WORKING PAPER 2024:10

The heterogeneous earnings impact of job loss across workers, establishments, and markets

Susan Athey
Lisa K. Simon
Oskar N. Skans
Johan Vikström
Yaroslav Yakymovych

IFAU
INSTITUTE FOR EVALUATION OF LABOUR MARKET AND EDUCATION POLICY

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala.

IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

Papers published in the Working Paper Series should, according to the IFAU policy, have been discussed at seminars held at IFAU and at least one other academic forum, and have been read by one external and one internal referee. They need not, however, have undergone the standard scrutiny for publication in a scientific journal. The purpose of the Working Paper Series is to provide a factual basis for public policy and the public policy discussion.

More information about IFAU and the institute's publications can be found on the website www.ifau.se

ISSN 1651-1166

The heterogeneous earnings impact of job loss across workers, establishments, and markets[*]

Susan Athey[†] Lisa K. Simon[‡] Oskar N. Skans[§] Johan Vikström[¶]
Yaroslav Yakymovych[||]

April 29, 2024

Abstract

Using generalized random forests and rich Swedish administrative data, we show that the earnings effects of job displacement due to establishment closures are highly heterogeneous. We find as much heterogeneity within as across closing establishments, and within as across worker types defined by age and schooling. We display the potential of market-based policy interventions by showing that much of the heterogeneity across establishments is shared within markets. Several results suggest that the effect heterogeneity disfavors already vulnerable workers. Thus, targeted policy interventions may be justified to a larger extent than suggested by estimated average earnings effects.

Keywords: Plant closures, heterogeneous effects, GRF
JEL-codes: J65, J21, J31, C45

[*]We thank Stefan Eriksson, Peter Fredriksson, Vitor Hadad, Martin Huber, Michael Lechner, Nicolaj Mühlbach, Stefan Pitschner, Kjell Salvanes, David Strömberg, Erik Sverdrup and seminar participants at Aalto, Cattolica Milan, CREAM, EALE, IfFS, Ratio, SKILS, SOLE, SSE, Stanford HAI, Swedish Economics Meeting, Umeå, Uppsala and Örebro. Skans and Yakymovych were supported by Vetenskapsrådet (2018-04581). Athey and Simon acknowledge support from the Golub Capital Social Impact Lab at Stanford GSB.

[†]Stanford University, GSB
[‡]Revelio Labs, New York
[§]Dep. of Economics, Uppsala University oskar.nordstrom_skans@nek.uu.se
[¶]IFAU; Dep. of Economics, Uppsala University
[||]Institute for Housing and Urban Research, Uppsala University

1 Introduction

Structural change and job reallocation drives economic growth (Bartelsman and Doms, 2000) while causing large persistent earnings losses for individual workers (Jacobson et al., 1993). These earnings losses are sizeable enough to adversely affect workers' health and well-being in multiple dimensions.[1] Consequently, governments spend considerable resources on social programs designed to mitigate the economic consequences of structural change (OECD, 2019). The policy mix includes *ex post* economic transfers and active labor market policies, as well as *ex ante* policies such as employment protection legislation, short-time work, and life-long learning schemes. To minimize costs and distortions, policymakers often restrict these policies to only cover segments where they are perceived to be needed the most, i.e. segments where workers' earnings capacities are tied to their current jobs. As a consequence, eligibility and generosity vary by worker, firm, industry, and location characteristics. But even though available empirical evidence, and economic theory, strongly suggest that job loss leads to heterogeneous earnings effects, the literature still lacks systematic evidence that could guide policy makers towards the workers who suffer the largest losses when jobs are destroyed.

This paper aims to fill this gap by showing how a rich set of policy-relevant characteristics interact to predict workers' economic losses when jobs are destroyed due to establishment closures. Our analysis relies on exceptionally rich Swedish administrative data that allow us to characterize individuals, establishments, and market conditions at the time of displacement events. Using closures instead of individual unemployment spells or exposure to adverse demand trends allows us to study negative shocks which are well-identified in time, unrelated to changes in personal circumstances, and observed even for the most resilient workers who move to a new job without an intermission. In line with a vast literature on the impact of mass layoffs and establishment closures pioneered by Jacobson et al. (1993), we rely on a selection-on-observables strategy, comparing earnings trajectories for displaced workers to a carefully matched control group of workers who did not experience an establishment closure in the same year.

In the paper, we outline a stylized theoretical framework to highlight how several well-documented labor market processes will generate heterogeneous displacement effects related to a broad set of worker, establishment, and market characteristics. Job loss should not affect workers' earnings if labor markets are competitive, without frictions or unemployment. However, many types of deviations from the competitive baseline

[1] See, e.g., Eliason (2014a) on increased drinking, Black et al. (2015) on increased smoking and Eliason (2012) on increased divorces.

can explain why displaced workers experience earnings losses, and these distortions are unlikely to play an equal role for all workers.

Previous research has examined how displacement effects vary across various dimensions. On the supply side, the literature has emphasized the role of human capital, gender, and age (see Davis and von Wachter, 2011, for an overview).[2] Other dimensions include white- and blue-collar workers (e.g., Schwerdt et al., 2010), immigrants from different countries (Bratsberg et al., 2018), education and skills (Seim, 2019), within-family inter-dependencies (Halla et al., 2020), life-cycle patterns (Salvanes et al., 2023), interactions between age and gender (Ichino et al., 2007), and race and gender (Hu and Taber, 2011). On the demand side, articles emphasize the role of job content (Blien et al., 2021; Yakymovych, 2022), occupation-specific human capital (Huckfeldt, 2022; Braxton and Taska, 2023), worker-employer match effects (Lachowska et al., 2020), sector or industry (Eliason, 2014b; Helm et al., 2023), establishment-specific wage premia (Bertheau et al., 2023; Gulyas and Pytka, 2021; Schmieder et al., 2023), regional structural change (Arntz et al., 2022), the size of the displacement event (Gathmann et al., 2020; Cederlöf, 2019) and aggregate business cycle conditions (Eliason and Storrie, 2006; Farber, 2011; Davis and von Wachter, 2011; Schmieder et al., 2023).[3]

We synthesize most aspects of heterogeneity highlighted in the job-loss literature within one unified flexible framework while retaining a focus on variables that in principle can be observed and used by central policy makers or local caseworkers. Using rich Swedish data, we characterize displaced workers, establishments, industries and locations. At the worker level, we capture age, gender, general and specific human capital (years and type of schooling, general and specific labor market experience), detailed family status, as well as internal and external migration history. Establishments and jobs are characterised by task content, education-industry match quality, establishment wage premium, establishment size, and size relative to the local market. We characterize the worker's pre-displacement industry through its industry wage premium, dynamism and employment trends. Locations are characterized by population density, unemployment rates, and industry composition and concentration.

We estimate displacement effects using the *Generalized Random Forest* (GRF) (Athey et al., 2019), which allows the effects to vary flexibly with our set of variables. The GRF combines a large number of "causal trees", each of which is constructed by sequentially

[2] During 2000–2022, 25 NBER and IZA papers reported displacement effect estimates. Most of them estimate heterogeneous effects, with age and gender as the two most common dimensions.

[3] Other related papers include Britto et al. (2022), who use GRF to study increased crime rates following job loss, and Mueller and Spinnewijn (2023), who study predictors of long-term unemployment, and conclude that job seekers' employment history is a key factor.

splitting the data at covariate values selected to maximize treatment-effect heterogeneity at each split. This is repeated in multiple bootstrapped iterations, and the endpoints of all the trees ("leaves") are then combined to obtain Conditional Average Treatment Effect (CATE) estimates. These CATEs are non-parametric estimates of treatment effects as functions of the covariates. To minimize the risk of overfitting, we split the data into five establishment-clustered folds and estimate a GRF for each fold using data from the four remaining folds.[4] We then rank the workers (within fold) by their CATE estimates, split the data into quantiles based on this ranking, and estimate average differences in outcomes between treated and controls (ATE estimates) for each group. Thus, the ATE estimates for each quantile group do not reuse any of the outcome data that the GRF relied on when defining the quantile groups.

Our results show that job loss leads to large, persistent, and highly heterogeneous earnings losses. On average, displacement reduces annual earnings by 24 percent, and employment by 15 percentage points, one year after a closure. One third of the earnings effect remains 10 years later.[5] Before turning to the GRF, we estimate a sequence of linear earnings regressions where each regression model interacts the displacement dummy with one variable of interest. More than 30 of our variables, capturing very different processes, are significant if interacted with displacement dummies in these regressions. We use the GRF to handle this multidimensional heterogeneity in a flexible way. The resulting CATE estimates capture extensive and complex heterogeneity. The 10 percent of workers with the lowest CATE estimates lose almost 50 percent of pre-displacement earnings during the calendar year after displacement. This is 2.5 times larger than the median, and 8 times larger than the least affected decile. Even though the model is trained to estimate effects in the first year after job loss, the underlying patterns are stable enough that the short term estimated CATEs also accurately predict longer term outcomes; they further predict outcomes for years outside the period used for estimation.

Policy targeting in this setting is non-trivial since the job-loss effects are heterogeneous within and across worker-types, closing establishments, and markets. Age and schooling are two of the most important predictors. Losses grow with age regardless of schooling, in particular if workers are older than 50, and shrink with years of schooling, regardless of age.[6] However, our CATE estimates exhibit substantial variation within "buckets" defined by age and schooling combinations. For almost every bucket, the

[4]Clustering on closure events is crucial to ensure that different data folds can be treated as independent.

[5]These estimated earnings effects are consistent with previous Swedish studies, see e.g., Eliason and Storrie (2006) for men in the private sector, and Eliason (2014b) for women in the public sector.

[6]Workers above 50 are disregarded in most previous studies. An important recent exception is Salvanes et al. (2023), finding similar results to us.

least resilient quartile within the bucket suffers earnings losses in excess of 30 percent, whereas the most resilient quartile lose less than 20 percent. Much of the remaining predicted heterogeneity is related to characteristics of industries (e.g., manufacturing) and locations (e.g., population density).

To show that effects vary both within and across establishments, we rank workers (by CATE) within establishments, rank establishments across events (using coworker CATEs), and estimate the treatment effects for these different cuts of the data. Our results show that effects are about as heterogeneous within as across closing establishments. Many recent studies (Schmieder et al., 2023; Bertheau et al., 2023; Helm et al., 2023; Gulyas and Pytka, 2021) point to lost rents from high-wage establishments as a key determinant of heterogeneous wage impact of job loss.[7] Our results puts some nuance to these results by showing evidence of substantial heterogeneity *within* closure events. It should also be noted that we focus on effects on earnings (rather than wages), where the role of firm-effects might be more limited.[8]

We further show that much of the heterogeneity across closure events is related to factors that are shared across different events within the same market (industry x location). This may be useful for policy makers who are unlikely to target specific establishments, but who may be willing to use place-based or industry-based policies. We show that market conditions matter the most for workers with worse-than-average predicted individual effects. Locations with large losses (net of industry and worker aspects) are primarily characterized by low population density. They also have high unemployment rates and a less diverse industry structure characterized by declining industries. Industries (3-digit) associated with large losses are all found in manufacturing. Many attributes associated with manufacturing (a high wage premium, low dynamism, negative employment trends) are also associated with larger displacement effects across different industries within (and outside of) manufacturing.

Several results indicate that job-loss has more severe economic consequences for already vulnerable workers. Workers with large predictable short-run losses also experience larger losses in the long run. Furthermore, they have lower pre-displacement earnings, a higher displacement propensity, and worse predicted counterfactual (non-

[7] Diverging results are found by Lachowska et al. (2020), instead pointing to employer-worker match quality, and Braxton and Taska (2023) emphasizing technology and occupational switching.

[8] In related parallel work, Gulyas and Pytka (2021) use GRF to study the wage impact of mass layoffs, using a set of characteristics related to wage setting theories. Their results indicate that wage losses primarily are related to firm-level pre-displacement rents. Our results, focusing on relative earnings effects, instead point to substantial heterogeneity both within and across closure events. Falling employment rates explains much of the earnings losses in the most affected groups, leaving a smaller role for firm/establishment rents. A plausible contributing factor is that we retain all workers up to age 60.

displaced) earnings trajectories. The fact that earnings consequences of job loss are systematically worse for economically vulnerable workers may be important for policy purposes as it suggest an added rationale for targeted policy interventions.

To synthesize our results, we study feasible targeting strategies. We assess whether policymakers can identify workers with large earnings losses using only one or two observed attributes and show that targeting on age, education, manufacturing industries or population density is substantially better than random targeting, but policymakers would need at least two variables to get close to the prediction of the GRF. The best two-dimensional rule, according to optimal policy trees (Athey and Wager, 2021), targets old workers in routine occupations. However, even with these two variables, the performance is significantly worse than non-parametric targeting with the GRF model. Finally, we use the GRF-estimates to evaluate the targeting of the existing Swedish redistribution system, showing that workers in the most affected groups are better insured than other displaced workers. But gradual policy changes have have eroded these differences over time, causing a larger pass-through to disposable income among vulnerable workers.

The targeting exercises reaffirm the overall conclusion that the economic consequences workers face if they lose their current job are far from uniform, and that the underlying heterogeneity cannot easily be attributed to a single factor. Some groups of workers – the young, well-educated, urban, non-manufacturing workers – recover from job loss with only marginal earnings losses, as if they faced a fully competitive labor market (i.e. their specific jobs were not associated with any economic rents). In contrast, other groups of workers, sometimes even within the same establishments, suffer very large earnings losses if their job disappears, a result which points to substantial labor market imperfections. This second group is characterized by a combination of important within-event and across-market heterogeneity related to, e.g., age, schooling, population density and routine job content. This extensive heterogeneity makes policy targeting both challenging and potentially useful. Our results point to a strong association between the magnitude of job loss effects and other adverse economic outcomes. Furthermore, adverse market conditions tend to increase the cost of job loss more for vulnerable workers. These patterns jointly suggest that policymakers who manage to target policies towards settings where the job-loss effects are particularly large also may succeed in mitigating the impact of job destruction on economic inequality.

The paper is structured as follows. Section 2 outlines a stylized model. Section 3 describes the data. Section 4 introduces heterogeneous effects and the GRF. Section 5 relates the effects to other economic outcomes. Section 6 discusses the predictors. Section 7 shows targeting results. Section 8 concludes. All appendices are online.

2 A Stylized Theoretical Framework

To illustrate why earnings effects of displacement are likely to be heterogeneous along multiple dimensions, we outline a stylized theoretical model. An important starting point is that on a fully competitive market, without search frictions or involuntary unemployment, jobs have no intrinsic value and workers' earnings will not depend on the fate of their employing firms. We therefore need a richer model to rationalize the well-documented negative earnings effects of displacement. Since a broad class of economic theories portray individual wages as a weighted sum of outside options and establishment- or job-specific factors (see, e.g., Card et al., 2016, for a discussion), we model the wage of worker i employed in job j as:

$$W_{ij} = \Omega_i + \beta_{ij} \times p_{ij}, \tag{1}$$

where Ω_i represents the outside option. p_{ij} is job-specific productivity and β_{ij} is a bargaining parameter capturing the share of the surplus extracted by the worker.[9]

The outside option can be modelled as the weighted average of the earnings as unemployed, b_i, and the agents' expectations regarding wages in the worker's alternative jobs $E(W_i)$, where the weights are given by the probability q_i of finding a new job:

$$W_{ij} = q_i \times E(W_i) + (1 - q_i) \times b_i + \beta_{ij} \times p_{ij}. \tag{2}$$

Our main analysis will, in line with conventions in the empirical literature, not consider non-labor income sources, but instead focus on how job loss affects gross labor earnings.[10] We therefore define the treatment effects Δ_i as the difference between labor earnings if displaced and the wage as employed in the initial job ($W_i j$). Displaced workers are either re-employed in a new job k, or non-employed (earning nothing). As a consequence, $\Delta_i = W_{ik} - W_{ij}$ if the worker is re-employed in job k, and $\Delta_i = -W_{ij}$ if the worker stays non-employed.[11]

We will consider treatment effects as a function of a vector of observable characteristics. To simplify the exposition in this section, we may assume that there exists a set of characteristics x that capture heterogeneity across all the key parameters, such

[9]Job-specific productivity, p_{ij}, is defined as the productive value of the match minus the outside option. Many models, such as the AKM model, focus on the special case where p and β can be treated as firm-specific constants, shared by all workers employed by the same firm, in log wage regressions.

