fixed effects without worker-job interactions, so that diversion is proportional to destination job shares. The thin line indicates perfect fit. See Appendix Figures A.6 and A.8 for the underlying diversion matrices.

6.3 Diversion Across Occupations

Appendix Figure A.7 provides the analogous comparison at the 4-digit occupation level. The empirical matrix again shows a strong diagonal and broad block structure by skill groups. The full model captures both features reasonably well, while the simple logit benchmark in Appendix Figure A.9 misses much of this structure.

The same pattern appears in Figure 5. The full model tracks the steep empirical relationship between observed and predicted occupational diversion much more closely, while the simple logit benchmark again remains too flat because it largely reallocates workers in proportion to destination shares rather than to worker-specific substitution patterns.

Figure 5: Model Fit: Diversion Across Occupations



Notes: Each marker represents a bin of origin-destination four-digit occupation (SSYK96) pairs, with the x-axis showing model-estimated diversion ratios and the y-axis showing empirical reallocation shares observed after plant closures. The empirical shares are computed from the plant closure sample described in Section 6.1, excluding transitions into unemployment. The full model uses the estimated preference parameters from Table 2. The simple logit benchmark uses only job fixed effects without worker-job interactions, so that diversion is proportional to destination job shares. The thin line indicates perfect fit. See Appendix Figures A.7 and A.9 for the underlying diversion matrices.

6.4 Aggregate Diversion Shares

The close correspondence between estimated and empirical diversion matrices at both geographic and occupational levels provides direct evidence that the structural model recovers realistic substitution patterns. This is out of sample: estimation does not use plant closure transitions.

Table 3 makes the comparison more transparent by aggregating both objects to four summary margins. The Plant Closures column reports observed transition shares for workers displaced by jobs that actually close between November 2010 and November 2011. The Model column aggregates the corresponding job-level diversion ratios over those same

closing jobs using the employment weights in (8). The first three rows summarize reallocation across commuting zones and occupations, while the final row reports transitions into unemployment.

The estimated diversion ratios closely track the closure data along all four margins: the model predicts that 83% of displaced workers remain in the same commuting zone, compared to 81% in the closure sample, and that 71% remain in the same occupation, compared to 70% empirically. The joint margin is slightly higher in the model than in the data, but the overall fit remains close.

Table 3: Aggregate Diversion Shares

	Plant Closures	Model
Switchers staying in same CZ	0.81	0.83
Switchers staying in same occupation	0.70	0.71
Switchers staying in same CZ and occupation	0.60	0.64
Switchers becoming unemployed	0.04	0.04

Notes: Plant Closures reports observed transition shares among workers displaced by jobs that close between November 2010 and November 2011. Model reports model-implied diversion shares aggregated over those same closing jobs using employment weights. The first three rows summarize reallocation across commuting zones and occupations; the last row reports transitions into unemployment.

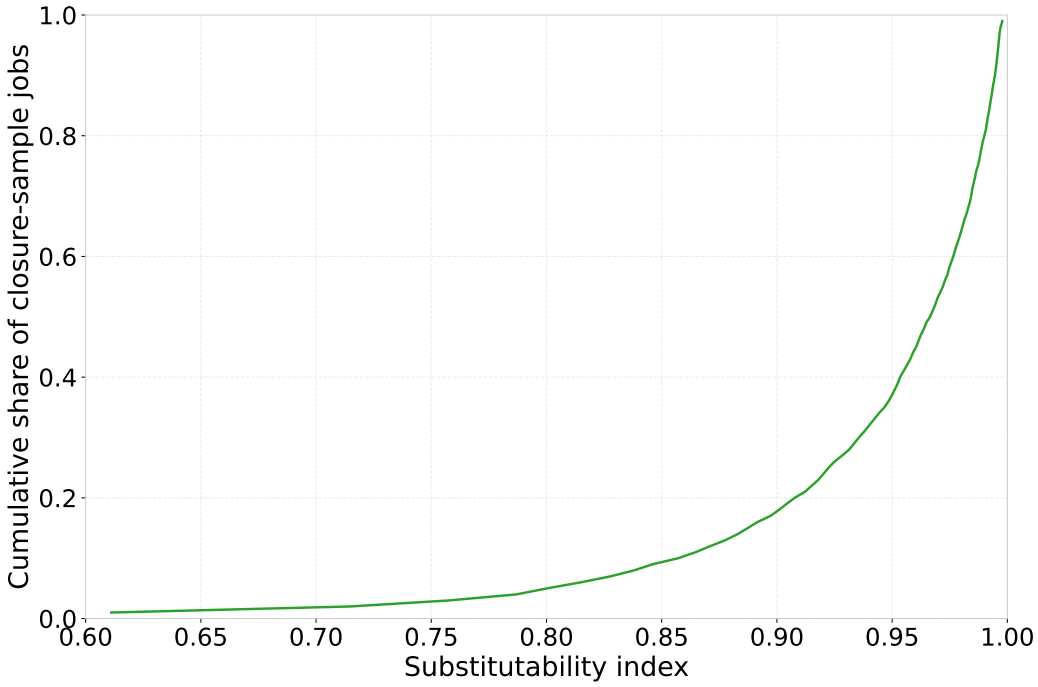
6.5 Earnings Recovery by Substitutability

We next use the closure sample for a second validation exercise. The question is not whether the substitutability index identifies the causal effect of labor market power on earnings. Rather, it is whether an index constructed entirely from the pre-closure model predicts which displaced workers adjust more successfully after their job disappears. We define displaced workers as employees aged 24–60 with at least three years of tenure at establishments that closed between November 2010 and November 2011.[14] For each displaced worker, we compute the substitutability index of the lost job. We construct a control group by drawing approximately three times as many workers from surviving establishments, subject to the same demographic and tenure restrictions. We follow both groups from $t = -5$ to $t = +7$ relative to closure and normalize earnings to $t = -3$. We then estimate an event-study specification comparing workers in the bottom versus top quartile of substitutability, using the control group to difference out common time effects.

[14]We classify an establishment as closed if it disappeared entirely or shrank by at least 90%, excluding cases consistent with relabeling or internal reorganization. We further restrict the displaced sample to workers with earnings above the 10th percentile, excluding military and senior government officials. See Appendix A.6 for full details.

Figure 6 shows the cumulative distribution of the substitutability index among jobs involved in plant closures. The distribution is concentrated near one, but the closure sample still spans a meaningful range of substitutability, allowing us to compare workers who lose jobs with very diffuse alternatives to workers whose reemployment options are much more concentrated.

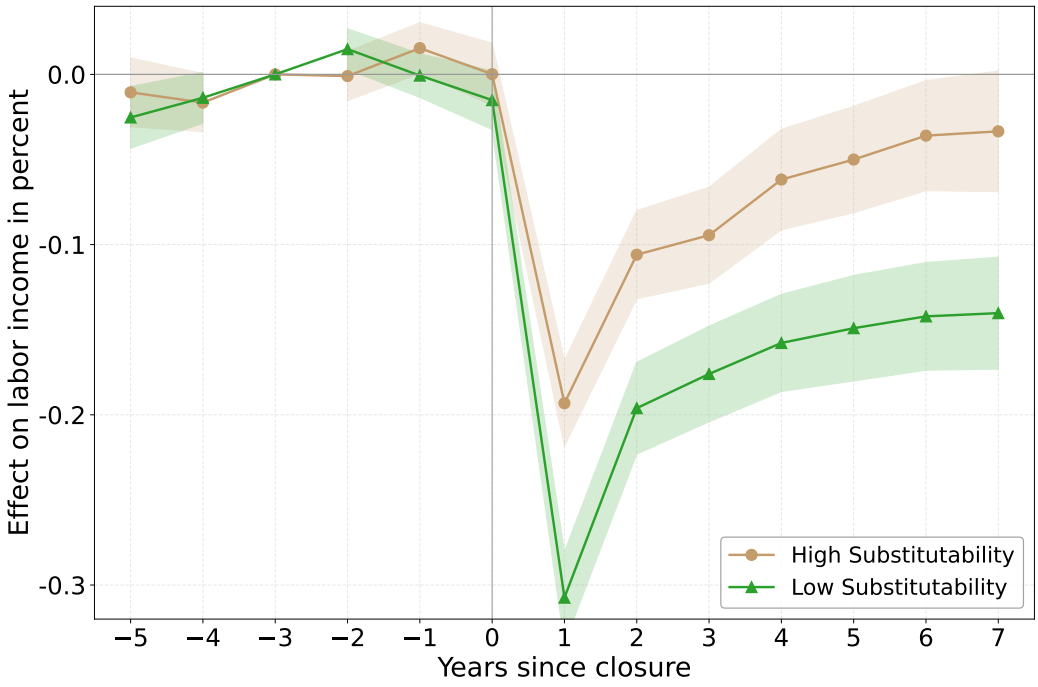
Figure 6: Distribution of Substitutability in the Closure Sample



Notes: Cumulative distribution function of the substitutability index for jobs involved in plant closures (establishments that closed between November 2010 and November 2011). The index equals $\text{Sub}_j = 1 - \sum_k D_{jk}^2$, where D_{jk} is the diversion ratio from job j to alternative k. Higher values indicate more dispersed diversion across alternatives.

Figure 7 presents the results. There is no clear evidence of differential pre-trends between the high- and low-substitutability groups before closure. After displacement, however, the two groups separate sharply. Workers from high-substitutability jobs experience a noticeably smaller initial dip and nearly fully recover by the end of the horizon. Workers from low-substitutability jobs, by contrast, suffer a much deeper drop and remain well below their pre-closure earnings even seven years later. In the first year after closure, workers from low-substitutability jobs lose roughly 30% of pre-closure earnings, compared to about 20% for workers from high-substitutability jobs. Intuitively, a janitor in Stockholm may have many nearby reemployment options after a layoff, whereas a nurse in rural Sweden may face a much narrower set of realistic alternatives.

Figure 7: Earnings Effects of Displacement



Notes: Event study of earnings around plant closures. The y-axis shows the average treatment effect on normalized earnings relative to $t = -3$. Workers split by the substitutability index of their lost job: top quartile (Q4, high substitutability) vs. bottom quartile (Q1, low substitutability). The substitutability index equals $1 - \sum_k D_{jk}^2$, where D_{jk} is the diversion ratio from job j to alternative k. Shaded areas indicate 95% confidence intervals. Control group consists of workers at surviving establishments.

Because we observe both earnings and employment status, we can decompose the total effect into an extensive margin (probability of non-employment) and an intensive margin (earnings conditional on employment). Appendix A.6 reports both event studies. The intensive margin accounts for the bulk of the gap: even conditional on being employed, low-substitutability workers experience persistently larger earnings losses. The extensive margin contributes a meaningful but transitory differential that vanishes by six years after displacement (Appendix Figures A.5 and A.4).

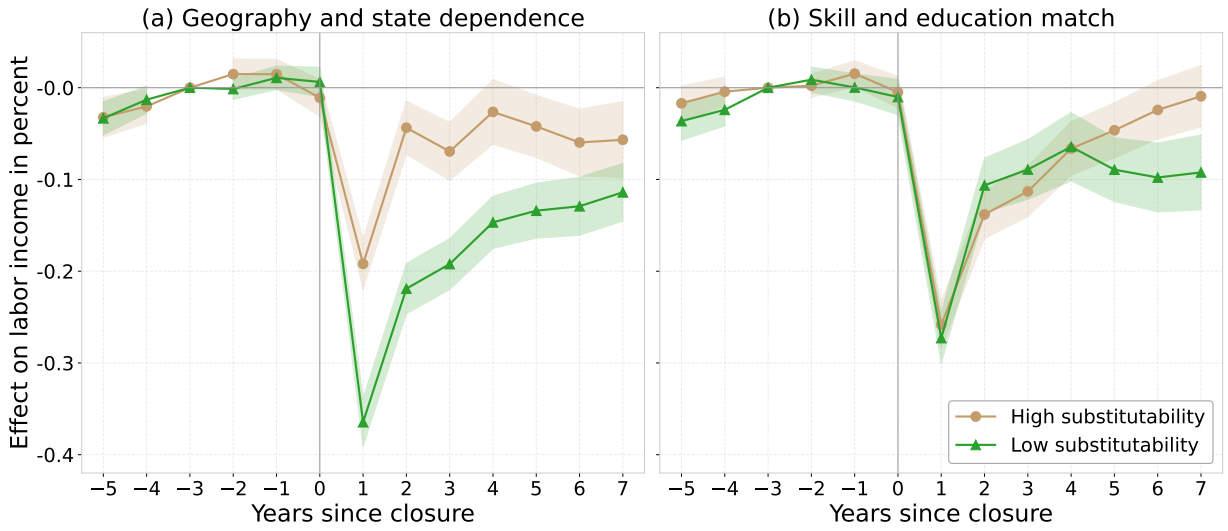
6.6 Which Dimensions of Substitutability Matter?

The substitutability index aggregates multiple dimensions of the systematic utility component μ_{ij} defined in Section 2 into a single scalar. To understand which dimensions drive the displacement patterns documented above, we construct two counterfactual indices. The first retains only the geographic and inertial terms in μ_{ij}: the commuting distance splines, family distance splines, and the state-dependence indicator. The second retains only the skill and match terms: occupation similarity and the education-occupation mismatch indicators (overqualification, subject mismatch). Each counterfactual index measures substitutability as if workers chose jobs based on only one set of factors, holding job fixed effects fixed. We then repeat the quartile-split event study for each index separately.

Figure 8 presents the results. Panel (a) stratifies workers by the geography-only index, and panel (b) by the skill-only index. The two panels reveal a clear division of labor between the two channels. Geographic substitutability governs the initial recovery: workers with many nearby alternatives (high geography-only index) experience a substantially smaller earnings drop in the first year and recover more quickly in years one through three. Skill-based substitutability, by contrast, matters most for long-run scarring. Workers whose occupational profile gives them access to many comparable jobs converge toward their pre-displacement trajectory, while those with concentrated skill-based alternatives remain well below their counterfactual even seven years later. Together, these patterns suggest that geographic frictions determine how quickly displaced workers find new employment, while human capital portability determines how well those new jobs match the old ones.