[10]As a complement, we show estimates of the impact on disposable income, which also accounts for taxes and benefits.

[11]Empirically, we estimate this effect by comparing the earnings trajectory of displaced workers relative to a matched control group.

as the probability of re-employment, unemployment benefits, the distribution of re-employment wages, pre-displacement rents and bargaining strength. The conditional average treatment effect $\Delta(x)$ then becomes:

$$\Delta(x) = -W_{ij}(x) + q(x)E(W_{ik}(x)),$$

which, using equation (2), can be reformulated as

$$\Delta(x) = -(1 - q(x)) \times b(x) - \beta_j(x) \times p_j(x). \tag{3}$$

The simple model leading to equation (3) highlights that earnings losses will be largest in settings where the outside options as non-employed are particularly important and/or in settings with large idiosyncratic pre-displacement rents. Without unemployment or idiosyncratic rents, workers will immediately recover their pre-displacement wage at another firm. Factors related to variations in *re-hiring probabilities* and the *utility while unemployed* on the one hand and *bargaining power* and *productivity in the pre-displacement job* on the other, can therefore generate heterogeneous displacement effects.

Rehiring probabilities can vary with the aggregate cycle. But as long as workers are imperfectly mobile, rehiring probabilities will also vary with regional and industry-specific market conditions. The importance of these conditions depends on mobility costs, which are inherently related to family status and the specificity of the worker's human capital. In addition, the size of the specific market will matter, in particular if multiple displaced workers compete with each other on a the market with few jobs relative to the number of displaced workers (Gathmann et al., 2020; Cederlöf, 2019). In our setting, unemployment benefits are determined by a universal replacement rate with a fixed cap and minimum. In addition, the utility as unemployed may vary with the value of home production, possibly related to family obligations.

If different jobs or firms pay different wages, as in the AKM framework (Abowd et al., 1999), workers will lose more if they were employed by a higher paying employer before displacement. Several economic processes suggest that workers, on average, have better paying jobs before than after displacement. In a frictional market with on-the-job search, workers gradually move to higher paying firms or better matches. Firm closures leads to a restart of this process, suggesting that experienced workers will have more to lose when displaced. Similarly, workers with long tenure can be more productive in the pre-displacement job due to non-transferable firm-specific human capital accumulated through on-the-job training.[12] In addition, late-career workers may receive compensation based on early-career productivity if firms use tournaments or deferred compensation as performance incentives, which generates another source of rents that cannot be transferred across jobs. If wages are (downward) rigid within employment relationships, current wages may also reflect *past* outside options, suggesting that earnings losses are more pronounced when workers are displaced from declining firms, industries, and occupations. Our stylized set-up abstracts from *counterfactual* dynamics but, in the medium run, earnings losses should be lower if the counterfactual displacement risk is very high, i.e. in markets where dynamics matter more.

The main take-away from this brief discussion is that heterogeneous displacement effects will arise due to a set of well-documented economic phenomena such as regional variation in unemployment rates, mobility costs arising from family obligations, differences in labor demand across industries, gradual accumulation of firm-specific human capital, deferred compensation to long-tenured workers, and productivity dispersion across firms. As a consequence, it is unsurprising that different studies have documented heterogeneous treatment effects in relation to different observable characteristics. Below we show that treatment effects indeed are heterogeneous in relation to a large set of observable characteristics and interactions between them.

3 Data and the Average Effects of Displacement

We rely on establishment closures to study earnings losses after job displacement. Closures are well-identified in time, unrelated to individual choices, and observed even for workers who find new jobs directly after being displaced. Estimating the effects on earnings of displacements is challenging because affected workers are a non-random subset of the overall population of workers. We follow the literature pioneered by Jacobson et al. (1993) and compare the earnings trajectories of the displaced to a matched control group of workers from surviving establishments. This section describes our data and estimated average effects of displacement. The approach to estimating CATEs is described in Section 4 below.

[12] The literature has shown that specific human capital can be tied to various dimensions of jobs such as firms or industries, or occupations and task content (Huckfeldt, 2022; Braxton and Taska, 2023)

3.1 Displaced Workers and Control Workers

Our main data source is the Swedish linked employer-employee register *RAMS* during 1985–2017. These annual tax-based records contain all transfers from establishments to employees. An establishment is a production unit with a physical location belonging to a firm or organization.[13] We use the term "firm" for all legal entities. These data are linked to other records through person, establishment and firm identification numbers. We start by selecting a panel of all individuals in Sweden aged 16 to 64. For each year, we keep each employed individual's highest-earning job.[14]

We use establishment closures during 1997–2014. Starting in 1997 gives us 10 years of labor market trajectories before displacement for all subjects. Ending in 2014 gives us sufficient post-displacement outcome years. For every year t, we define *closing establishments* as establishments with at least 5 employees in year $t-1$ that $i)$ disappear completely by year $t+1$ *or* $ii)$ reduce total employment by at least 90 percent until $t+1$. We further require that the establishment had *some economic activity* during year $t = 0$. Partial layoff events are not coded as closures. We remove *false closures*, defined as cases where at least 30 percent of workers involved in an apparent closure moved to a single new establishment, or to other establishments within the original firm (see, e.g., Kuhn (2002)). Workers from such false closures are also excluded from the control group.

In our *displaced worker* sample, we keep workers aged 24 to 60 with at least three years of tenure at a closing establishment in year $t-1$.[15] The age restriction ensures meaningful pre- and post-displacement characteristics for all workers. Importantly, we use a higher upper age threshold than many previous studies since our conjecture (see also Salvanes et al., 2023) is that older workers suffer particularly large earnings losses due to job loss. The tenure restriction ensures that workers are sufficiently connected to the closing establishment, see, e.g., Davis and von Wachter (2011).

The earnings of displaced workers are compared to *control workers*. Apart from the closures, we impose identical restrictions on these workers. Thus, workers in the control group for closures in year t worked at establishments with at least 5 employees in $t-1$, were aged 24 to 60, and had at least three years of tenure. In contrast to the displaced,

[13] Around 10 percent of workers are not employed at a physical establishment due to the nature of their work (e.g. home-care workers). These are treated as working in single-worker establishments.

[14] Individuals are employed if earning at least 3 times the minimum monthly wage from a single employer during a year. Sweden does not have a legislated minimum wage so we use the 10th percentile of the wage distribution, following earlier studies on Swedish register data.

[15] Note that we include the few workers who remain within the original establishment in the cases when the establishment did not fully disappear, but where employment did decline by more than 90 percent. Also, we do not place any restrictions on what the workers do during year $t = 0$. This reduces potential endogenous selection due to early leavers from declining establishments.

their establishments must *survive* until year $t + 1$, with at least 10 percent of the original size. As with the displaced workers, we do not impose any restrictions on future outcomes. From this sample, we select control workers through propensity score matching, i.e. we assume selection-on-observables/unconfoundedness. We match on all variables used for our heterogeneity analysis. As a consequence, we match on a much richer set of variables describing workers, establishments and locations (details below) than previous studies. We estimate year-specific logit regressions to get propensity scores, drop individuals outside of the region of common support, and match three (to gain precision) control workers to each displaced worker. We match without replacement to ensure that we can split the sample into separate independent sub-samples (folds). Assuming unconfoundedness, the matched control group identifies the counterfactual (non-displaced) earnings trajectories of displaced workers. We show that our main results are robust if we use alternative matching strategies.

The main outcome of interest is annual *earnings* normalised by the worker's own earnings in $t - 1$. This outcome allows us to measure relative earnings response without conditioning our sample on future outcomes. The normalization also removes fixed (over time) earnings differences across workers. We also consider a binary outcome for whether workers are *employed* (earn more than three times the monthly minimum wage during the year). For completeness, we also show some results on earnings for the endogenous sub-sample of (re-)employed workers. Outcomes are measured from $t - 3$ until $t + 10$.[16] Geographical and industry mobility is measured by whether the worker lives in a different local labor market or works in a different industry compared to $t - 1$. Industry mobility can only be measured for workers who find new employment.

3.2 Worker, Industry and Location Characteristics

This section gives an overview of the broad set of industry, location, establishment and worker characteristics used in our analyses, focusing on variables often available to policy makers. Details for all variable are in Appendix A. All characteristics are measured in year $t - 1$ unless noted otherwise. *Basic demographics* include age, gender and indicators for first and second generation immigrants. *Family* include indicators for marital status, number of children (total, and school aged), and the worker's share of total household earnings and measures of internal migration history (born outside of the current region

[16] All displaced and controls are observed from $t - 3$ to $t + 1$. From $t + 2$ on, the outcomes are missing if they have died or moved abroad. From $t + 4$ onwards, information for the latest years is missing and the oldest individuals in our sample reach retirement age. Due to sampling restrictions, employment equals one in $t - 3$ through $t - 1$.

and number of cross-location moves). These are intended to capture mobility costs and the value of home production.

General human capital is measured by years of schooling, labor market experience (years employed since $t-10$), and annual earnings separately for the years $t-3$ to $t-1$. *Specific human capital* is captured by establishment and industry tenure before the closure (capped at 10), and variables capturing how specific the field of education is. Field specificity is measured through the share of workers that work in the ten most common industries (by field), and through separate indicators for STEM education and for training in a licensed occupation (e.g. nurses).

Establishment and job characteristics include plant size in the year before displacement, trend in plant size, wage premium at the closing plant[17], a manager dummy, the lost job's routine task component, size of the displacement event as a share of total employment in its industry-location cell, as well as an industry-education match indicator (equal to one if a worker is employed in one of the 10 main industries for their educational field). With regard to *industry* characteristics, we include the industry wage premium, churning rate, excess reallocation rate (Burgess et al., 2000), long run industry employment trend, current industry-specific business cycle conditions, and dummies for manufacturing and publicly funded industries (education, health and public administration).[18] These variables should, for instance, capture industry-level dynamism and growth trends. All industry variables are measured at the three-digit NACE level.

Local characteristics are measured at the level of local labor markets, which are clusters of municipalities created by Statistics Sweden based on commuting patterns. Local characteristics include unemployment rates, population density, exposure to industry-specific trends, cycles, churning and reallocation rates, as well as the share of manufacturing jobs. Local industry concentration is measured by an HHI index across 3-digit industries. Furthermore, we include aggregate characteristics (year of displacement and the aggregate unemployment rate in the year $t+1$) in our analysis.

[17]To avoid making the structural assumptions of the AKM model (exogenous mobility, static wage setting, and additive separability between person and firm effects), we measure the wage premium as coworkers' residual wage in the spirit of Card et al. (2013). The wage is residualized from age, gender and immigration status, as well as education (level and field). We measure the establishment wage premium as the deviations from the industry mean of this residual, and add the industry mean as a separate variable below. For completeness, we verify that including AKM effects in addition to the information we already use does not improve our estimates of heterogeneity in post-displacement losses.

[18]For research on displacements in the public sector, see Eliason (2014b).

3.3 The Sample

About $180,000$ workers in our sample were displaced due to $21,000$ closure events during 1997-2014. Over $4,000,000$ workers at $200,000$ establishments are eligible as controls. Table B.1 in Appendix B presents sample statistics. Comparing the statistics in Columns 1 and 3 for the displaced workers and the eligible controls (before the matching), we see clear differences in terms of industry, size, and gender. Displaced workers are more concentrated in manufacturing and fewer workers are displaced within education, health and public administration. Displaced workers are also drawn from smaller establishments. As a consequence of the industry structure, men are overrepresented among the displaced. The statistics in Column 2 shows that the matched controls, as expected, are very similar to the displaced workers.

3.4 Estimated Average Effects

As a first step, we estimate the average effects of displacement on normalized annual earnings (earnings in $t+1$ relative to earnings in $t-1$). The effects are estimated as average differences in outcomes between displaced and matched controls. Figure 1 Panel A show earnings effects that resemble the existing literature. The short-run impact on earnings is large, and the effects are highly persistent. Displaced workers as a group are far from returning to the earnings and employment levels of control workers even ten years after displacement. On average, earnings of displaced workers drop 24% relative to controls by $t=1$, and are still 8% lower in $t=10$. Panel B shows that this earnings effect to a large extent is explained by a drop in employment, especially, in the short run.

Panel C revisits the earnings effects after conditioning the sample on a positive employment outcome (separately by outcome year). These results suggest that parts of the earnings losses are driven by lower wages. However, the time patterns are difficult to interpret, in particular if the effects of displacement are heterogeneous. The reason is that the identifying samples change across treatment time, both because of declining employment rates in the control group, and because of treatment effects on employment in the treated group. Changes in the estimation sample across time will affect the estimates even if the processes would not alter the balance between treated and controls if treatment effects are heterogeneous.[19] For these reasons, our analysis focuses on earnings, which we can measure for all workers.

[19] We remove individual heterogeneity directly by using differences in earnings relative to $t-1$ instead of using individual fixed effects. This makes the treatment-control contrast more explicit as our focus is on one outcome year. But neither of these approaches help with changes in the sample composition when outcomes are conditional on employment if time patterns, or treatment effects, are heterogeneous.

Figure 1: Average effects of displacement

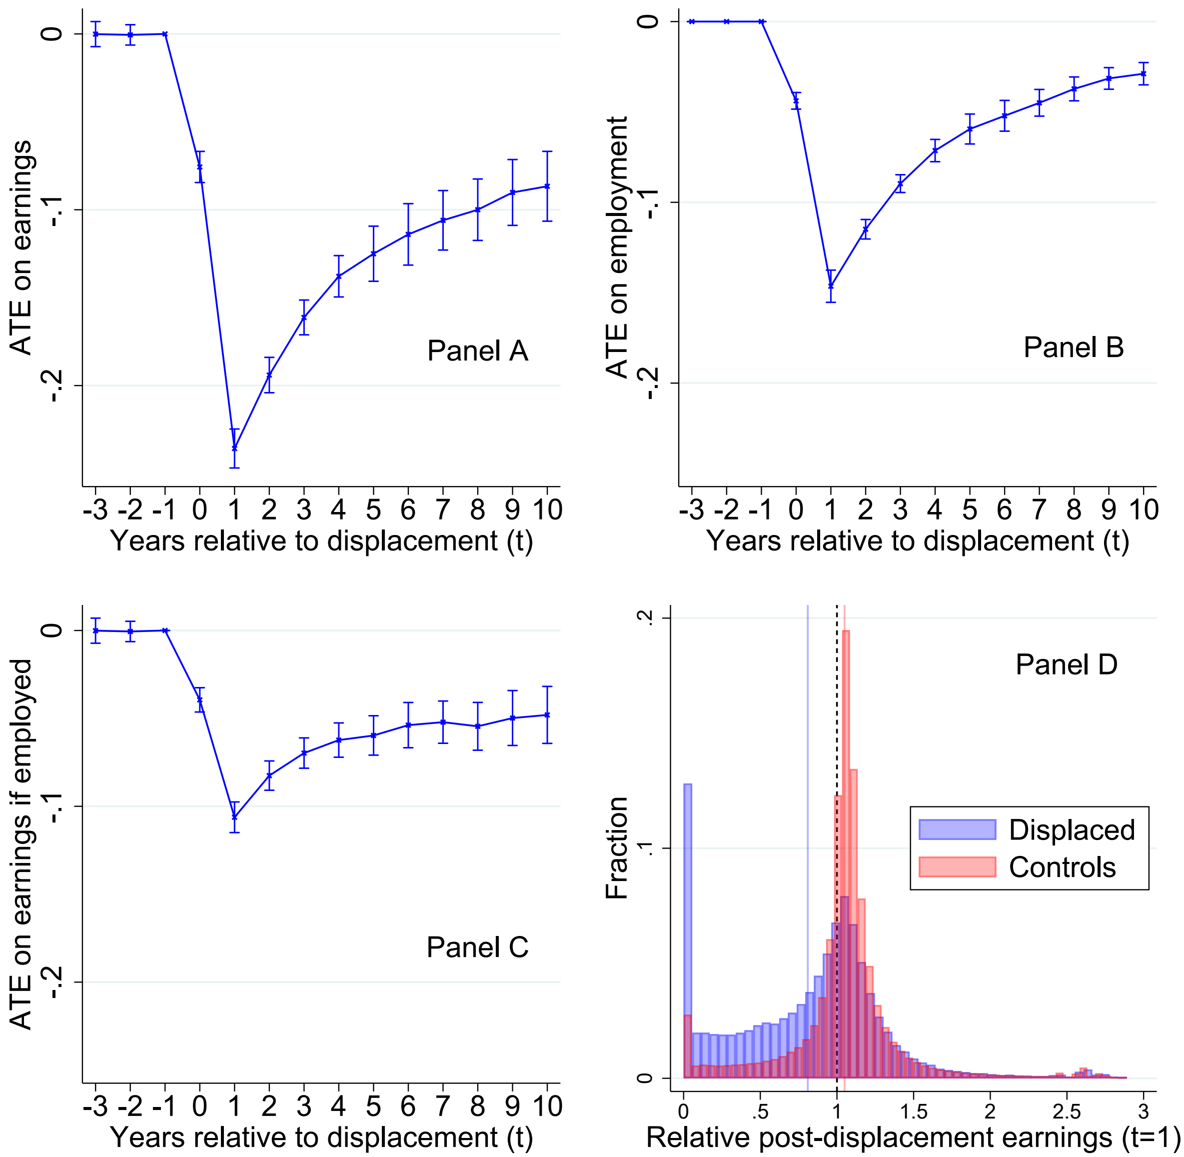

Note: Panels A–C show estimated differences between displaced workers and matched controls (ATE) (with 95 percent confidence intervals). Panel A shows estimates for labor earnings (normalized by earnings in the year before displacement), Panel B for employment status, Panel C for labor earnings conditional on being employed. Standard errors clustered at establishment (pre-displacement) level. Panel D shows distributions of relative earnings in the year after displacement for the displaced and the matched controls. Solid lines indicate group means, the dashed line indicates unchanged nominal earnings.