These results confirm that the substitutability index, constructed from cross-sectional employment patterns without any information about plant closures, predicts how workers fare when their job actually disappears. Workers in jobs with many close substitutes transition more smoothly to alternative employment; workers in jobs with concentrated diversion face more difficult adjustments. The decomposition further shows that both geographic access and occupational fit contribute to this prediction, but through distinct temporal channels. This validation also illustrates a potential policy application: author-

Figure 8: Decomposing the Sources of Displacement Recovery



Notes: Event studies of earnings around plant closures, splitting displaced workers by quartiles of two counterfactual substitutability indices. Panel (a) uses an index constructed from geographic and state-dependence components of μ_{ij} only: commuting distance splines, family distance splines, and the state-dependence indicator. Panel (b) uses an index constructed from skill and match components only: occupation similarity and education-occupation mismatch (overqualification, subject match). In each panel, workers are split into bottom (Q1, low substitutability) and top (Q4, high substitutability) quartiles. The y-axis shows the average treatment effect on normalized earnings relative to $t = -3$. Shaded areas indicate 95% confidence intervals. Control group consists of workers at surviving establishments.

ities evaluating a proposed merger could use the substitutability index to identify which affected workers would face the greatest adjustment costs if the merger leads to job losses.

7 Conclusion

This paper develops a structural approach to measuring employee diversion ratios in labor markets. Using Swedish administrative data covering more than 2.6 million workers and 615,000 jobs, we recover rich worker-job substitution patterns without imposing arbitrary labor market boundaries and without relying on observed job-to-job flows, which conflate voluntary and involuntary transitions. The resulting diversion matrix shows that competition for workers is highly heterogeneous. Most jobs face diffuse competition from many alternatives, but an important tail of jobs has much more concentrated substitution patterns.

We validate these predictions out of sample using plant closures. The model reproduces the main reallocation patterns across commuting zones and occupations and substantially outperforms a simple logit benchmark, which reallocates displaced workers too mechanically toward large destinations. The substitutability index also predicts realized adjustment after displacement: workers who lose high-substitutability jobs experience a smaller earnings dip and nearly fully recover, whereas workers who lose low-substitutability jobs suffer larger and much more persistent losses. Decomposing the index into geographic and skill-based components reveals that these two channels operate on different time horizons: geographic access to alternatives governs the speed of re-employment, while occupational fit determines the persistence of earnings losses.

These diversion ratios are partial-equilibrium objects, but they provide an operational measure of worker-side competitive pressure. They help identify the jobs and workers that are most exposed when employment options disappear, whether because of layoffs, restructuring, or employer consolidation. More broadly, the framework offers a way to connect labor market competition to the realized costs workers bear when jobs are lost.

References

Agarwal, Nikhil and Paulo Somaini, "Demand Analysis under Latent Choice Constraints," *Review of Economic Studies*, 2025. Forthcoming.

Arnold, David, "Mergers and acquisitions, local labor market concentration, and worker outcomes," 2021.

Athey, Susan, Lisa K Simon, Oskar N Skans, Johan Vikstrom, and Yaroslav Yakymovych, "The heterogeneous earnings impact of job loss across workers, establishments, and markets," *arXiv preprint arXiv:2307.06684*, 2023.

Baker, Stuart G, "The multinomial-Poisson transformation," *Journal of the Royal Statistical Society: Series D (The Statistician)*, 1994, *43* (4), 495–504.

Beuschlein, Jakob, Jósef Sigurdsson, and Horng Chern Wong, "The Labor Market Consequences of Acquisitions," Discussion Paper 25076, RFBerlin 2025.

Caldwell, Sydnee and Oren Danieli, "Outside Options in the Labor Market," *The Review of Economic Studies*, 2024, *91* (6), 3286–3315.

Conlon, Christopher and Julie Holland Mortimer, "Empirical properties of diversion ratios," *The RAND Journal of Economics*, 2021, *52* (4), 693–726.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin, "Fast Poisson estimation with high-dimensional fixed effects," *The Stata Journal: Promoting communications on statistics and Stata*, March 2020, *20* (1), 95–115.

Eliason, Marcus and Donald Storrie, "Lasting or Latent Scars? Swedish Evidence on the Long-Term Effects of Job Displacement," *Journal of Labor Economics*, October 2006, *24* (4), 831–856.

Fan, Tianyu, "The Labor Market Incidence of New Technologies," Technical Report, Yale University 2025. Job Market Paper.

Farrell, Joseph and Carl Shapiro, "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition," *The B.E. Journal of Theoretical Economics*, 2010, *10* (1).

Fogel, Jamie and Bernardo Modenesi, "What is a Labor Market? Classifying Workers and Jobs Using Network Theory," Technical Report, arXiv 2023. arXiv:2311.00777.

_ , _ , **and Benjamin Scuderi**, "Market Definition Bias in Studies of (Labor) Market Power," Technical Report, Working Paper 2024.

Frejinger, E., M. Bierlaire, and M. Ben-Akiva, "Sampling of alternatives for route choice modeling," *Transportation Research Part B: Methodological*, 2009, *43* (10), 984–994.

Geurts, Karen and Johannes Van Biesebroeck, "Employment Growth Following Takeovers," *The RAND Journal of Economics*, 2019, *50* (4), 916–950.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, *83* (4), 685–709.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin, "Granular Search, Market Structure, and Wages," *The Review of Economic Studies*, 2024, *91* (6), 3569–3607.

Nimczik, Jan Sebastian, "Job Mobility Networks and Data-driven Labor Markets," 2023.

Prager, Elena and Matt Schmitt, "Employer Consolidation and Wages: Evidence from Hospitals," *American Economic Review*, February 2021, *111* (2), 397–427.

Qiu, Y Jeff, Masayuki Sawada, and Gloria Sheu, "Win/loss data and consumer switching costs: Measuring diversion ratios and the impact of mergers," *The Journal of Industrial Economics*, 2024, *72* (1), 327–355.

Seim, David, "On the incidence and effects of job displacement: Evidence from Sweden," *Labour Economics*, April 2019, *57*, 131–145.

Sorkin, Isaac, "Ranking firms using revealed preference.," *The Quarterly Journal of Economics*, 2018, *133* (3).

Train, Kenneth E, Daniel L McFadden, and Moshe Ben-Akiva, "The demand for local telephone service: A fully discrete model of residential calling patterns and service choices," *The RAND Journal of Economics*, 1987, pp. 109–123.

United States District Court for the District of Oregon, "Opinion and Order on Motion for Preliminary Injunction, FTC v. The Kroger Co. and Albertsons Cos., Inc.," 2024. Case No. 3:24-cv-00347-AN (D. Or. Dec. 10, 2024).

A Appendix

A.1 Education-Occupation Matrix Details

Statistics Sweden's education-occupation matching combines normative and statistical methods. The qualitative method defines in advance which education programs fit which professions, independent of actual distributions. The statistical method analyzes how education backgrounds are distributed within occupations to detect over- or under-represented combinations. Uncommon combinations can still be accepted if statistically over-represented.

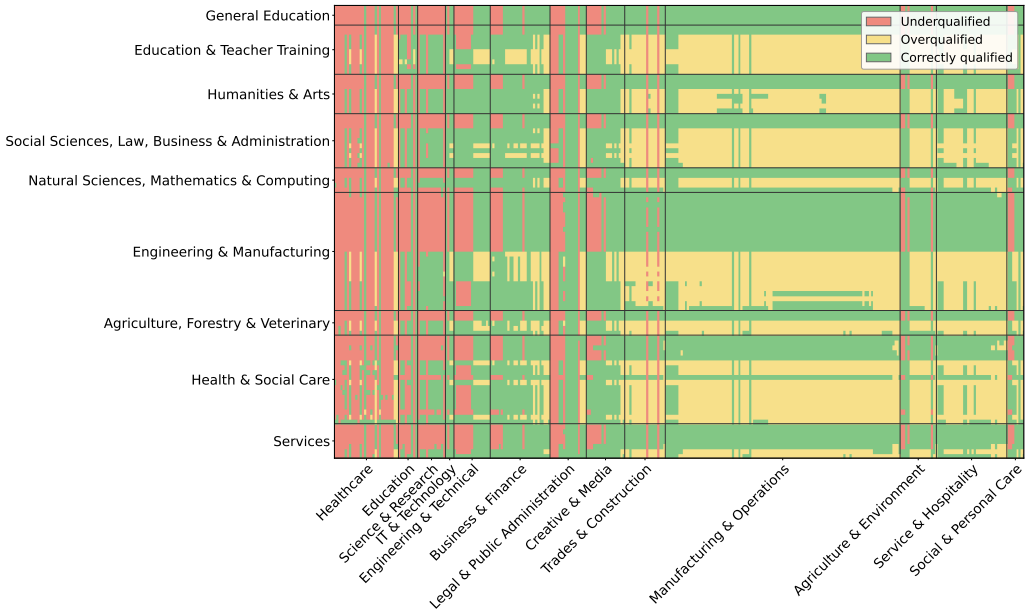
The classification is based on approximately 120 education groups and 140 occupation groups, forming several thousand combinations. It uses data from education-occupation distributions, job advertisements, and training providers.

Each education-occupation pair is evaluated along two dimensions. Subject matching captures how well the education's field fits the job, classified as complete, partial, or no match. Level matching captures how well the education's level fits job requirements, classified as lower, correct, or higher than required.

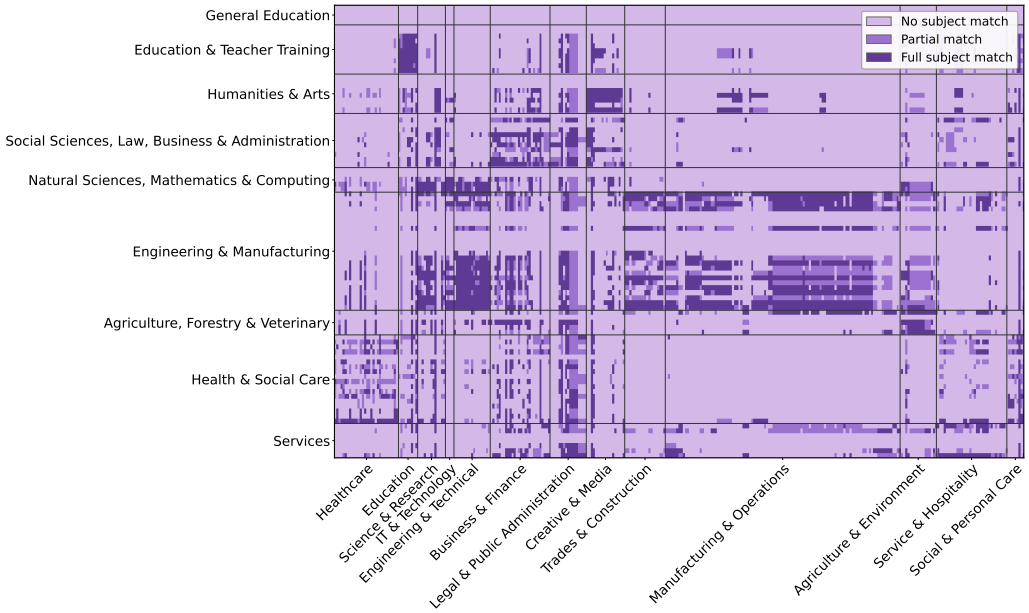
A.2 Education and Occupation Group Definitions

Table A.1 lists the nine education field groups used in Figure A.1, with the individual SUN2000 education codes assigned to each group. Table A.2 lists the thirteen occupation groups, with the individual SSYK96 four-digit codes assigned to each group.