Figure 1D shows that the *shape* of the earnings-change distribution (in $t = 1$) is altered by displacement, implying that the effects are heterogeneous. The distribution shifts to the left and many workers drop out of the labor market entirely (zero annual earnings). Around 13 percent of the displaced (as compared to 3 percent among controls) have no labor earnings at all in year $t+1$. In addition, there is a uniformly shaped upward shift in the number of workers with earnings above 0 but below 80 percent of pre-displacement annual earnings (the line at 1 indicates unchanged earnings). The increased number of workers with very low (but not quite zero) earnings illustrate that many displaced workers work a limited set of hours during the year after displacement. The number of workers with earnings *increases* above 50 percent is almost unaffected, suggesting that the number of workers who move to a much better job remains unchanged.[20]

4 Heterogeneous Treatment Effects

A conventional approach to study heterogeneous effects is to estimate linear interaction models. We start by estimating such models where our variables, one-by-one in separate regressions, are interacted with a displacement dummy. Estimated interaction terms are reported in Figures 2 and 3. Dummies are kept as they are, whereas all other variables are standardized to mean 0 and standard deviation 1. The estimated interaction terms are, with few exceptions, statistically and economically significant as predictors of heterogeneous displacement effects. The continuous variables with the largest estimates (in absolute values) are age (older workers lose more), schooling (educated workers lose less), routine intensity (routine workers lose more), population density (workers in urban areas lose less), tenure (longer-tenured workers lose more) and industry cycle (losses are smaller in growing industries). The dummy variables with the largest effects are the industry indicators for manufacturing (larger negative effects) and education, health and public administration (smaller effects in publicly funded industries). All of these estimates are in line with previous research, and broadly in line with the theoretical framework we outlined above. In some cases, different variables may plausibly capture the same causal processes (e.g., age and tenure), but in other cases it is more likely that the causes are different (e.g., age and publicly funded industries).

[20] Appendix Figure B.1 presents distributions for small (to avoid spillovers) and stable (to avoid early leavers) closures, with similar results. The same figure also presents distributions for earnings changes in $t+5$ and $t+10$. For these longer horizons, differences between displaced and matched controls are smaller, but the qualitative patterns persist with a large share among the displaced having zero earnings and many workers have relative earnings below one.

Figure 2: Displacement effects interacted with standardized continuous variables

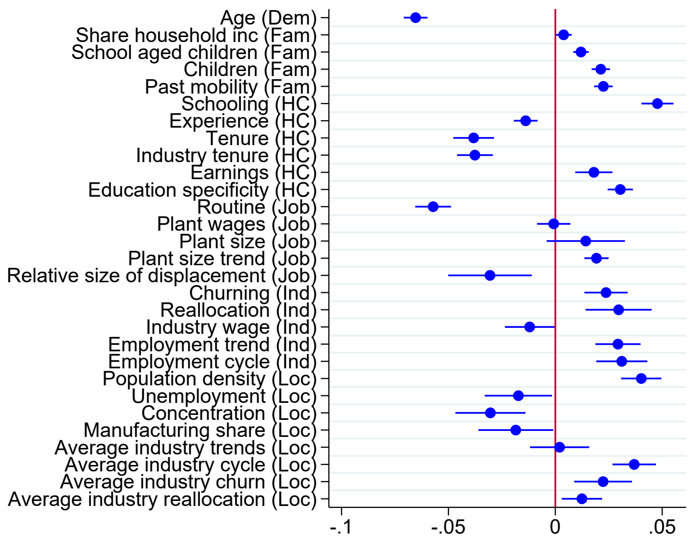

Note: Regression coefficients from separate regressions, one for each variable. Outcome is labor earnings in the year after displacement normalized by earnings in the year before displacement. Explanatory variables are an indicator for displacement, the specific variable (standardized), and the interaction of the two. The graph displays the estimated interaction terms with 95 percent confidence intervals. Standard errors are clustered at the (pre-displacement) establishment level.

Figure 3: Displacement effects interacted with dummy variables

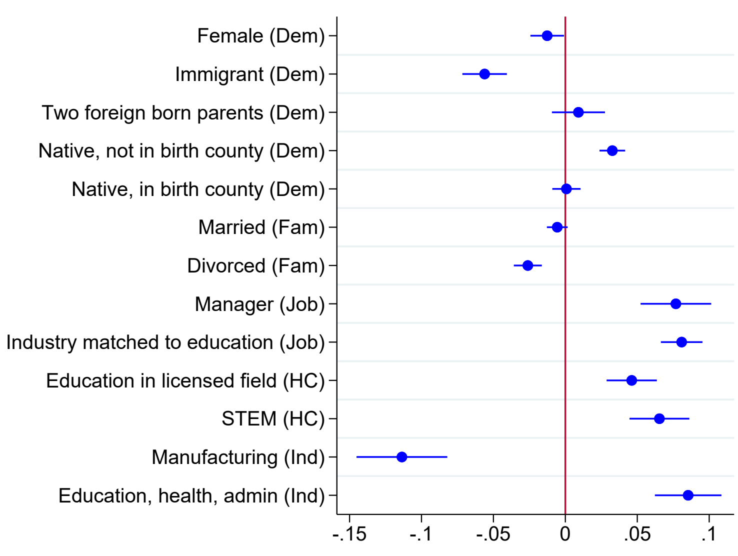

Note: Regression coefficients from separate regressions, one for each variable. Outcome is labor earnings in the year after displacement normalized by earnings in the year before displacement. Explanatory variables are an indicator for displacement, the specific (dummy) variable, and the interaction of the two. The graph displays the estimated interaction terms with 95 percent confidence intervals. Standard errors are clustered at the (pre-displacement) establishment level.

A key challenge is how to handle so many dimensions of heterogeneity. Furthermore, the underlying heterogeneity may be even more complex due to higher order interactions and non-linearities. We could gradually make the linear model richer by introducing more interaction terms and higher order interactions in the same model. To provide a systematic approach to the data, we instead move to the GRF.

4.1 The GRF

To study multidimensional treatment-effect heterogeneity in a flexible way, we rely on the GRF developed by Athey et al. (2019). The forest iterates across random subsets of data, estimating a *causal tree* (Athey and Imbens, 2016) in each subset. Each causal tree is a sequence of splits. Each split partitions the data using one of the x variables. The algorithm chooses the variable and cutoff value which maximize the difference in treatment effects between workers on either side of the split. In more detail, the splitting rule is based on a regression of residual $t+1$ wages on residual treatment assignment, where the residuals come from predictive models of outcomes and treatment that depend on the same covariates used to model heterogeneity. This residual-on-residual regression produces doubly-robust estimates of the treatment which would in principle be valid even if our matching of control workers to treatment workers was imperfect.[21] After each split, the data in the resulting nodes are split again until the workers have been grouped into "leaves" with similar treatment effects. To avoid overfitting, the tree is constructed using part of the selected subset, while treatment effects in each leaf are estimated using the other part. The estimates from a single tree can be non-robust and the GRF therefore computes many trees, with each tree using a random subset of workers and a random subset of x variables at each split. The complete forest is based on an ensemble of the estimated trees, ensuring that estimates are robust across subsamples, and providing consistent treatment-effect estimates.

The GRF allows us to consider heterogeneity across a large number of covariates within a unified model while reducing the risk of overfitting. It can flexibly account for nonlinear effects and high-order interactions (e.g., between industry-specific human capital and industry trends). We use the covariates described in Section 3.2 and estimate the forest with outcome equal to annual earnings in the year $t+1$ after displacement. We show below that our short-run CATE estimates also predict long-run effects.

[21] See (see Nie and Wager, 2021). We use two separate regression forests in the spirit of Breiman (2001) to estimate the conditional propensity score $\hat{e}_i = W_i|x_i$ and marginal response function $\hat{m}_i = y_i|x_i$. The treatment status and the outcome are residualized to obtain $\tilde{W}_i = W_i - \hat{e}_i$ and $\tilde{y}_i = y_i - \hat{m}_i$. The GRF is then estimated using these residualized values.

All steps of the estimation are *clustered* at the establishment level, which is important for avoiding dependencies (leakage) across sub-samples which could lead to overfitting if outcomes are similar for workers who were displaced in the same event. We set aside a test dataset containing 20% of the closing establishments and their associated matched controls for use in the final targeting exercises in Section 7. The remaining 80% of the data is divided into 5 *folds*, containing an equal share of closing establishments and matched controls. We then leave one fold of the data out at a time and estimate a GRF on the remaining folds. The GRF gives us estimates of the conditional average treatment effects (CATE), but we only retain estimates for workers in the *left-out* fold. As a consequence, we do not use any information from the observation's own establishment closure when estimating that observation's CATE.[22]

4.2 GRF Estimates of Heterogeneous Effects

Figure 4A illustrates substantial variation in CATE estimates. All CATE estimates are negative. The distribution is skewed to the right, suggesting that many workers cope reasonably well, but there is a long left tail of workers estimated to suffer very large earnings losses. Since this distribution incorporates both true effects and sampling variation in CATE estimates, we next proceed to our preferred analysis, which evaluates heterogeneity by creating data-driven groupings and then estimating treatment effects within the groupings using distinct, held-out data sets, resulting in consistent estimates of treatment effects within the groups. Importantly, even if our CATE estimates are imperfect, our approach still provides valid estimates of treatment effects for the groups we define, so long as our unconfoundedness assumption holds.

Figure 4B sorts workers into decile groups based on their estimated CATEs from the GRF, with the hardest-hit workers in Decile 1.[23] We then calculate the average effect for the workers in each decile as the difference in outcomes between treated and controls (by design, with similar values of x). We refer to this as the *Average Treatment Effect* (ATE). By splitting workers according to their ranked CATE, we get groups with monotonically improving ATEs. The differences in effects are very large. Displaced workers

[22] All of the forests we estimate contain 2000 trees. The test set used for the targeting analysis is also divided into five folds; each of the GRFs estimated using the 5-fold procedure in the training set is used to predict CATEs for one test set fold. We use the default parameters of the `grf` package in R. These are: fraction of data sampled into each tree = 0.50, number of variables randomly available for each split = $\sqrt{Total \quad number \quad of \quad x \quad variables} + 20$, minimum leaf size = 5 treated and 5 control workers, fraction of sample used for determining splits = 0.5, prune leaves which end up empty when determining treatment effects = `TRUE`, maximum split imbalance = 0.05, soft imbalance penalty = 0.

[23] For each fold of the data, CATEs and the mapping from characteristics to deciles is estimated using only data from the other folds; thus the decile mappings are constructed separately for each fold.

in the least-affected decile (Decile 10) lose only 5.9 percent of their earnings on average. These small effects are (nearly) consistent with a fully competitive labor market without unemployment or job-specific wages. On such a market, job loss should leave earnings completely unaffected. In contrast, losses for workers in Decile 1 amount to 46 percent of earnings. These effects, which are almost 8 times larger than in Decile 10, suggest that labor market imperfections are very prominent for workers in Decile 1.

The relationship between ATE and CATE implies that the GRF captures heterogeneous effects out-of-sample. Figure 4B also allows us to compare the average CATE to the ATE within each decile as a measure of how well-calibrated the GRF is. The ATE grows monotonically across CATE deciles, but the variation across deciles is larger in the actual data (the ATEs) than predicted by the GRF (the CATEs).[24] The GRF thus manages to accurately rank workers across deciles, but underestimates the quantitative differences across these ranks. This pattern caries over to other cuts of the data based on CATE as analyzed below, and we therefore focus on the ATE estimates whenever we discuss magnitudes. Figure 4B also plots "doubly robust" Agumented Inverse-Probability Weighted (AIPW) estimates of the average effect for each CATE decile. The small differences between the AIPW estimates and the unadjusted ATE estimates suggest that our basic matching procedure is sufficient to adjust for differences in observables, and we therefore primarily rely on the unadjusted ATE estimates in the remainder of the paper.

We use a set of auxiliary exercises to assess the robustness of the estimated heterogeneity. In Appendix Figure B.2, we show that the estimates are robust if we use alternative matching protocols to generate the counterfactuals.[25] To handle concerns about spillover effects from large closures, we re-estimate the ATE:s in each decile using only events with less than 20 workers in Figure 4C, with unchanged results. Our main strategy uses matching on a very rich covariate vector, and a sample selection rule that is lagged one year, to handle potential pre-closure selection out of the closing plants. But to further address concerns regarding early leavers selected on unobservables, we focus on sudden closure events affecting plants with stable employment in the 3 years prior to the closure (Figure 4D) with negligible changes in results.[26] Finally, Figure 4E shows that the underlying processes are stable enough to predict heterogeneous effects of displacements that occurred after our estimation period. Here, we apply the decile

[24]Formally, the best linear predictor test of Chernozhukov et al. (2020) tests if the GRF can predict the correct mean effect, and variation around this mean. In our case, the estimated mean effect is well calibrated (in the full sample, the average CATE is similar to the ATE), but the variation across deciles is underestimated. For the mean we estimate $\alpha = 1.06$ and for the variation, $\beta = 1.57$. Both are significantly different from zero at conventional levels.

[25]The variations use GRF for the initial matching and place more weight on aggregate variables.

[26]See e.g. Gathmann et al. (2020) on large closures and Seim (2019) on early leavers.

assignments to workers impacted by displacements during 2015 – 16. Recalling that we only consider closures in 1997 – 2014 for estimation of the GRF, and that 2015 is a full decade after the median year in our training sample.

5 Heterogeneous Effects and Other Economic Outcomes

In this section, we show that workers with large one-year displacement losses also suffer from other adverse economic outcomes. Panel F of Figure 4 relate the displacement loss deciles to the estimated *risk of displacement* (using the estimated displacement propensity discussed in Section 3). The figure plots the average risk by CATE decile after normalising the risk to 1 in CATE Decile 1. The relationship is monotonically negative and differences across deciles are substantial; the most resilient workers (in terms of earnings CATE) have a 40 percent lower estimated risk of displacement relative to the least resilient decile.

Figure 5 shows how the effects of displacement evolve over time across the distribution of CATEs. We use displacement events before 2010 to ensure a balanced panel, but this does not affect our conclusions. Panel A plots ATEs over time for CATE deciles 1 (least resilient) and 10 (most resilient), as well as for the 10 percent of workers straddling the median. Short-run CATEs are strongly predictive of displacement losses throughout our 7-year follow-up period, suggesting that a short-run model is sufficient to capture most of the medium-run heterogeneity. After 5 years, earnings losses in Decile 1 are four times larger than in Decile 10, as compared to nine times larger in the first post-displacement year. Aggregating across the 7 years, the worst-affected decile is estimated to lose labour income corresponding to 1.98 years of pre-displacement labor earnings. This is five and a half times more than the loss of the least affected decile (0.36). To further illustrate the tight link between heterogeneous effects at different horizons, Appendix Figure B.3 show ATE estimates for $t + 3$ in each decile (estimated from $t + 1$ data) and the estimates grow monotonically across deciles.[27]

Panels B to D in Figure 5 show outcomes over time relative to the event for treated and controls, separately for the top and bottom deciles as well as the median group. As the controls provide the counterfactual outcome for the treated, the differences between these series are identical to the ATEs. Annual earnings are normalized relative to the median CATE group in $t - 1$ to highlight pre-existing earnings differences. Pre-trends are perfectly matched within each decile even though they differ across deciles.

[27] Estimating a separate GRF for $t + 3$, gives CATE estimates that are highly correlated with the CATE estimates for $t + 1$ we use to classify workers here ($\rho = 0.83$).

Figure 4: CATE, ATE, and the risk of displacement

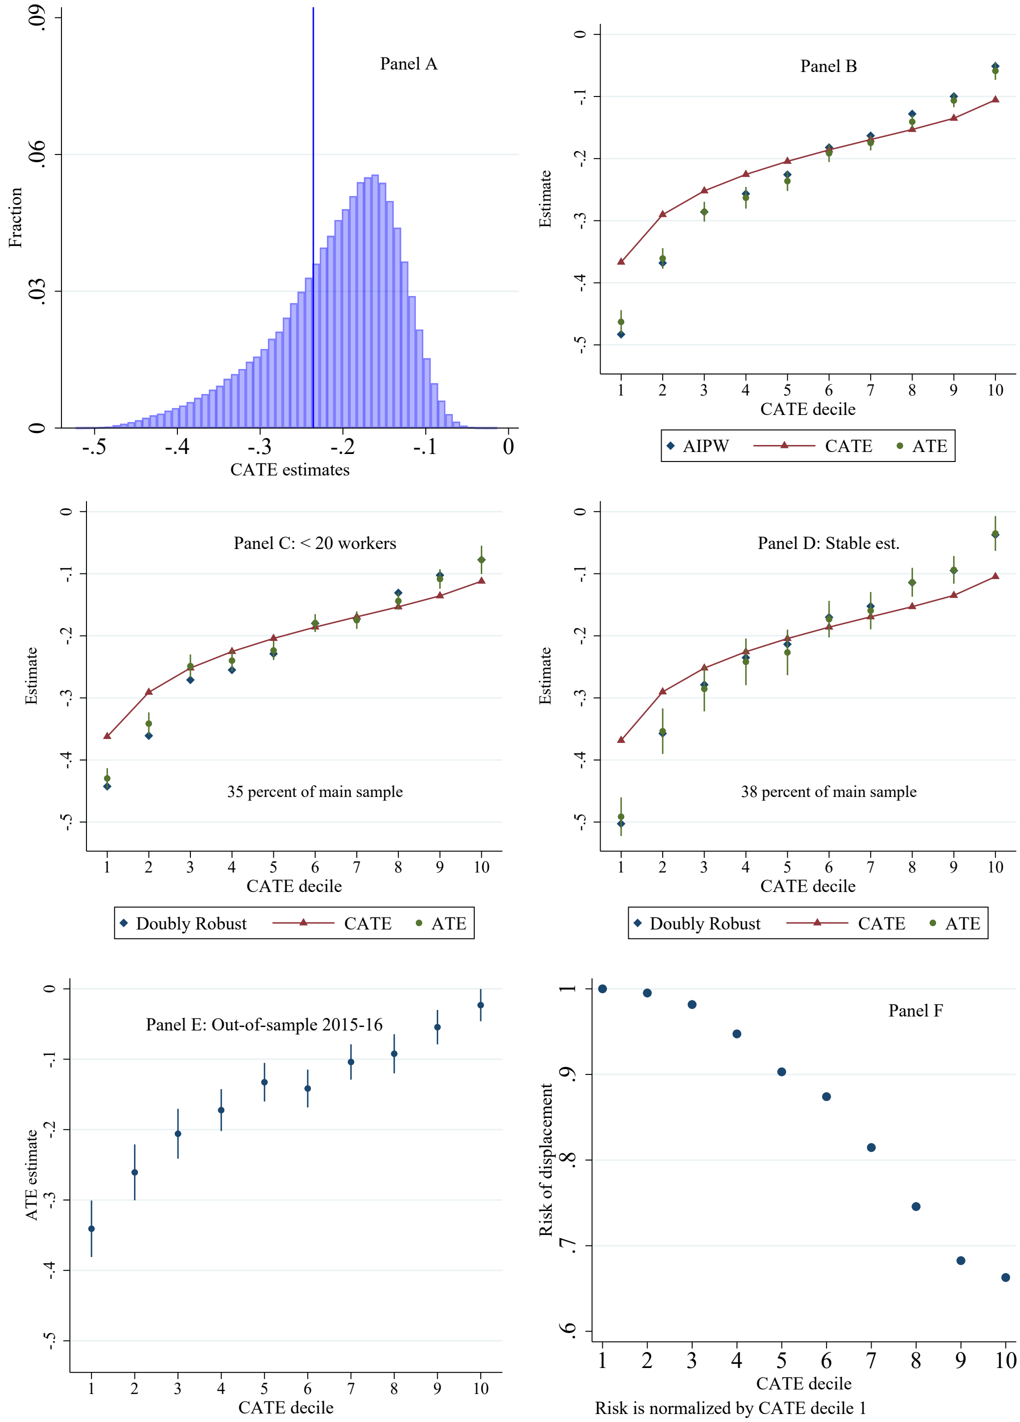

Note: CATE:s for the main data set estimated using 5-folds estimation. Ranking of CATE:s for each fold separately. Panel A shows a histogram of CATE:s (line at estimated ATE for our sample). Panel B shows average CATE:s, ATE:s and AIPW scores in deciles of CATE:s for relative earnings in t+1. For the ATE:s, we report point estimates and 95 percent confidence intervals with standard errors clustered at the establishment level. Panels C and D show analogous calibration results across CATE deciles in sub-samples consisting of establishments with fewer than 20 workers (Panel C) and with absolute employment growth between t-3 and t-1 below 10 percent (Panel D). Panel E shows ATE estimates for those displaced during two years *after* our estimation period (2015-16). Panel F shows the average estimated propensity of displacement by CATE decile. Scores normalized relative to CATE decile 1 to highlight relative risk. Propensity scores estimated by logit on the full sample of workers.

Figure 5: Effects of displacement across time and CATE deciles

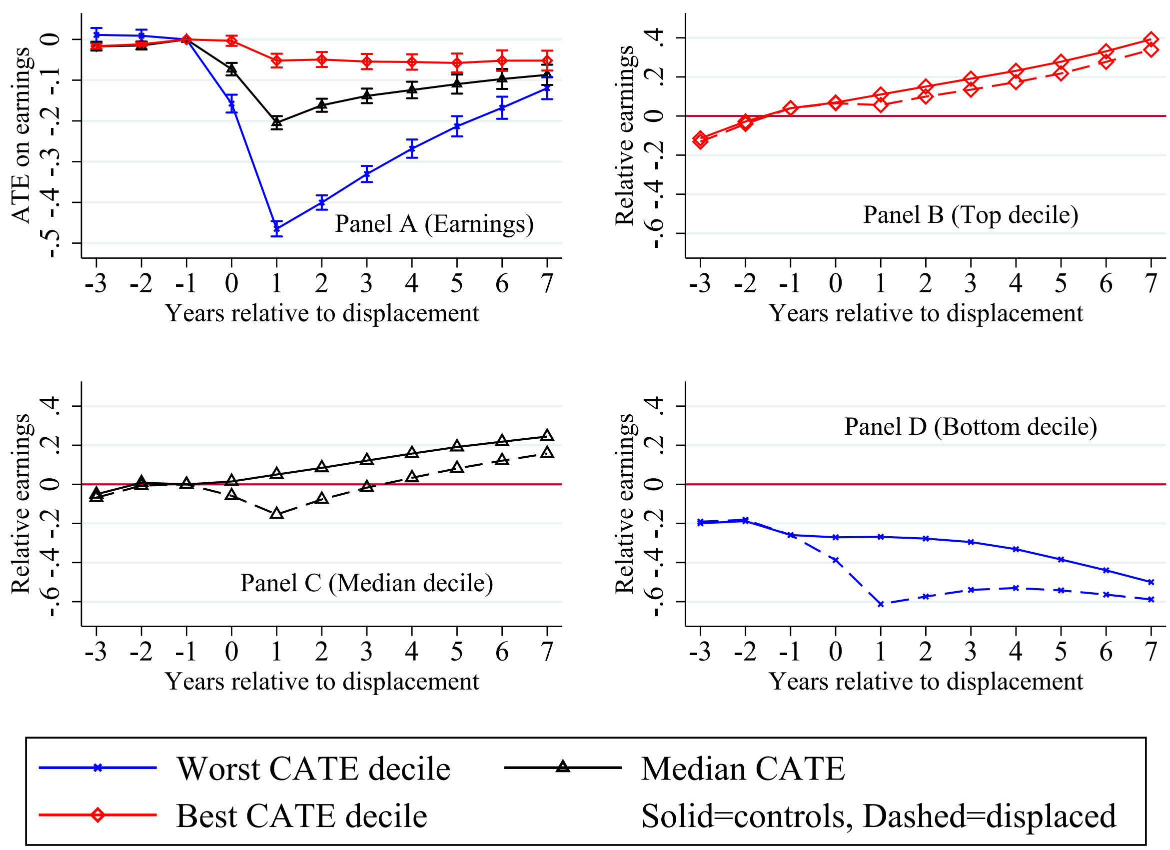

Note: The figure shows statistics for three decile groups of the CATE distribution. The "median" group straddles the median (i.e. it contains the 10*th* and 11*th* ventile). Panel A shows ATE effect over time, similar to figure 1, but separately for the decile groups. Point estimates and 95 percent confidence interval with standard errors clustered at the establishment in $t-1$. Panel B–D show the underlying earnings trajectories for displaced and matched controls within each decile group for the top, bottom and median deciles respectively. The series are normalized to reflect differences relative to the median group in $t-1$. Only workers observed in each of the periods $t-3$ to $t+7$ are included (thus we exclude 2011-2014).

A priori, we could have expected workers with a more positive counterfactual trajectory to suffer larger effects since they have more to lose, but the results point in the very opposite direction. The large effects for low-decile workers arise because of poorer outcomes for the treated, not because of better outcomes among controls. Workers in Decile 1 already had lower earnings than the median before displacement *and* a much less favorable counterfactual (non-displaced, control) earnings trajectory. Decile 10, on the other hand, experiences smaller effects, but also above-median pre-displacement earnings and more positive counterfactual earnings trajectories. The figure also illustrate a key difference in terms of the catch-up over time – much of the declining effects for Decile 1 is related to falling earnings in the control group, which is not the case in the other deciles.

Figure B.4 in the appendix replicates Figure 5 with employment as the outcome. The results are very similar, which highlights the importance of the employment margin.[28]

The results of Figures 4 and 5 jointly suggest that job loss episodes lead to accelerated gross earnings inequality by causing larger earnings losses for workers who already had lower wages from before the event, who would have had worse earnings trajectories without displacement, and who have a higher risk of displacement. In Section 7, we discuss the role of insurance policies that may mitigate this heterogeneity.

6 Predictors of Heterogeneous Treatment Effects

This section presents estimates of how the impact of displacement varies across groups of workers and market conditions, with the goal of characterizing vulnerable workers. The estimated earnings loss within each split of the data should be considered as an estimated causal effect of displacement for that worker group, but the splits themselves should not be given a causal interpretation.[29]

Figure 6 displays differences in characteristics between workers in the lowest and highest quartiles of estimated CATE:s. Continuous variables, shown to the left, are standardized to mean zero and a standard deviation of unity. Dummies are reported at their true mean values to the right. Solid blue bars are sorted according to the magnitudes of the difference between the bottom and the top quartile. In Appendix B (Figure B.5) we compare the top and bottom deciles instead, with similar (but starker) results.

[28]A key difference when studying employment is, however, that the sample restriction ensures that everyone is employed before the event, which implies that the trends *have to be negative* for all groups in Panels B to D. The declining employment rates also illustrate why we do not want to study outcomes that are measured conditional on employment, as the sample changes across treatment time.

[29]For instance, when comparing displacement effects among workers with different levels of schooling, we will *i*) interpret the estimates within each education group as causal effects and *ii*) claim that the differences in estimates between the education groups describe differences in causal effects. However, we will not claim that the differences only arise because of education *per se*, as education may be correlated with other important attributes, whereof some may be unobserved.

Figure 6: Differences in characteristics across CATE quartiles

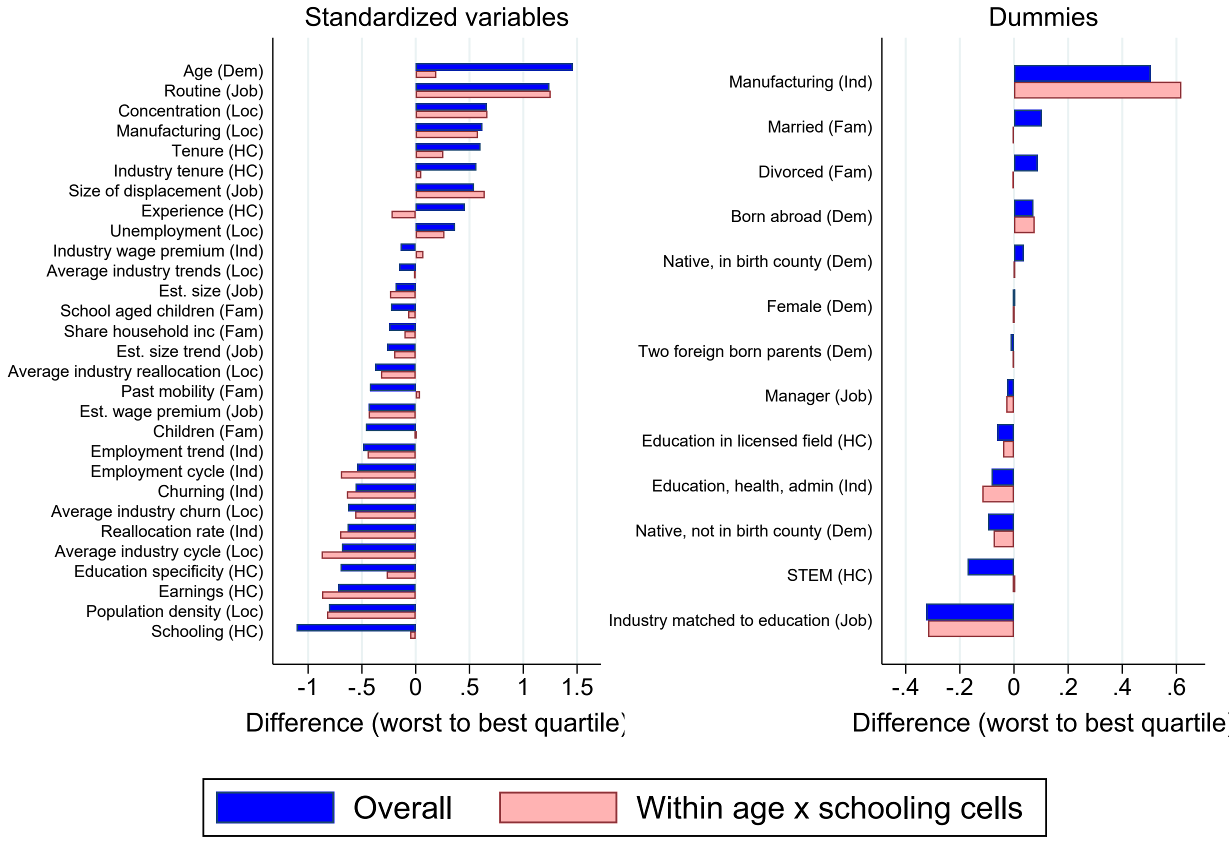

Note: The figure shows differences in characteristics between the lowest and highest quartile of CATE:s, using the main data set. CATE:s estimated with 5-fold estimation and ranking done within each fold. The left-hand panel presents standardized (mean 0 and standard deviation 1) continuous variables and the right-hand panel presents dummy variables. Blue bars are for the overall quartiles and red bars cover the highest and lowest quartiles of CATE within each combination of 8 schooling and 10 age categories (see Figure 7 below).