Figure A.1: Education-Occupation Match Quality



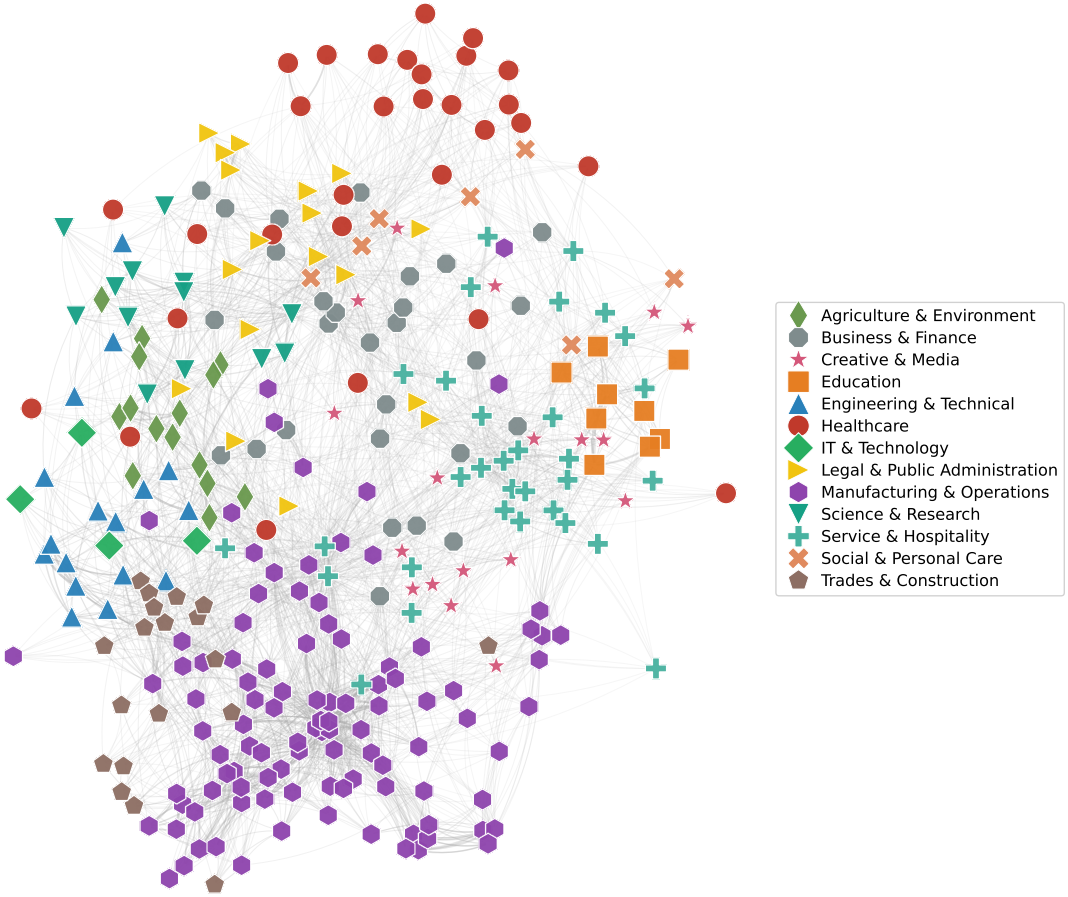
(a) Feasibility



(b) Subject Fit

Notes: Each row corresponds to a detailed education code and each column to a detailed occupation code, ordered by broader education fields and occupation groups. Panel (a): Feasibility. Red indicates under-qualified, yellow over-qualified, and green correctly qualified. Panel (b): Subject Fit. Darker purple indicates a stronger subject fit. Thick lines separate groups.

Figure A.2: Occupation Similarity Network



Notes: Network visualization of occupation similarity based on skill overlap in job advertisements. Each node represents a 4-digit occupation (SSYK96), colored by occupation group. Edges connect each occupation to its ten most similar occupations for visualization only; the estimation uses the full similarity matrix. Edge thickness reflects similarity score. Spring layout positions similar occupations closer together.

Table A.1: Education Field Groups

Group	SUN2000 Codes
General Education	01Z, 02Z, 03A, 03N
Education & Teacher Training	13Z, 14P, 15B, 15F, 15G, 15H, 15P, 15S, 15V, 15X
Humanities & Arts	23E, 23M, 24P, 25H, 25K, 25M, 25T, 25X
Social Sciences, Law, Business & Admin.	33E, 33H, 34P, 35B, 35E, 35F, 35J, 35M, 35P, 35S, 35X
Natural Sciences, Math. & Computing	43Z, 44P, 45D, 45N, 45X
Engineering & Manufacturing	53A, 53B, 53E, 53F, 53K, 53L, 53P, 53R, 53T, 53V, 53X, 54P, 55A, 55C, 55D, 55E, 55F, 55G, 55H, 55I, 55J, 55K, 55L, 55X
Agriculture, Forestry & Veterinary	63Z, 64P, 65J, 65S, 65V, 65X
Health & Social Care	73B, 73O, 73T, 73X, 74P, 75A, 75B, 75D, 75F, 75H, 75J, 75L, 75N, 75O, 75P, 75T, 75V, 75X
Services	83R, 83T, 83X, 84P, 85M, 85P, 85T, 85X

Table A.2: Occupation Groups

Group	SSYK96 Codes
Healthcare	2221, 2222, 2223, 2224, 2225, 2229, 2231, 2232, 2233, 2234, 2235, 2236, 2491, 3134, 3221, 3222, 3223, 3224, 3225, 3226, 3227, 3228, 3229, 3231, 3232, 3233, 3234, 3235, 3239, 5132, 5135
Education	2321, 2322, 2323, 2330, 2340, 2351, 2359, 3310, 3320
Science & Research	2111, 2112, 2113, 2114, 2121, 2122, 2211, 2212, 2310, 2442, 2443, 3111, 3240
IT & Technology	2131, 2139, 3121, 3122
Engineering & Technical	2141, 2142, 2143, 2144, 2145, 2146, 2147, 2148, 2149, 3112, 3113, 3114, 3115, 3116, 3117, 3118, 3119
Business & Finance	2411, 2412, 2413, 2414, 2419, 2441, 3411, 3412, 3413, 3415, 3416, 3417, 3418, 3419, 3421, 3423, 3429, 3431, 3433, 3476, 4111, 4112, 4120, 4190, 4212, 4214, 4215, 5227
Legal & Public Admin.	2421, 2422, 2423, 2429, 2470, 2480, 3151, 3152, 3432, 3441, 3442, 3443, 3449, 3450, 5151, 5152, 5153, 5159
Creative & Media	2431, 2432, 2444, 2451, 2452, 2453, 2454, 2455, 2456, 3131, 3132, 3133, 3471, 3472, 3473, 3474, 4140, 5210
Trades & Construction	7121, 7122, 7123, 7124, 7129, 7131, 7132, 7133, 7134, 7135, 7136, 7139, 7141, 7142, 7143, 7241, 7242, 7243, 9310
Manufacturing & Operations	3141, 3142, 3143, 3144, 3145, 3422, 4131, 4132, 4150, 7111, 7112, 7211, 7212, 7213, 7214, 7215, 7216, 7221, 7222, 7223, 7224, 7231, 7232, 7233, 7311, 7312, 7313, 7321, 7322, 7323, 7324, 7330, 7341, 7342, 7343, 7344, 7411, 7412, 7413, 7421, 7422, 7431, 7432, 7433, 7434, 7435, 7441, 7442, 8111, 8112, 8121, 8122, 8123, 8124, 8125, 8130, 8141, 8142, 8143, 8144, 8150, 8160, 8170, 8211, 8212, 8221, 8222, 8223, 8224, 8229, 8231, 8232, 8240, 8251, 8252, 8253, 8261, 8262, 8263, 8264, 8265, 8269, 8271, 8272, 8273, 8274, 8275, 8276, 8277, 8278, 8279, 8281, 8282, 8283, 8284, 8285, 8290, 8311, 8312, 8321, 8322, 8323, 8332, 8333, 8334, 8340, 9141, 9150, 9320, 9330
Agriculture & Environment	2213, 2214, 3211, 3212, 6111, 6112, 6113, 6121, 6122, 6129, 6130, 6140, 6151, 6152, 6153, 8331, 9210
Service & Hospitality	3414, 3462, 3475, 4211, 4213, 4221, 4222, 4223, 4224, 5111, 5112, 5113, 5121, 5122, 5123, 5141, 5142, 5149, 5221, 5222, 5223, 5224, 5225, 5226, 5227, 7137, 9110, 9121, 9122, 9123, 9130, 9142, 9143, 9190
Social & Personal Care	2460, 2492, 3461, 3480, 5131, 5133, 5134, 5139

A.3 Aggregating Diversion Ratios

For occupation/region m and l:

$$D_{jk} = \frac{1}{s_j} \int \frac{s_{ij} s_{ik}}{1 - s_{ij}} dF(i) \tag{13}$$

$$= \frac{1}{s_j} \left(s_k(\mathcal{J}) - s_k(\mathcal{J} \setminus j) \right) \tag{14}$$

$$D_{ml} = \frac{1}{\sum_{j \in m} s_j} \sum_{j \in m} \sum_{k \in l} \left(s_k(\mathcal{J}) - s_k(\mathcal{J} \setminus j) \right) \tag{15}$$

$$= \frac{\#\text{ switches from } m \text{ to } l}{\#\text{ workers at } m} \tag{16}$$

A.4 Wage Coefficient Estimation

Table A.3 reports OLS and IV estimates of the wage coefficient from regressing estimated job fixed effects on log wages, controlling for municipality and 4-digit occupation fixed effects.