The characteristics that stand out are the same as in the simple one-dimensional heterogeneity analysis. The quartile with the largest losses contains older workers with lower levels of schooling; these two variables have the largest cross-quartile differences among all continuous variables. Human capital variables (e.g., tenure) and occupation related factors (e.g., job routineness) also differ markedly across the CATE distribution. More vulnerable workers have less prior mobility and are less likely to hold an education in STEM fields. As already documented, vulnerable workers had lower earnings before becoming displaced. But, they also have more sector- and firm-specific human capital, as evidenced by their longer industry and job tenures. Vulnerable workers are also more concentrated in manufacturing industries and more exposed to unfavorable industry characteristics such as bad long-term industry trends and low churn rates. They are more likely to live in rural areas with high unemployment rates and high manufacturing shares and their closure events displace a larger share of workers in their specific industry-location cell.

6.1 Heterogeneity Across and Within Age and Years of Schooling

Motivated by the finding that age and years of schooling differ substantially across quartiles in Figure 6, this section shows that older workers have larger losses due to displacement conditional on schooling, while lower-educated workers have larger losses conditional on age. Since these variables are likely to be correlated, they could, in principle, capture the same underlying factor(s). In Panel A of Figure 7, we plot treatment-control differences (i.e. ATEs) estimated separately for 80 combinations of age and schooling. As is evident from the figure, a higher age is associated with a larger ATE regardless of the level of schooling, and longer schooling is associated with a smaller ATE regardless of age. Young workers with at least a bachelor's degree lose less than ten percent of their earnings, whereas old workers who have not completed high school suffer losses of over 40 percent. Panel B shows AIPW estimates instead. These estimates adjust for any imbalances related to X-variables within each cell but, as expected, the results are almost identical to the ATEs. In the appendix (figure B.6), we show that the corresponding CATE estimates are similar, but less noisy.

In panels C and D we exploit GRF predictions *within* each of these 80 age-schooling combinations. They show ATE estimates among the most (least) resilient CATE quartile *within* each combination. In (almost) all 80 cases, the GRF identifies at least a quarter of workers with ATEs below the overall grand mean, and a quarter with ATEs above the overall grand mean. The worst-affected quartile (as identified by the GRF) of 30 year-old university graduates has larger earnings losses than the least affected quartile of 55-year old compulsory school graduates.

These results illustrate that our CATE estimates accurately uncover heterogeneity (across data folds) even within narrowly defined subsets of the data. It also highlights that policy targeting is complex as job-loss effects vary both within and across worker-types. The red (shaded) bars of Figure 6 illustrate treatment-effect heterogeneity within age and schooling combinations. The bars show differences in means across CATE quartiles defined *within* age-schooling cells. Much of this "within heterogeneity" is related to industry and location characteristics. A more direct illustration of the importance of industry and location characteristics is shown in Figure B.7 in the appendix; it illustrates a strong relationship to manufacturing intensity and population density for the extreme age-education groups of those younger than 30 with at least a bachelor's degree and older than 50 with only 9 years of schooling.

Figure 7: Displacement effects on earnings across and within combinations of age and schooling

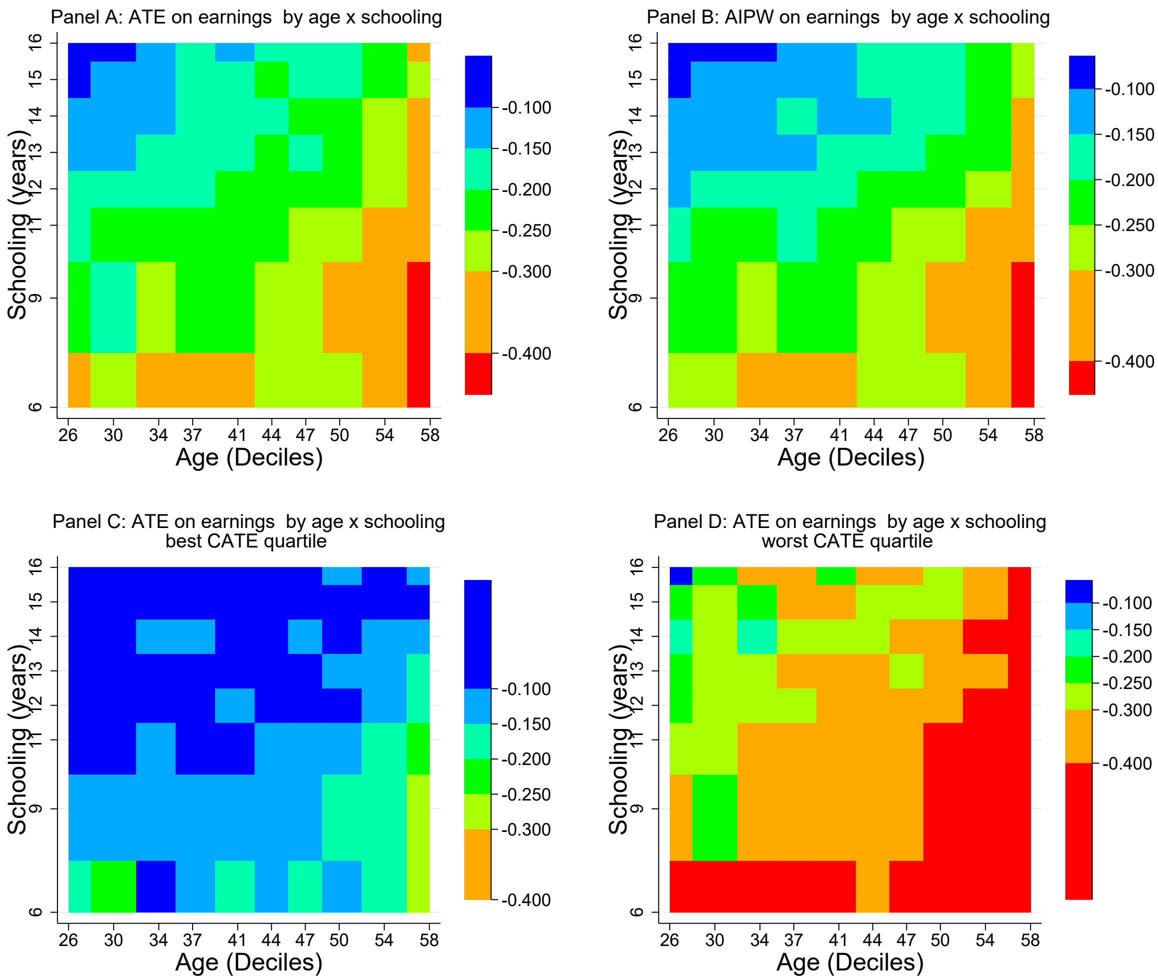

Note: The figure divides training set workers into cells by age and schooling. Schooling has been aggregated to 8 groups by pooling the few with 10 years of schooling together with those with 9 years of schooling, and by letting the top group include all with 16 or more years of schooling. Age is defined in deciles among the displaced and the x-axis shows the median in each age group. Colors indicate the size of point estimates. Panel A shows estimated ATE:s by combinations of age and schooling. Panel B shows AIPW estimates separately within each of the 80 groups. Panel C replicates panel A, but only uses the highest quartile of 5-fold CATE:s within each group. Panel D repeats the exercise for the lowest quartile.

6.2 Heterogeneity Across and Within Establishments

We perform a number of exercises to assess the importance of *establishments* for our estimated heterogeneous earnings effects, and we find very robust evidence for heterogeneous effects both within and across establishments. To this end, we divide the sample of workers in two ways. Initially, we focus on the variation within establishments and divide the sample by their CATE decile *within* each establishment. We also examine the variation *across* establishments, where we for each displaced worker's establishment compute average CATE estimates (using only the CATEs of coworkers) and divide the sample of workers into deciles based on these establishment-level CATEs. We then estimate the average effect of displacement for these within-establishment deciles and across-establishment deciles. Results are shown in Figure 8A (within establishment) and 8B (across establishment). To ensure comparability, we focus on events with at least 10 displaced workers in both of these panels.[30] The amount of heterogeneity documented in panel A is very similar to panel B, which suggests that there is about as much heterogeneity within as across establishments.

Appendix B reports results from corroborating exercises. The results are identical if we use doubly robust AIPW scores (Figure B.8). The patterns are similar if we use longer term outcomes and/or if we study employment instead (Figure B.9). For completeness, we show (see Figure B.10) that effects are identical for high vs. low AKM-establishment effects (Abowd et al., 1999) *within each decile*; this should not be surprising as we include establishment- and industry-specific residual wages in the GRF (i.e. when defining deciles). We further exploit the fact that CATE estimates should be correlated across coworkers from the same event if establishment-specific factors are important. Figure B.11 shows that about half of the variation in CATEs across displaced individuals is shared with displaced coworkers, while the other half is specific to individuals within the same event. Finally, we show that the across-establishment ranking is more important for workers at the bottom of the within-establishment ranking (Figure B.12).

[30] We need to make one further adjustment relative to the main analysis. We let the individually matched controls follow the displaced workers as we do not have estimates of CATE for all workers in the establishments of the control workers. We have verified that we get very similar estimates of average treatment effects by overall decile with this strategy, see Figure B.2A in the appendix.

Figure 8: Displacement effects across and within establishments

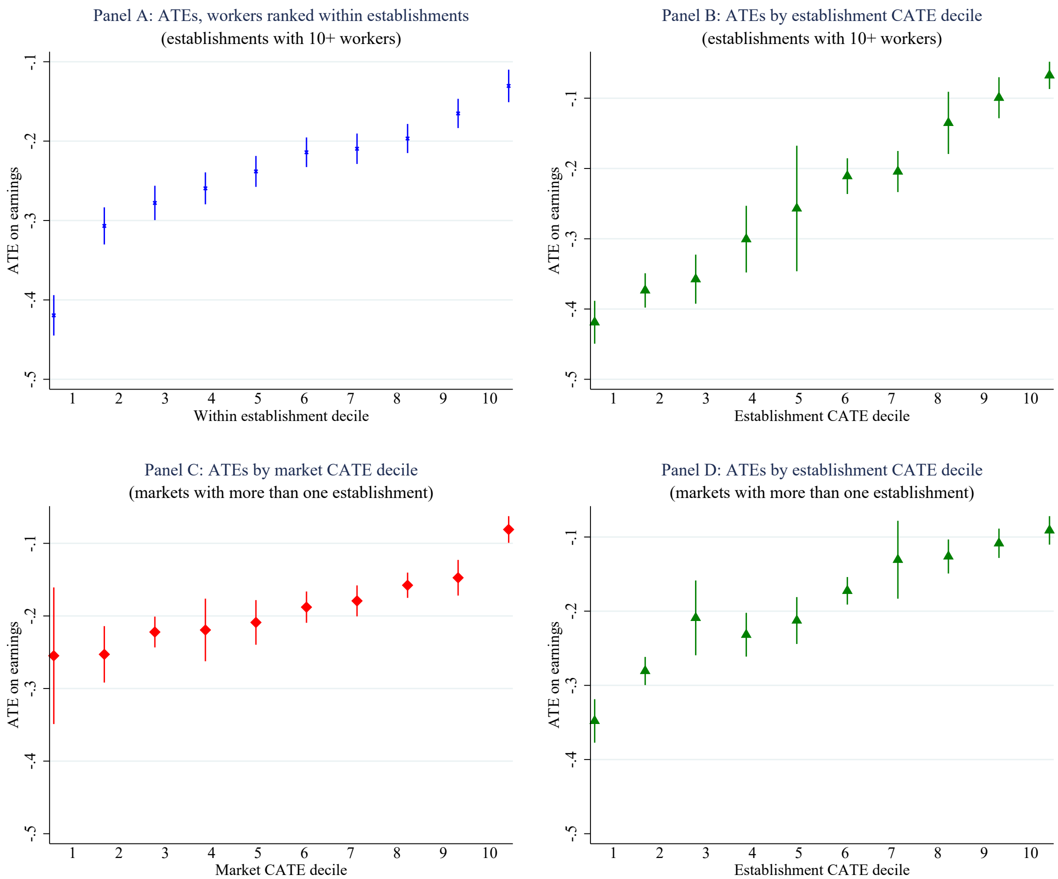

Note: Panel A shows ATE estimates when displaced workers are ranked based on their within-establishment CATE. Panel B shows ATE estimates when displaced workers are ranked based on the CATE of their co-workers (defined as the leave-out mean for the workers at the establishment and then averaging over the individuals in the CATE decile.) Panels A and B exclude establishments with fewer than 10 displaced workers. Panel C ranks the workers by the average CATE in their market (location-industry cell). Only markets with at least two establishment closures are included in panel C. Panel D ranks workers by the CATE of their co-workers like Panel B, but only includes markets with at least two establishment closures. Control workers are allocated to the same sample as the displaced worker they were matched to in all panels.

The predictable heterogeneity we document between workers displaced from the same establishment in Figure 8A *must* be attributed to individual-level factors that vary within establishments. On the other hand, CATE differences across establishments (Figure 8B) are based on all variables that are shared among workers in the same establishment. Establishment-level CATEs will therefore be good predictors of treatment effects if *either* establishment-level variables (the wage premium, size etc.) are important, *or* if important individual/job-level variables are correlated within establishment (sorting), *or* if market-level factors matter (since industry and location are fixed within event).

In Panels C and D of Figure 8, we show that the variation in displacement effects across markets (industry times location) appear to be of similar magnitude as the variation *across* establishments. When characterizing markets, we use the average CATE among workers who are displaced in *other* events in the same market. Results for splits by market CATE are shown in panel C. The exercise forces us to exclude singleton events, i.e. events where just one establishment is closed down in a market, and to facilitate a fair comparison we recalculate the across-establishment estimates for this sample in panel D.[31] The results show similar heterogeneity when ranking workers by the market CATEs and when using the across-establishment ranking of the workers. Note that this holds even though we rank workers by the CATEs of *other* events on the same market.

6.3 Heterogeneity Across and Within Markets

We have seen that much of the heterogeneity within age and education combinations is related to local labor markets or industries, and that market-level factors appear to be highly predictive of displacement effects. We now show that market conditions matter more for vulnerable workers. To this end, we select the best, worst and median market deciles from Panel C of Figure 8, and then examine how the displacement effects vary across workers deciles (ranking workers within establishments as in Panel A of Figure 8) for these three market conditions. The results in Figure 9 show that vulnerable workers experience worse displacement effects under all market conditions, but the differences across worker deciles are substantially larger when market conditions are unfavorable.

[31]Sample differences explain why results in Panels B and D are not identical.

Figure 9: Displacement effects across workers for different market conditions

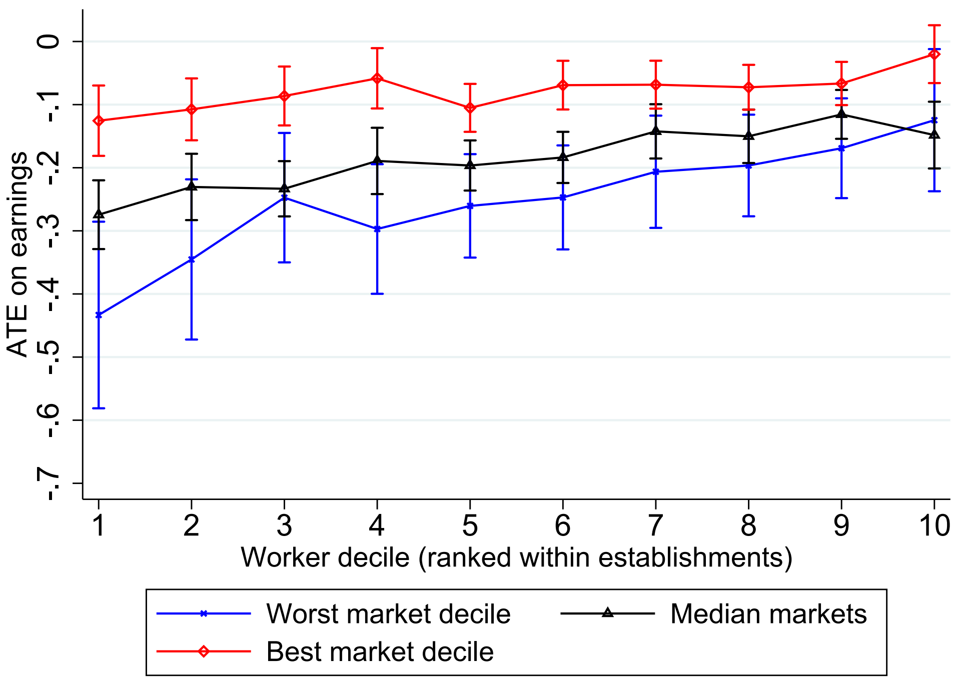

Note: The figure shows ATE estimates when displaced workers are ranked based on their within-establishment CATE for different markets, where markets are ranked by the average CATE in their market (location-industry cell). Establishments with fewer than 10 displaced workers are excluded. Control workers are allocated to the same sample as the displaced worker they were matched to in all panels.

The importance of market-level factors is important from a policy perspective since policymakers are more likely to use place-based or industry-related policies than policies targeting establishments. We can also shed light on the most important industry and location variables. However, isolating the importance of market-level factors is not trivial since workers are likely to be sorted across regions and industries. For that reason we use the GRF to estimate partial dependence functions, holding the individual characteristics at their empirical levels and sequentially rotating across all observed sets of market characteristics within the same year (methodological details are in Appendix C). This way, we capture the predicted role of aggregate conditions across the full distribution of displaced workers, accounting for nonlinearities.

The results from this partial procedure, presented in Table C.1 in Appendix C, confirms that market conditions are important, and that market conditions are particularly important predictors for workers whose individual characteristics are predictive of large negative displacement effects. We also see that workers are more likely to change location if displaced under worse conditions, and workers are more likely to move away from their industry if they are displaced from an industry with large displacement effects. However, even though we find that industry mobility responds to displacement during distressful conditions, the fact that industry characteristics in general correlate so strongly with displacement effects suggests that the degree of industry mobility is insufficient to offset the negative effects of being displaced in a bad industry.

Table 1: Location and industry characteristics for the worst and best locations/industries

	(1) Worst Quartile	(2) Best Quartile	(3) Interquartile Difference	(4) Standard Error of Difference
Panel A: Location characteristics, by location quartiles				
Population density	20.410	147.913	-127.503	3.026
Unemployment rate	0.105	0.062	0.043	0.003
Industry concentration (HHI)	0.040	0.025	0.015	0.001
Share manufacturing	0.230	0.106	0.124	0.016
Average industry trend	0.058	0.113	-0.055	0.005
Average industry cycle	0.005	0.011	-0.006	0.001
Average industry churn	0.203	0.268	-0.065	0.004
Average industry reallocation	0.144	0.161	-0.017	0.002
Panel B: Industry characteristics, by industry quartiles				
Employment trend	-0.150	0.215	-0.366	0.061
Employment cycle	-0.033	0.024	-0.057	0.006
Industry wage premium	0.067	-0.049	0.116	0.027
Churning	0.177	0.253	-0.076	0.021
Reallocation rate	0.116	0.164	-0.048	0.008
Manufacturing	1.000	0.000	1.000	0.000
Education, health, admin	0.000	0.240	-0.240	0.113

Note: Panel A displays average location characteristics for the best and the worst quartiles of locations, and Panel B displays average industry characteristics for the best and the worst quartiles of industries. For details on the classification method, see Appendix C. Both panels use the partial-effects procedure described above. The location and industry characteristics are described in Section 3.2. Standard errors in Panel A clustered at the location level, and standard errors in Panel B clustered at the industry level. All estimates use the main data set.

Table 1 describes the characteristics of the best and worst locations and industries (top and bottom quartiles) after we have partialed out differences in worker characteristics across locations and industries (see Appendix C). Locations in the bottom-quartile are in particular characterized by much lower population density. These locations also have high unemployment rates and a more concentrated industry structure dominated by declining industries and manufacturing jobs. Industries with large predicted earnings losses are exclusively found in manufacturing, while industries with small predicted losses are in non-manufacturing sectors. Bad industries also have higher wage premia, are less dynamic (lower churning and reallocation rates), and experience declining employment trends over both the short and the long run. All of these attributes are typical of manufacturing, but they are also related to effect heterogeneity if we only compare different manufacturing industries to each other, or if we only compare non-manufacturing industries to each other (see Table C.3 in Appendix C).[32]

7 Policy Targeting

Given the severe average impact of displacement, policymakers have designed various forms of policies to ameliorate the consequences of job loss and/or to prevent it from occurring (e.g. through short-time work schemes or employment protection legislation). But supporting all groups of workers equally may be misguided if some groups of workers suffer only very small losses. As shown by our analysis above, targeting based on a single variable is unlikely to capture the complex heterogeneity pattern uncovered by the GRF. But for policy purposes we would like to find usable (i.e. simplified) targeting principles that would direct policymakers toward the right set of workers.

As a first step, we consider key one-dimensional predictors at the worker level (age, schooling), industry level (manufacturing) and location level (population density), and pairwise combinations of these. To make further progress, we then select targeting variables using optimal policy trees (Athey and Wager, 2021). We consider trees of depth two, selecting the two best targeting variables from our full set of variables. Each set of variables defines a targeting rule which we use to extract a fraction of targeted workers (varied between 5 and 25 percent of the overall sample) and estimate the displacement effects on earnings one year after displacement (ATEs) for the selected group of workers.[33] For targeting using the full GRF, this implies selecting the workers with the most negative CATEs, for age it means selecting the oldest workers, and so forth. To obtain a fair comparison between the GRF and the simple targeting rules we use our held-out test set, i.e. 20 percent of the closing establishments and their matched controls not used for producing any of the results shown so far.

The results are reported in Figure 10. The blue lines in all panels show the displacement effect among a randomly selected set of workers. This results in an average earnings loss of 22.2 percent regardless of how many workers we select. On the other extreme, when using the full GRF model which uses all covariates, we find, as expected, a group of workers with much larger displacement effects, as shown by the red line.

[32] Appendix C also discusses partial effects for worker attributes, net of the industry and location conditions. A main result is that the small gender differences shown in our main analysis arise because women tend to work in markets with smaller effects. Holding market conditions fixed, the GRF predicts larger effects for women.

[33] Whenever there is a tie, we select a random set within this tie.

Figure 10: Displacement effects for workers selected by different targeting policies

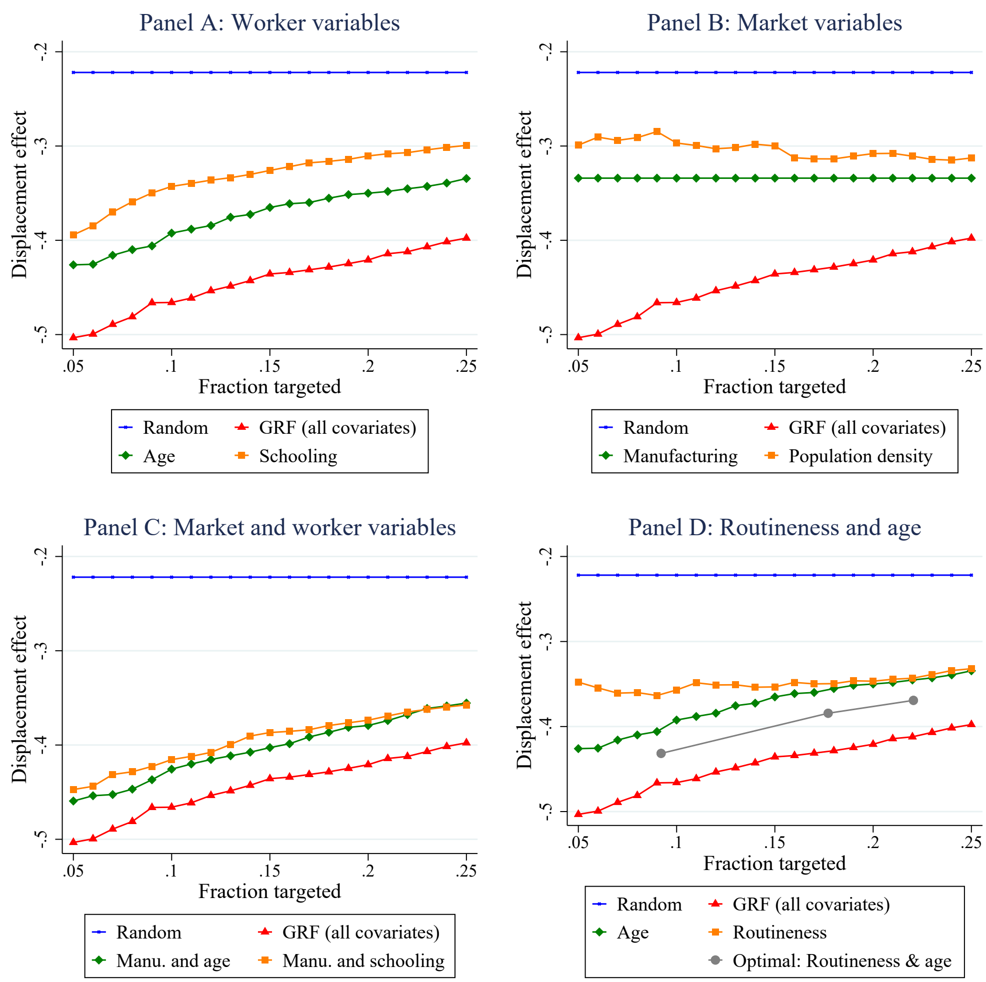

Note: The figure reports estimated observed displacement effects on earnings one year after displacement (calculated as displaced-control differences, "ATEs") when selecting a fraction of the workers using different targeting rules. The random rule selects a random set of workers. The GRF selects workers using the estimated CATEs. Panel A: the age rule targets the oldest workers first and the schooling rule the least educated first. Panel B: the manufacturing rules targets a random set of manufacturing workers, and the population density selects workers in the least dense locations first. Panel C targets the oldest respectively the least educated manufacturing workers. Panel D reports targeting results when targeting the oldest workers (age) and workers in the most routine occupations. It also reports results when targeting on age and routineness using optimal policy rules (Athey and Wager, 2021).

The five percent of workers who are hit the hardest according to the GRF experience displacement losses of 50 percent. Among the hardest-hit 25 percent, earnings losses are 40 percent. Panel A also includes targeting based on age (target the oldest first) and schooling (target the least educated first). The green line shows that average earnings losses are 43 percent among the oldest five percent of workers, and 33 percent among the oldest 25 percent. Targeting based on schooling, shown by the orange line, reaches workers who are on average less hard-hit than the oldest workers (e.g., losses of 39 percent when targeting the five percent least educated). We also see that targeting one variable (age or education) is substantially better than random, but also clearly worse than targeting based on the output of the GRF. For presentation reasons, Figure 10 does not include standard errors, but GRF-based targeting is significantly better than random, age-based and schooling-based targeting at conventional significance levels. The same holds when GRF-based targeting is compared to other simple targeting rules.[34]

Panel B of Figure 10 considers targeting on industry and location characteristics. If manufacturing workers are selected at random, the losses of the workers reached by the policy are around 33 percent (green line). Since manufacturing workers constitute more than 25 percent of our sample, this average loss is the same when targeting the hardest-hit 5 percent and the hardest-hit 25 percent. However, Panel C shows that it is possible to improve targeting by selecting the oldest or least educated manufacturing workers. The first five percent of individuals selected according to these rules suffer earnings losses of 46 percent; if a quarter of the individuals in the test sample are selected, their losses amount to 36 percent. Targeting workers in locations with low population density (shown by the orange line in Panel B) does slightly less well in terms of identifying the hardest-hit individuals than targeting based on manufacturing.

In Panel D, we explore targeting rules based on optimal policy trees (Athey and Wager, 2021). Specifically, we assume that we have a fixed per-person cost of the policy and that the benefits of the policy are proportional to the income losses after displacement. We consider different costs of the policy in order to generate different fractions of targeted workers (the figure shows the preferred targeting of three optimal policy trees which assume different costs of the policy). When choosing between all of our variables, the preferred rule involves targeting based on age and the job's routine content. For a policy cost that results in about 10 percent treated, the optimal policy is to target everyone older than 59, and to target workers aged 51–58 if the level of routineness in their jobs is above the 80th overall routineness percentile (corresponding to the 69th percentile in the sample of displaced workers). Most likely, age is selected because age captures losses related to occupation-specific and establishment-specific human capital, as well as lost job-ladder related rents in the pre-displacement job. Routineness may capture exposure to both declining demand (e.g., related to the manufacturing industry) and low education. The gray line in Panel D shows that using information on both routineness

[34] To examine these differences in more detail, we have also explored the rank-weighted average treatment effects (RATE) metrics from Yadlowsky et al. (2021), with results for the Qini coefficient metric in Table B.2 in Appendix B. It shows that prioritization based on the CATE:s from the full GRF performs significantly better than a random allocation, confirming the presence of heterogeneous treatment effects. Targeting based on the GRF model also outperforms targeting on age or schooling, but these simple targeting rules perform significantly better than a random allocation. We obtain similar results for all the other simple targeting policies evaluated below.

and age improves targeting compared to targeting based on age or routineness only, but still performs worse than targeting based on the entire GRF.

We conclude that simple targeting policies give better results than randomization, but do not quite reach the performance of the GRF model. This confirms our key finding that the size of displacement losses cannot be explained by any single variable in isolation as it takes at least two variables to come close to the GRF. The GRF's ability to take many characteristics into account in a flexible way is key in this context. If GRF estimation based on all characteristics in our analysis is unfeasible, a targeting rule that combines information on routineness and age works best.

7.1 Targeting and Existing Redistribution Policies

Public tax and transfer systems are designed to mitigate the pass-through from gross labor earnings to after-tax income by means of, e.g., unemployment insurance and social assistance, and progressive taxation. The degree of insurance therefore tends to be larger for workers who experience large gross earnings losses. In addition, as outlined in our theoretical framework, workers could experience larger earnings losses because they are well-insured in case of unemployment. Both of these channels suggest that the degree of insurance should be positively correlated with the size of the earnings effect. To illustrate this, we examine the impact on disposable income by labor-earnings CATE decile in Appendix Figure B.13.[35] As expected, disposable income losses are much smaller than labor income losses, in particular in Decile 1 where the degree of implied insurance is 71 percent (as compared to around 50 percent in Deciles 6 to 10).[36] Thus, insurance responses mitigate the impact on inequality, but instead shift parts of the financial burden from the most affected workers onto public finances.

Because insurance policies tend to focus on responses at the bottom, it is not obvious that the groups that suffer from the largest gross earnings losses always lose more in terms of disposable income. To study this issue, we estimate a separate GRF for disposable income. The correlation between CATEs for the two outcomes is positive, but moderate ($\rho = 0.37$). In Appendix Figure B.14 we show the interquartile ranges of characteristics of workers as a function of disposable income losses. Overall, large disposable

[35]Disposable income is calculated by Statistics Sweden at the household level as some support systems give money to households, not individuals. The income is then attributed to household members according to a fixed formula. Thus, the variable also captures within-household income pooling with other household members, but our results are very similar if we focus on singles.

[36]Adding taxes to our theoretical framework, we can write the disposable income effect as $\tilde{\Delta} = (1-q)b + qE((1-tax_k)W_k) - (1-tax_j)W_j$ and the degree of insurance as $(\Delta - \tilde{\Delta})/\Delta$ where Δ is the treatment effect on gross earnings.

income losses are relatively more prevalent higher up in the income distribution.