Table A.3: Recovering the Wage Coefficient

	OLS	IV
Log wage	0.313***	3.447***
	(0.018)	(0.172)
Municipality FE	Yes	Yes
4-digit occupation FE	Yes	Yes

Notes: Dependent variable: estimated job fixed effect $\hat{\delta}_j$. Full-time adjust wage at job j. The IV specification instruments log wage with log value added per worker. All specifications include municipality and 4-digit SSYK96 occupation fixed effects. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

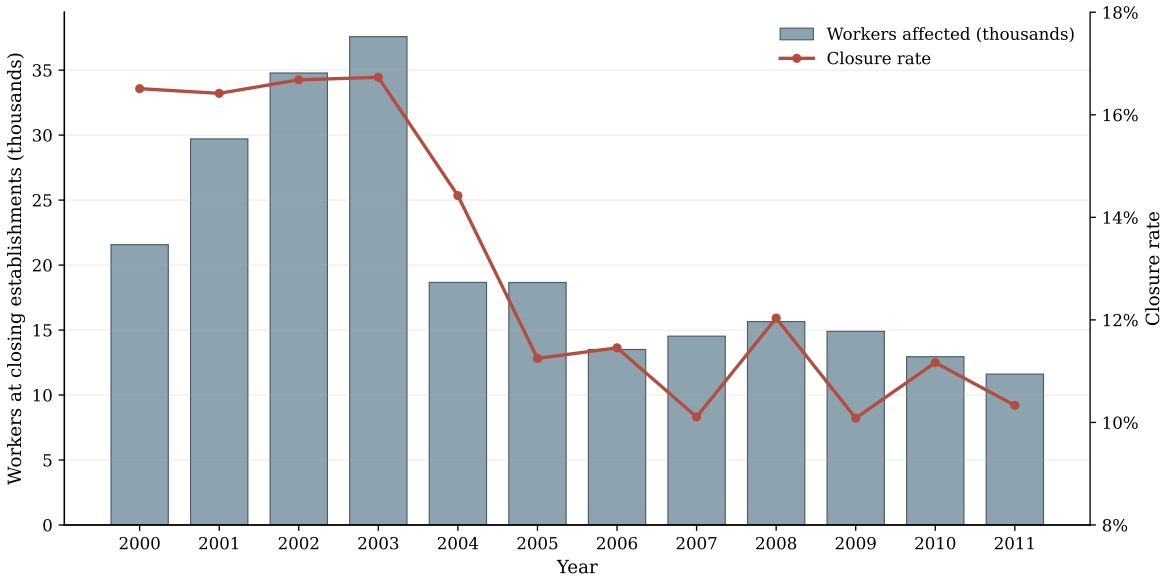
A.5 Merger Sample Construction

We identify mergers using administrative ownership changes in the Swedish establishment register. Each establishment (CFAR) is linked to an organizational unit (PeOrg). A merger event is defined as a year-to-year change in an establishment's PeOrg identifier across consecutive years. All establishments sharing the same old PeOrg, new PeOrg, and year are grouped into a single merger event.

We apply several filters to exclude false or economically insignificant transitions. The transitioning establishment must have average employment of at least five across all observed years. At least 50 percent of the old firm's total employment must be involved in the transition, and the old firm must lose at least 50 percent of its employment at the time of the transition. Events where the old and new firm already share more than 50 percent of their establishments are dropped, as these likely reflect internal reorganizations rather than mergers.

We define an establishment as closed if its employment falls to 10 percent or less of its pre-merger level (measured at $t-1$) within two years of the merger, or if it disappears from the register entirely. Affected workers are measured as the establishment's $t-1$ employment.

Figure A.3: Merger-Related Closures in Sweden



Notes: This figure summarizes mergers identified from year-to-year changes in establishment ownership in the Swedish establishment register between 2000 and 2011. An establishment is classified as closed if employment falls to 10% or less of its pre-merger level within two years of the merger, or if it disappears from the register entirely. Affected workers are measured by the establishment's employment in the year before the merger. See above for details on sample construction.

A.6 Displacement Sample and Decomposition

Sample construction. We identify plant closures using the November 2010 and November 2011 establishment snapshots. An establishment is classified as closed if it either disappeared entirely or shrank by at least 90% of its workforce. We exclude closures consistent with relabeling (the same workers reappearing under a new establishment identifier)

or internal reorganization within the same firm. Only establishments with at least five employees prior to closure are included. Displaced workers are defined as individuals aged 24–60 with at least three years of continuous tenure at the closing establishment, whose annual earnings exceed the 10th percentile of the earnings distribution. We exclude military personnel and senior government officials.

Control group. We draw control workers from establishments that survive throughout the sample period, subject to the same age, tenure, and earnings restrictions as the displaced sample. We randomly sample approximately three control workers per displaced worker. Both groups are followed in a balanced panel from $t = -5$ to $t = +7$ relative to the closure year, using November employment snapshots. The outcome variable is annual earnings normalized to the worker's own earnings at $t = -3$. Workers who are non-employed in a given year receive zero earnings.

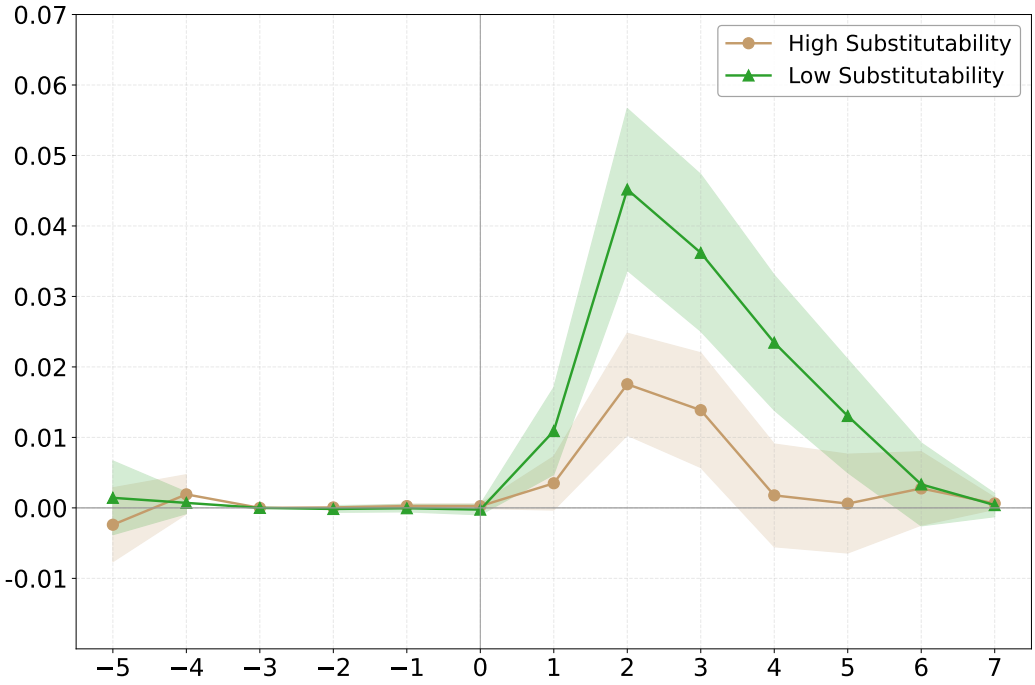
Extensive margin. Figure A.4 decomposes the total earnings effect into the extensive margin: the differential probability of non-employment between displaced and control workers. In the first two years after closure, low-substitutability workers (bottom quartile) are roughly 4.5 percentage points more likely to be non-employed than their controls, compared to a 1.8 percentage point gap for high-substitutability workers (top quartile). This differential narrows steadily and converges to zero by approximately six years after displacement, indicating that the extensive margin effect is transitory.