Figure 11: Displacement effects on earnings and disposable income over time and across earnings CATE deciles

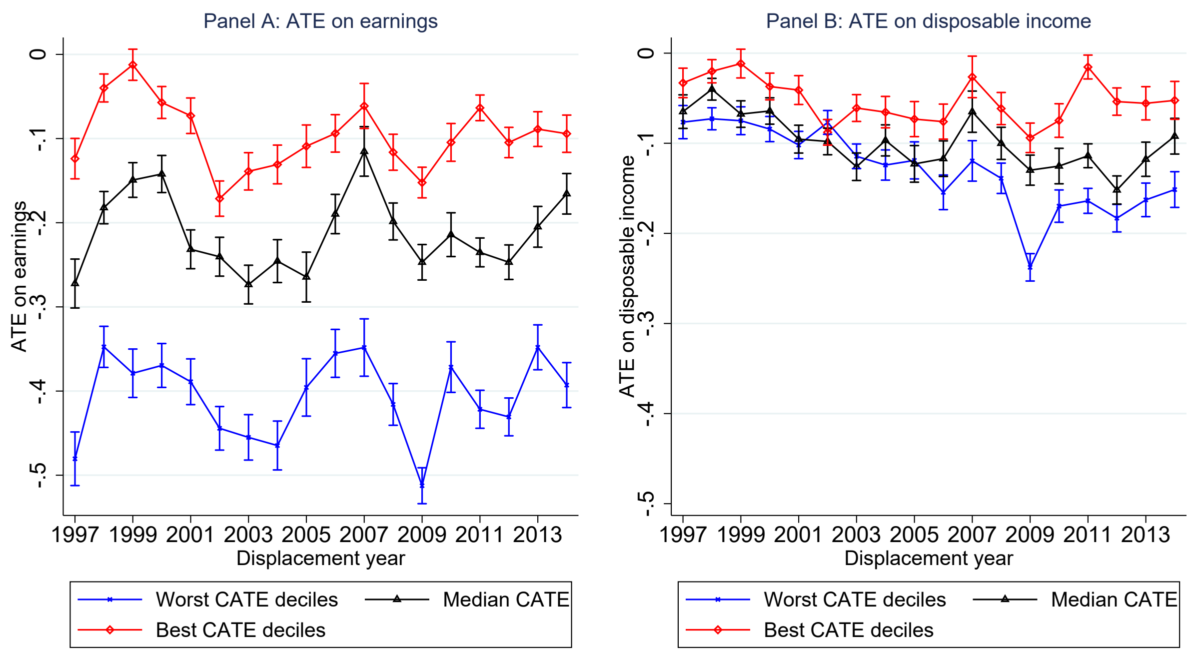

Note: The figure shows estimates by displacement year and deciles of the CATE on earnings in $t+1$. CATE deciles defined within years. Panel A shows ATE estimates on labor earnings in $t+1$ and Panel B ATE estimates for disposable income in $t+1$. Worst CATE deciles includes the first and second CATE deciles, Median CATE the fifth and sixth CATE deciles, and the Best CATE deciles include the ninth end tenth CATE decile.

An interesting evolution, which our data allow us to shed new light on, is the gradual erosion of the generosity within the Swedish Public Insurance System (Skans et al., 2017). The generosity declined as benefit caps and benefit floors remained fixed in nominal terms over most of our sample period, while nominal (and real) wages grew considerably. In Figure 11, we show how the impact of job loss on labor earnings and disposable income evolved over time across the distribution of predicted earnings effects (CATEs for earnings in $t+1$). The effects on labor earnings are relatively constant across the period, but the patterns related to effects on disposable income diverge very strongly. As is evident from the graph, the pass-through to disposable income has increased substantially in the group with the largest predicted earnings losses. This implies that the gradual policy changes that occurred during the study period eroded the degree of insurance particularly among workers who experience large negative effects on gross earnings when displaced.

8 Conclusions

The results presented in this paper contribute to the vast displacement literature by dissecting how heterogeneous earnings effects can be used by policy makers or caseworkers if they wish to target workers who suffer from particularly adverse effects when their jobs are destroyed. Our results show that groups with large negative effects are vulnerable in other economic dimensions as well: they have lower pre-displacement earnings, a higher displacement risk, and much worse counterfactual earnings trajectories than other displaced workers. This suggests that job loss leads to accelerated gross earnings inequality, which may motivate well-targeted policy interventions.

The existing literature has discussed a host of explanations for why post-displacement losses may vary, highlighting factors such as firm, industry and occupation-specific human capital, location mobility and aggregate conditions. Our results confirm that determinants of earnings losses after displacement are multidimensional, as substantial systematic heterogeneity remains even after conditioning on crucial predictors such as age and schooling, or on the closing establishment. While older workers with lower levels of general human capital lose more on average, even old and low-educated workers can do better than the average worker if other circumstances are favorable. Aggregate demand-side conditions at the local and industry level play an important part in predicting the size of losses as workers in urban areas and outside of manufacturing consistently fare better than other displaced workers.

The complexity of heterogeneous effects makes policy targeting challenging, but also implies that policy-design details may affect whether or not a policy reaches the intended target group. Our results show that targeting on age, education, manufacturing indicators or population density is a substantial improvement over a random allocation if the aim is to target workers with large effects. But the overall heterogeneity is too complex to be fully approximated by one-dimensional targeting rules. The targeting rule which comes closest to the targeting recommended by the GRF model involves focusing on older workers in routine occupations, which is a group which suffers from both detrimental individual and market-level conditions.

At a more general level, our results indicate that the economic value of jobs is highly heterogeneous. Some groups of workers, primarily the young, well-educated and those who work in non-manufacturing industries, are able to move on to a new job with only short-term and marginal losses if their firm closes down. This pattern is consistent with an (almost) fully competitive labor market without unemployment or job-specific wage rents. Other groups of workers, such as the older and less educated, instead suffer very large and prolonged losses if their jobs disappear, suggesting that these workers are exposed to highly imperfect labor markets. The interaction between worker and market characteristics should also be understood in this context. Workers with favorable individual characteristics experience small effects of job loss regardless of market conditions, whereas workers with large average effects are much more sensitive.

Our analysis also points to a number of interesting avenues for future research where predictions from the GRF model can help us understand other aspects of the economic and social consequences of job loss. The predictions can be used to analyse a wider set of labor market outcomes beyond earnings, which have been at the focus of this paper. Furthermore, the predicted heterogeneity is likely to be useful when analysing different back-to-work pathways such as further education or self-employment that members of different groups may rely on. Most notably, the GRF model can be used to scrutinize the targeting embedded in existing support policies, as illustrated by our analysis of redistribution policies in the Swedish context. This approach can be extended to asses how policy targeting varies across specific policies such as benefit schemes, short-time work schemes and life-long learning policies.

References

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

Altonji, J. G., E. Blom, and C. Meghir (2012): "Heterogeneity in human capital investments: High school curriculum, college major, and careers," Working paper 17985, National Bureau of Economic Research.

Arntz, M., B. Ivanov, and L. Pohlan (2022): "Regional Structural Change and the Effects of Job Loss," Discussion Paper 22-019, ZEW - Centre for European Economic Research.

Athey, S. and G. Imbens (2016): "Recursive partitioning for heterogeneous causal effects," *Proceedings of the National Academy of Sciences*, 113(27), 7353–7360.

Athey, S., J. Tibshirani, and S. Wager (2019): "Generalized random forests," *The Annals of Statistics*, 47(2), 1148–1178.

Athey, S. and S. Wager (2021): "Policy Learning With Observational Data," *Econometrica*, 89, 133–161.

Autor, D. H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 103, 1553–97.

Bartelsman, E. J. and M. Doms (2000): "Understanding Productivity: Lessons from Longitudinal Microdata," *Journal of Economic Literature*, 38, 569–594.

Bertheau, A., E. M. Acabbi, C. Barcelo, A. Gulyas, S. Lombardi, and R. Saggio (2023): "The Unequal Consequences of Job Loss across Countries," *American Economic Review: Insights*, 5, 393–408.

Black, S. E., P. J. Devereux, and K. G. Salvanes (2015): "Losing Heart? The Effect of Job Displacement on Health," *ILR Review*, 68.

Blien, U., W. Dauth, and D. H. Roth (2021): "Occupational routine intensity and the costs of job loss: evidence from mass layoffs," *Labour Economics*, 68.

Bonhomme, S., T. Lamadon, and E. Manresa (2019): "A Distributional Framework for Matched Employer-Employee Data," *Econometrica*, 81(3), 699–739.

Bratsberg, B., O. Raaum, and K. Røed (2018): "Job loss and immigrant labour market performance," *Economica*, 85, 124–151.

Braxton, J. C. and B. Taska (2023): "Technological Change and the Consequences of Job Loss," *American Economic Review*, 113, 279–316.

Breiman, L. (2001): "Random forests," *Machine learning*, 45(1), 5–32.

Britto, D. G. C., P. Pinotti, and B. Sampaio (2022): "The Effect of Job Loss and Unemployment Insurance on Crime in Brazil," *Econometrica*, 90, 1393–1423.

Burgess, S., J. Lane, and D. Stevens (2000): "Job Flows, Worker Flows, and Churning," *Journal of Labor Economics*, 18(3), 473–502.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36, S13–S70.

Card, D., 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," *Quarterly Journal of Economics*, 131, 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.

Carlsson, M., I. Häkkinen Skans, and O. N Skans (2019): "Wage Flexibility in a Unionized Economy with Stable Wage Dispersion," Discussion Paper 12093, IZA.

Carlsson, M., J. Messina, and O. N. Skans (2016): "Wage adjustment and productivity shocks," *The Economic Journal*, 126, 1739–1773.

Cederlöf, J. (2019): "Saved by Seniority? Effects of Displacement for Workers at the margin of Layoff," *Job market paper*.

Chernozhukov, V., M. Demirer, E. Duflo, and I. Fernandez-Val (2020): "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India," Working paper No. w24678, National Bureau of Economic Research.

Davis, S. J. and T. von Wachter (2011): "Recessions and the costs of job loss," *Brookings Papers on Economic Activity*.

Eliason, M. (2012): "Lost jobs, broken marriages," *Journal of Population Economics*, 25, 1365–1397.

——— (2014a): "Alcohol-related morbidity and mortality following involuntary job loss: Evidence from Swedish register data," *Journal of Studies on Alcohol and Drugs*, 75, 35–46.

——— (2014b): "Assistant and auxiliary nurses in crisis times: Earnings, employment, and income effects of female job loss in the Swedish public sector," *International Journal of Manpower*, 35, 1159–1184.

Eliason, M. and D. Storrie (2006): "Lasting or latent scars? Swedish evidence on the long-term effects of job displacement," *Journal of Labor Economics*, 24, 831–856.

Farber, H. S. (2011): "Job Loss in the Great Recession: Historical Perspective from the Displaced Workers Survey, 1984-2010," Working Paper 17040, National Bureau of Economic Research.

Gathmann, C., I. Helm, and U. Schönberg (2020): "Spillover effects of mass layoffs," *Journal of the European Economic Association*, 18, 427–468.

Goos, M., A. Manning, and A. Salomons (2014): "Explaining job polarization: Routine-biased technological change and offshoring," *American Economic Review*, 104, 2509–26.

Gulyas, A. and K. Pytka (2021): "Understanding the Sources of Earnings Losses After Job Displacement: A Machine-Learning Approach," Conference version, at the jahrestagung des vereins für socialpolitik 2021: Climate economics, zbw - leibniz information centre for economics, kiel, hamburg.

Halla, M., J. Schmieder, and A. Weber (2020): "Job Displacement, Family Dynamics, and Spousal Labor Supply," *American Economic Journal: Applied Economics*, 12, 253–87.

Helm, I., I. Kügler, and U. Schönberg (2023): "Displacement Effects in Manufacturing," mimeo.

Hu, L. and C. Taber (2011): "Displacement, asymmetric information, and heterogeneous human capital," *Journal of Labor Economics*, 29, 113–152.

Huckfeldt, C. (2022): "Understanding the Scarring Effect of Recessions," *American Economic Review*, 112, 1273–1310.

Ichino, A., G. Schwerdt, R. Winter-Ebmer, and J. Zweimüller (2007): "Too old to work, too young to retire?" *Too Young to Retire*.

Jacobson, L. S., R. LaLonde, and D. Sullivan (1993): "Earnings Losses of Displaced Workers," *American Economic Review*, 83, 685–709.

Kuhn, P. J., ed. (2002): *Losing Work, Moving On: International Perspectives on Worker Displacement*, W.E. Upjohn Institute for Employment Research.

Lachowska, M., A. Mas, and S. Woodbury (2020): "Sources of Displaced Workers' Long-Term Earnings Losses," *American Economic Review*, 110.

Lamadon, T., M. Mogstad, and B. Setzler (2022): "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market," *American Economic Review*, 110, 169–212.

Leighton, M. and J. D. Speer (2020): "Labor market returns to college major specificity," *European Economic Review*, 128.

Mueller, A. I. and J. Spinnewijn (2023): "The Nature of Long-Term Unemployment: Predictability, Heterogeneity and Selection," Working Paper 30979, National Bureau of Economic Research.

Nie, X. and S. Wager (2021): "Quasi-oracle estimation of heterogeneous treatment effects," *Biometrika*, 108(2), 299–319.

OECD (2019): "Society at a glance: OECD Social Indicators," .

Salvanes, K. G., B. Willage, and A. L. Willén (2023): "The Effect of Labor Market Shocks Across the Life Cycle," *Journal of Labor Economics*.

Schmieder, J. F., T. M. von Wachter, and J. Heining (2023): "The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany," *American Economic Review*.

Schwerdt, G., A. Ichino, O. Ruf, R. Winter-Ebmer, and J. Zweimüller (2010): "Does the color of the collar matter? Employment and earnings after plant closure," *Economics Letters*, 108, 137–140.

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

Skans, O., L. Hensvik, and S. Eriksson (2017): *Åtgärder för en inkluderande arbetsmarknad*, SNS förlag.

Yadlowsky, S., S. Fleming, N. Shah, E. Brunskill, and S. Wager (2021): "Evaluating Treatment Prioritization Rules via Rank-Weighted Average Treatment Effects," .

Yakymovych, Y. (2022): "Consequences of job loss for routine workers," Working paper 2022:15, IFAU.

A Online Appendix: Data Details

Below, we explain the variables used in the GRF model and when matching displaced workers to non-displaced controls.

A.1 Demographics

Data on basic demographics are drawn from the administrative population register Louise produced by Statistics Sweden. Data in Louise is collected in November of each year. *Age* is measured in years. We include a *Female* dummy. A categorical *Native-immigrant* variable takes the value zero for natives, one for those with two foreign-born parents, and two for those born abroad (unless both parents are born in Sweden).

A.2 Family

Civil status is coded using two dummies. *Married* workers are either formally married, or cohabit with a partner with whom they have a common child. This definition is well in line with Swedish perceptions of ”marriages” and is used in most research on Swedish data. *Divorced* is a dummy for workers whose marriage has ended and who are not currently defined as married according to *Married*.

To quantify the economic value of the partnership and degree of dependence, our *Share of household income* variable measures the subject's labor earnings as a share of total household labor earnings.

We measure the number of children in the household using two different variables. *Children* contains the number of all children under the age of 18 in the household, whereas *School-aged children* counts children aged 7 to 17. We make this distinction to allow the second variable to measure the additional impediments to mobility that may arise when children start school at age 7.

Social ties to the current location are measured by two variables. First, we define a dummy for being *Born outside of the current county*. This variable naturally takes the value 1 for all foreign-born.[37] We also measure *Past mobility* as the number of moves across local labor market boundaries during the past 10 years. This variable is a lagged counterpart of our key indicator for geographical mobility after the event.

[37] The register information is based on the county of birth and cannot be disaggregated further.

A.3 General Human Capital

Years of schooling are based on the highest achieved level of education. *Labor market experience* is measured as the number of years employed during the 10 years prior to displacement. Because of the tenure restriction, all displaced and control workers must have been employed for at least three years prior to the year $t = 0$ and the variable therefore has a range from 3 to 10.

In this category, we also include pre-displacement earnings. We measure these by *Earnings* as the rank among the full population of displaced and controls in year $t = -1$, as well as two variables capturing *Earnings growth* from $t = -3$ and $t = -2$ respectively to $t = -1$. The rank form of the *Earnings* variable has been chosen to ensure orthogonality to time trends.