Intensive margin. Figure A.5 presents the intensive margin: earnings losses conditional on being employed. At $t = 1$, low-substitutability workers lose approximately 29% of their pre-displacement earnings compared to 19% for high-substitutability workers. While high-substitutability workers recover to within roughly 3% of their pre-displacement level by $t = 7$, low-substitutability workers still suffer earnings losses of approximately 14% at the same horizon.

The intensive margin dominates the overall gap between the two substitutability groups. The non-employment differential, while economically meaningful in the short run, is transitory; the conditional earnings gap, by contrast, persists throughout the seven-year follow-up window. This pattern suggests that the long-run cost of losing a low-substitutability job operates primarily through lower re-employment wages rather than through prolonged joblessness.

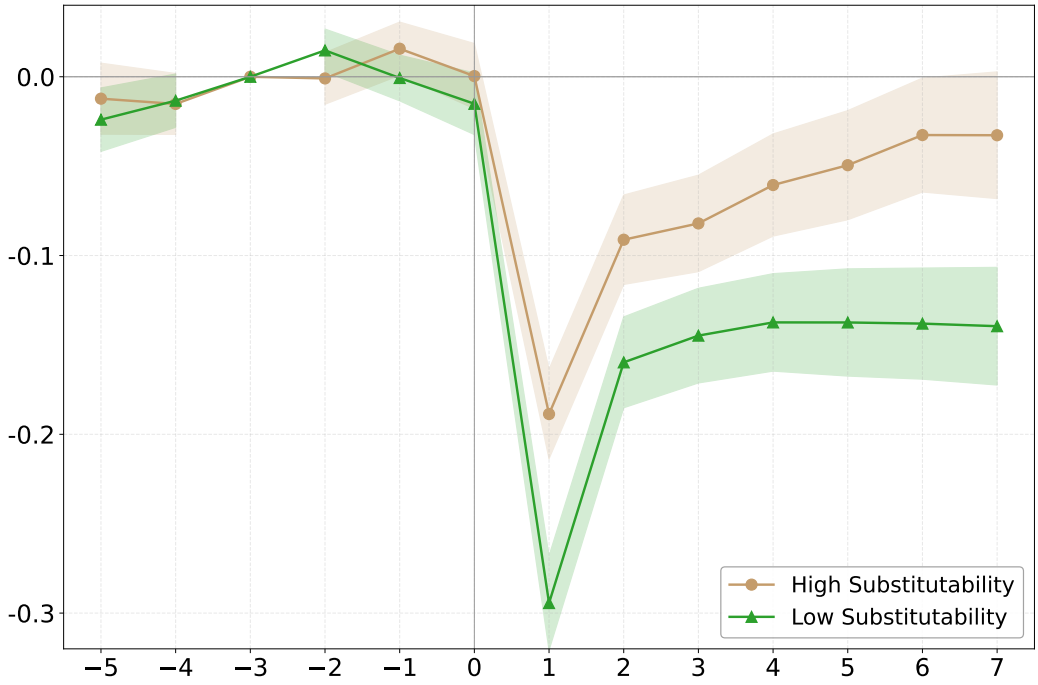
A.7 Additional Figures

Figure A.4: Extensive Margin: Non-Employment After Displacement



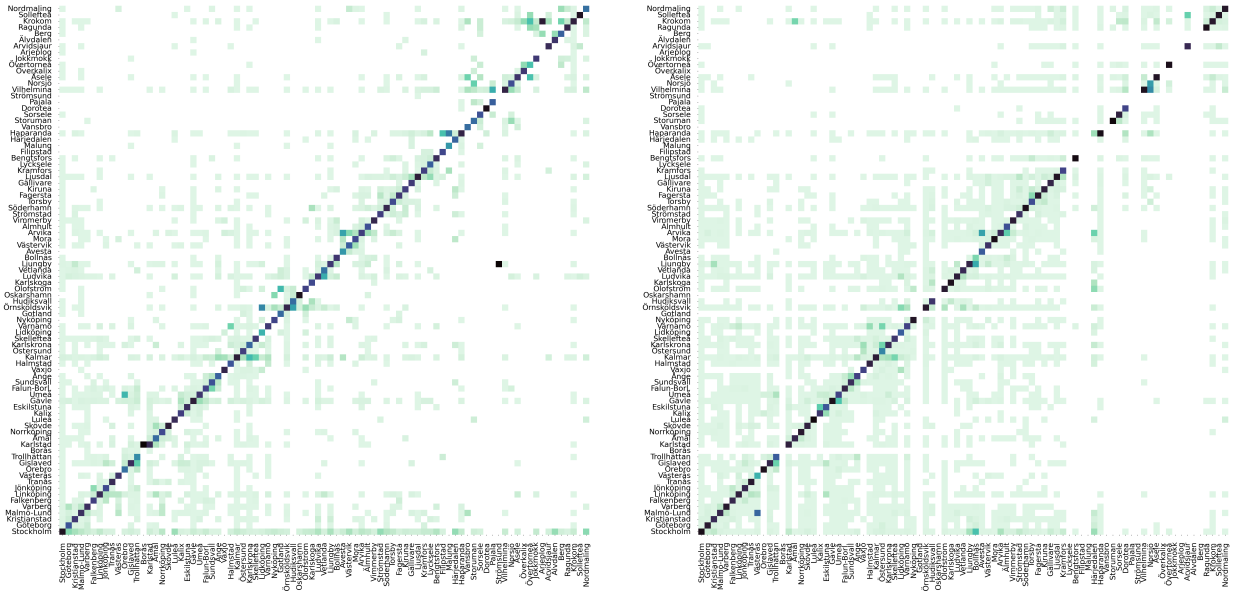
Notes: Event study of non-employment probability around plant closures, comparing displaced workers to controls. Workers split by the substitutability index of their lost job: top quartile (Q4, high substitutability) vs. bottom quartile (Q1, low substitutability). Shaded areas indicate 95% confidence intervals. The extensive margin gap between the two groups peaks at $t = 2$ and converges to zero by $t = 6$.

Figure A.5: Intensive Margin: Conditional Earnings After Displacement



Notes: Event study of earnings conditional on employment around plant closures, comparing displaced workers to controls. Workers split by the substitutability index of their lost job: top quartile (Q4, high substitutability) vs. bottom quartile (Q1, low substitutability). Shaded areas indicate 95% confidence intervals. Earnings normalized to three years before closure ($t = -3$).

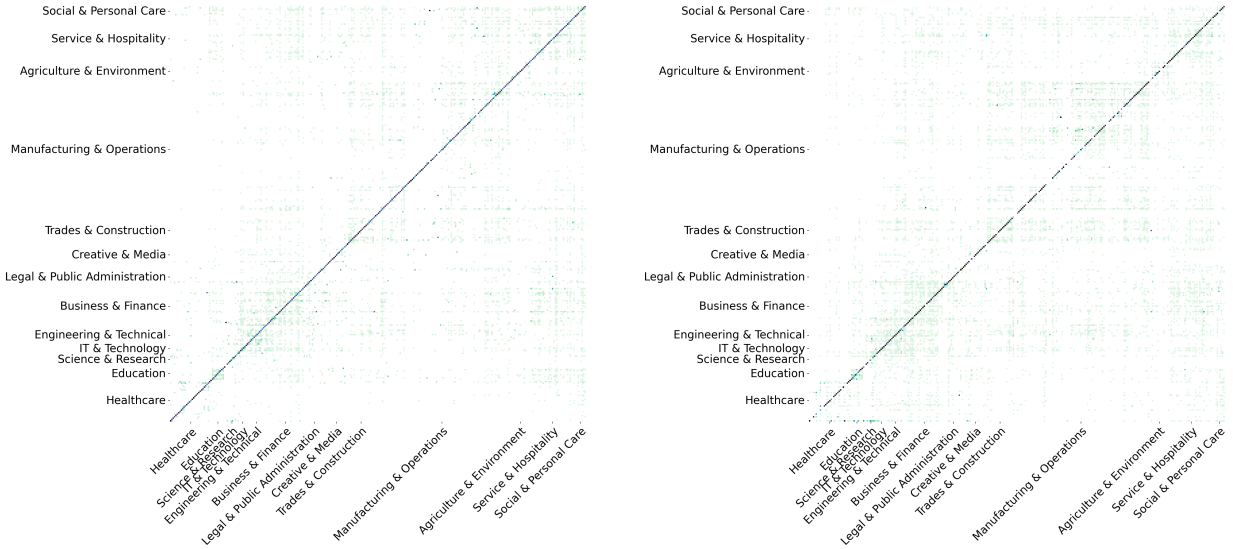
Figure A.6: Diversion Ratios Across Commuting Zones



(a) Plant Closures (b) Model Estimated

Notes: The x-axis shows origin commuting zones and the y-axis destination commuting zones. Panel (a) shows observed reemployment transitions after plant closures and panel (b) model-implied diversion ratios. Both panels exclude transitions into unemployment.

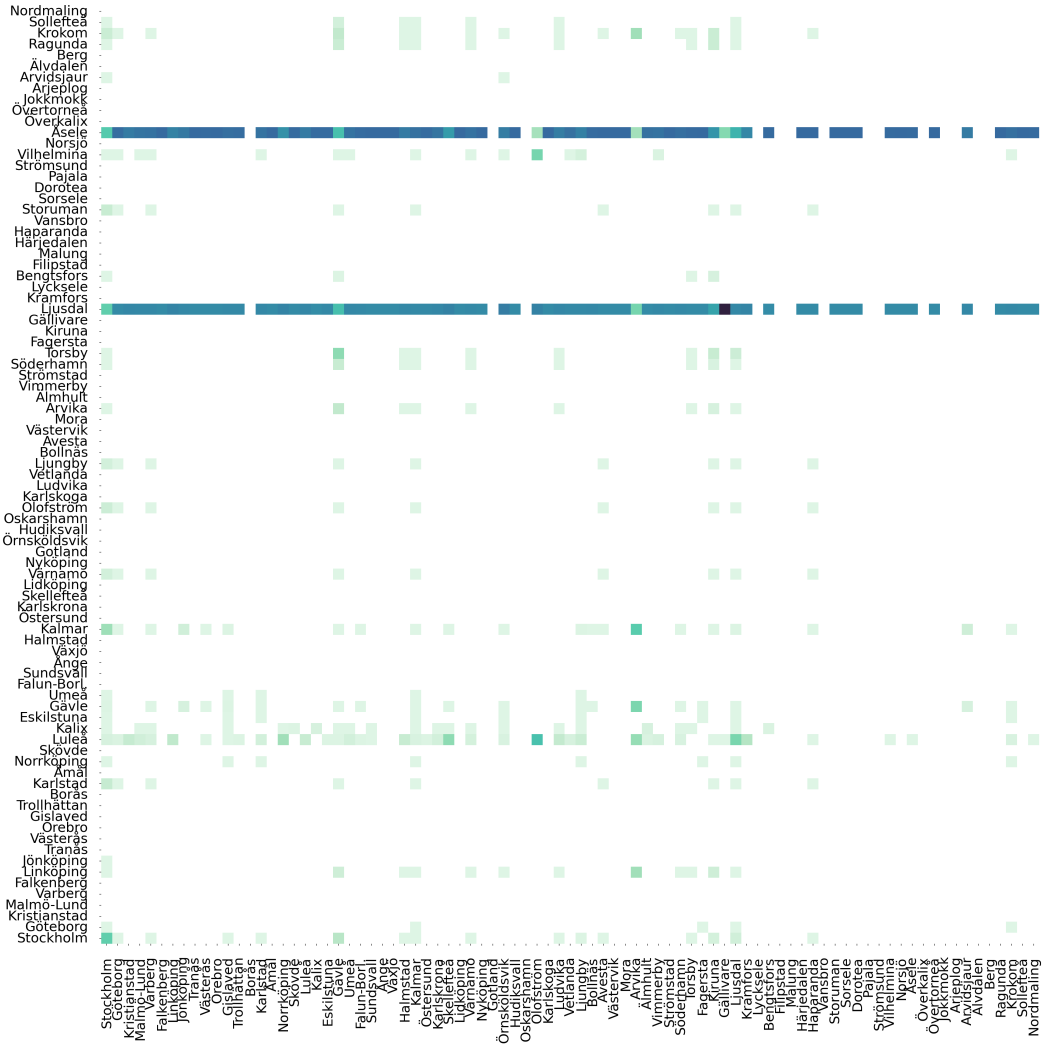
Figure A.7: Diversion Ratios Across Occupations



(a) Plant Closures (b) Model Estimated

Notes: The x-axis shows origin occupations and the y-axis destination occupations. Panel (a) shows observed reemployment transitions after plant closures and panel (b) model-implied diversion ratios. Both panels exclude transitions into unemployment.

Figure A.8: Simple Logit Diversion Matrix Across Commuting Zones



Notes: Diversion matrix across commuting zones implied by the simple logit benchmark used in Section 6.

Figure A.9: Simple Logit Diversion Matrix Across Occupations



Notes: Diversion matrix across occupations implied by the simple logit benchmark used in Section 6.