A.4 Specific Human Capital

It is likely that workers differ in terms of how costly or difficult it is for them to switch firm and industry if hit by a negative shock. The ability to adjust along this margin is potentially determined by the degree to which the worker is tied to the closing firm or sector. In order to measure the empirical relevance of these aspects, we include variables capturing pre-closure establishment *Tenure* (truncated at 10), and the number of years spent in the same industry as the closing establishment (*Industry tenure*, also truncated at 10).

Furthermore, we characterize the *Education specificity* using data on the 1-digit level and 3-digit field of the highest achieved education (thus, high school programs and college majors). Our metric uses the fraction of workers, by education cell, that is employed in the ten main 3-digit industries. The strategy follows Altonji et al. (2012).[38] Examples of fields with high levels of specificity are pharmacists and nurses.

We further include an indicator for *STEM* education beyond high school. We also provide a dummy for *Education in a licensed field*, that is a field that caters to the health or education sector.

A.5 Lost Job Characteristics

The causal impact of job loss for displaced workers does not only depend on their outside opportunities, but also on aspects of the job that they have lost. We therefore define a

[38] An alternative is provided by Leighton and Speer (2020) but their approach requires a full matrix where all types of education are present in all industries and is therefore less well suited to our granular data.

number of variables which capture key characteristics of the lost job.

We first characterize the lost job by *Plant size* in terms of employment in $t = -1$, and *Plant size trend* between $t = -3$ and $t = -1$, measured as the difference in log(*Plant size*).

Additionally, we measure the *Plant wage* premium at the closing establishment. This feature is shown to be particularly important in the case of Austria by Gulyas and Pytka (2021). They study mass layoffs and characterize the affected firms by firm wage effects as estimated conditional on person-fixed effects in the tradition of the AKM model (Abowd et al., 1999). Indeed, there is a large literature discussing the origins of firm-specific rents, and how they should be estimated and interpreted (see, e.g., Card et al., 2018; Bonhomme et al., 2019; Lachowska et al., 2020, in the context of plant closures). A drawback of the AKM model is that it requires structural assumptions that we do not want to impose, and is likely to result in biased premium estimates for dying firms. We therefore take a slightly different route. As a reduced form measure of the pay level of the closing firm at the time of displacement, we use the leave-one-out mean of residual earnings in year $t = -1$ as in Card et al. (2013). For each displaced worker and potential control, we characterize the closing establishment environment by the average Mincer residual earnings of all co-workers at the same establishment. Residual earnings are computed using year-specific regressions of log earnings on 3-digit industry indicators, years of schooling, field of education (2 digits), gender, immigration status and a full set of age dummies.[39]

We measure the *Routine* task component of the lost job. To this end, we use data on routine intensity by occupation from Autor and Dorn (2013) and Goos et al. (2014). We translate the occupational codes into the Swedish nomenclature. We have data on occupations for about half of our workers. We impute routine intensity for each worker based on the average routine score for those in the same detailed education-industry combination. [40]

We further construct a *Manager* dummy for workers who are employed as managers. This information is drawn from the occupational codes, which we do not have for all workers, and therefore contains false negatives. To mitigate the problem, we impute manager status based on data from the previous three years if the information is missing in $t = -1$.

[39]The R^2 in a regression of labor earnings on the plant wage premium measure among the displaced is around 22 percent.

[40]We define the combinations using 1-digit education level, 2-digit education field, and three-digit industry. For cells where we have fewer than 100 workers we use data on the 1-digit industry and education (field+level) instead. The leave-out correlation between our education-industry-based measure of routineness and routineness as measured directly based on occupation is 0.7. The reason for not using occupation-level routineness directly is the large number of missing values this would imply.

We calculate the *Relative size of the displacement event* as the share of total employment by 3-digit industry and local labor market combination. This is motivated by previous studies, e.g. Cederlöf (2019) and Gathmann et al. (2020), that have found that the impact of being displaced in a large event is particularly severe. One possible explanation for this finding is that workers from the same event may compete with each other for the same job openings, and this type of competition might be particularly problematic if the displacement event is large relative to the industry-specific local labor market.

Finally, we generate a variable for *Industry being matched to education*. This indicator takes on the value 1 for workers who were employed in one of the 10 main 3-digit industries for their 1-digit level and 3-digit field before displacement.

A.6 Industry Characteristics

There is ample evidence that, on average, workers have comparative advantage in their industry of employment and that shocks to this industry have an impact on their overall earnings prospects. Examples include Carlsson et al. (2016) for Sweden, who show that that technology shocks have a larger impact on workers' wages if the shocks are shared with other firms in the same industry and that this distinction is entirely driven by workers with fields of education where most job-to-job mobility is within the industry. Similar arguments are made in Lamadon et al. (2022). It is also likely that displacement has more lasting negative effects in industries with low labor turnover such as manufacturing than in fluid sectors such as restaurants. Overall, this suggests that workers who lose their jobs in declining low-turnover sectors will suffer more adverse consequences, in particular if much of their human capital is industry-specific.

In order to quantify the conditions in each industry, we first construct time-consistent industry indicators at the 3-digit level. This is a non-trivial endeavour as the codes changed three times during our sample period. We start from the SNI2002 codes that our raw data provide for the period 2002 to 2010. Next, we rely on the SNI2002 code reported in 2010 and use it for all later years for those establishments that continue to exist beyond 2010, and conversely use the data from 2002 for establishments that existed prior to that. We refer to these overlapping establishments as stayers. We then use the modal overlap between SNI2002 (as imputed for stayers) and the current codes (SNI69, SNI92 and SNI2007) to fill in SNI2002 codes for non-stayers. This works particularly well for the years and sample we study, as almost all closing establishments existed in 2010 (due to the 3-year tenure requirement) and since most codes remained unchanged between SNI92 and SNI2002.

We measure *Industry wage* premia in each 3-digit industry as average earnings in the industry, residual on years of schooling, 2-digit education field, gender, immigrant background and a full set of age dummies. Furthermore, we also include average *Churn* and *Reallocation* rates for each industry. We follow conventions from Burgess et al. (2000), and define establishment-level churning as the number of workers who are hired or separated beyond what was needed for the actual change in employment between two adjacent years. The churning rate is measured as churning relative to the average employment during the two years.[41] Similarly, we calculate excess reallocation in each industry as the excess creation and destruction of jobs over what was needed to adjust industry employment. The reallocation rate is measured relative to the average employment in the two years.[42] We then aggregate the establishment-year churning rates and the industry-year reallocation rates to industry-level averages which are constant over time. When computing these numbers, we take care to exclude workers that satisfy the conditions for false closures as discussed above. Note also that the scale of the measures (relative to average employment in the two years) is in approximate percentages, but with a maximum of 2 and a minimum of -2.

We also follow the conventions from this literature when computing *Employment growth over the past 10 years* in each industry. Thus, for displacement events in the year $t = -1$, we define $Trend_{Ind,t} = (Emp_{Ind,t-1} - Emp_{Ind,t-10})/(Emp_{Ind,t-1}/2 + Emp_{Ind,t-10}/2)$ which takes the value 2 for newly emerging industries and the value -2 for disappearing industries. This metric bounds cases where some very small industries experience extreme changes (including exits and entries) during the sample period. Similarly, we calculate the *Industry-specific business cycle* as the change in employment between year $t = -1$ and $t = 0$ using the same metric.

We also add dummy variables for closures in the *Manufacturing industry*, as this industry has been a focus of the job displacement literature since Jacobson et al. (1993). Furthermore, we add a dummy for closures within *Education, health and public administration*, as the labor markets in these (mostly public sector) industries tend to experience a constant shortage of workers.

[41] Thus, churning is [(Hires+Separations)-abs(Emp(t)-Emp(t-1))]/[Emp(t)/2+Emp(t-1)/2] at the establishment-year level.

[42] Thus, reallocation is [(Job Creation+Job Destruction)-abs(Emp(t)-Emp(t-1))]/[Emp(t)/2+Emp(t-1)/2] at the industry-year level.

A.7 Location Characteristics

As with industries, there is ample evidence suggesting that local labor market conditions have causal effects on workers' outcomes. This is also the case in Sweden, where wages seem to react to local labor market shocks (Carlsson et al., 2019).

Our local variables are measured at the level of local labor markets, which are an aggregation of municipalities. These are constructed by Statistics Sweden based on commuting patterns.

For each local labor market, we measure the *Local unemployment* rate as the number of residents who are registered with the public employment service (a prerequisite for receiving benefits from either the unemployment insurance system or from the municipal welfare system) divided by the size of the local labor force (sum of the registered unemployed and the number of employed as described above). *Population density* is also measured at the local labor market level. The *Concentration of local employment across 3-digit industries* is measured using an HHI index.

In addition, we measure local exposure to the industry characteristics discussed above. These shift-share/Bartík-style variables are calculated by multiplying each industry's employment share in the local labor market (by year) with the characteristics of that industry and then summing over the industries within the local labor market. In this way, we define the *Manufacturing employment share*, *Average long-term industry trend*, *Average industry business cycle*, *Average industry churn* and *Average industry reallocation*.

It is important to note the fundamental difference between these variables and their industry-level counterparts. Whereas the industry-level variables measure characteristics in the industry from which the worker was displaced (and is potentially tied to, if changing industry is costly), the local labor market counterpart measure how exposed the worker is to these attributes if searching at random at the local labor market (which should be more pertinent for workers who are restricted in their mobility).

A.8 Aggregate Conditions

To capture aggregate conditions in the economy, which may both impact the average size of earnings losses, as well as have heterogeneous effects on different worker groups, we include the calendar *Year* ($t = 0$) and the *National unemployment rate* in the year $t = 1$. The national unemployment rate captures the aggregate business cycle conditions displaced workers experience after job loss. These have been shown to be an important correlate of earnings losses by Davis and von Wachter (2011).

B Online Appendix: Additional Results

Table B.1: Sample statistics for displaced and non-displaced workers

	Controls (unmatched) (1)	Controls (matched) (2)	Displaced (unmatched) (3)
N workers	28,882,769	554,454	184,833
N establishments	230,617	119,830	21,785
Covariates			
Demographics			
Age	44.111	42.161	42.153
Female	0.478	0.347	0.347
Born abroad	0.097	0.129	0.130
Two foreign born parents	0.028	0.032	0.032
Native, not birth county	0.274	0.251	0.250
Family			
Married	0.644	0.588	0.587
Divorced	0.101	0.099	0.099
Children (# living at home)	0.780	0.746	0.745
School-aged children	0.524	0.466	0.466
Share household inc.	0.714	0.755	0.756
Past mobility (# location moves in last 10 years)	0.164	0.212	0.212
Human capital			
Schooling (years of education)	12.162	11.466	11.473
Experience (years employed in last 10 years)	9.229	8.918	8.910
Tenure	6.624	5.906	5.891
Industry tenure	7.922	7.134	7.127
Earnings (rank)	0.500	0.497	0.497
Earnings in t-2	0.994	1.014	1.014
Earnings in t-3	0.951	0.962	0.962
STEM education	0.097	0.109	0.109
Education in licensed field	0.238	0.070	0.071
Education specificity	0.565	0.470	0.470
Lost job characteristics			
Establishment size	521.051	179.935	186.073
Establishment size trend	0.025	-0.022	-0.024
Establishment wage premium	0.001	-0.032	-0.033

Job routineness	0.515	0.558	0.558
Industry matched to education	0.560	0.376	0.375
Size of displacement	0.170	0.157	0.159
Manager	0.033	0.027	0.027
Industry characteristics			
Employment trend	0.052	0.079	0.080
Employment cycle	0.005	-0.001	-0.001
Industry wage premium	0.006	0.028	0.027
Churning	0.223	0.209	0.209
Reallocation rate	0.131	0.139	0.139
Manufacturing	0.230	0.365	0.367
Education, health, admin	0.383	0.078	0.079
Location characteristics			
Population density	82.467	86.204	86.241
Unemployment	0.086	0.087	0.086
Concentration (HHI for industries)	0.031	0.031	0.031
Manufacturing (share of local emp.).	0.189	0.186	0.186
Average industry trends	0.081	0.079	0.079
Average industry cycle	0.007	0.007	0.007
Average industry churning	0.224	0.226	0.226
Average industry reallocation	0.147	0.148	0.149
Year	2005.881	2005.576	2005.576
National unemployment	7.451	7.432	7.432

Note: Statistics for the non-displaced (controls) before matching (Column 1) and after matching (column 2), and for the displaced workers (Column 3). Mean values of all covariates used in the analyses. Details for all covariates are in Appendix A.

Table B.2: RATE estimates of GRF model and different simple targeting rules on the test set

	Qini coefficient (x100)	SE (x100)
GRF	-3.95	0.3
Age	-2.06	0.2
Schooling	-1.66	0.2
Manufacturing	-1.51	0.2
Population density	-1.35	0.2
Manufacturing & age	-2.58	0.2
Manufacturing & schooling	-2.23	0.3

Note: *GRF*: workers with the lowest estimated CATEs; *Age*: oldest workers; *Schooling*: *Manufacturing*: manufacturing workers at random and then other workers; *Population density*: least dense locations; *Manufacturing and age*: oldest manufacturing workers; *Manufacturing and schooling*: least educated manufacturing workers.

Figure B.1: Distribution of relative earnings five and ten years after displacement

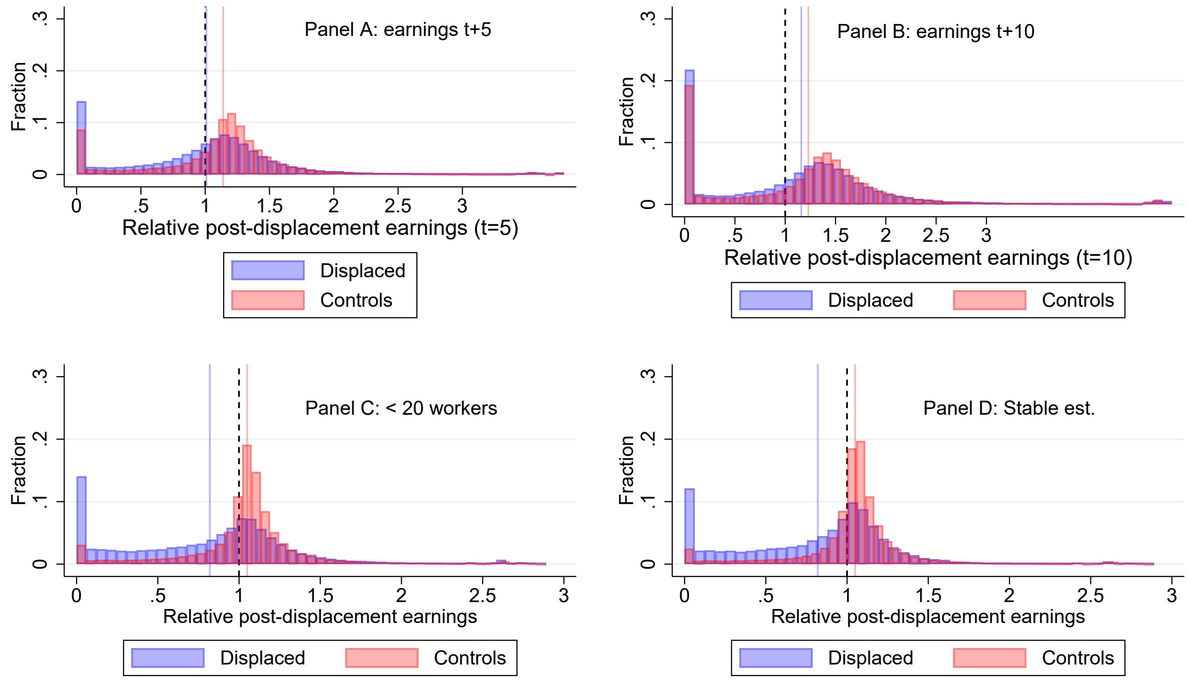

Note: Panels A and B show the distributions of relative earnings between one year before displacement and five years after displacement (Panel A) and ten years after displacement (Panel B). Panels C and D show distributions of relative earnings in $t = 1$ for displaced and matched controls in sub-samples with less than 20 workers (Panel C) and with less than 10 percent change in employment between t-3 and t-1 (Panel D), similar to Figure 1B in the paper. Solid lines indicate group means, the dashed lines indicate unchanged nominal earnings.

Figure B.2: Robustness related to Figure 4: ATE estimates across the CATE distribution using alternative control groups

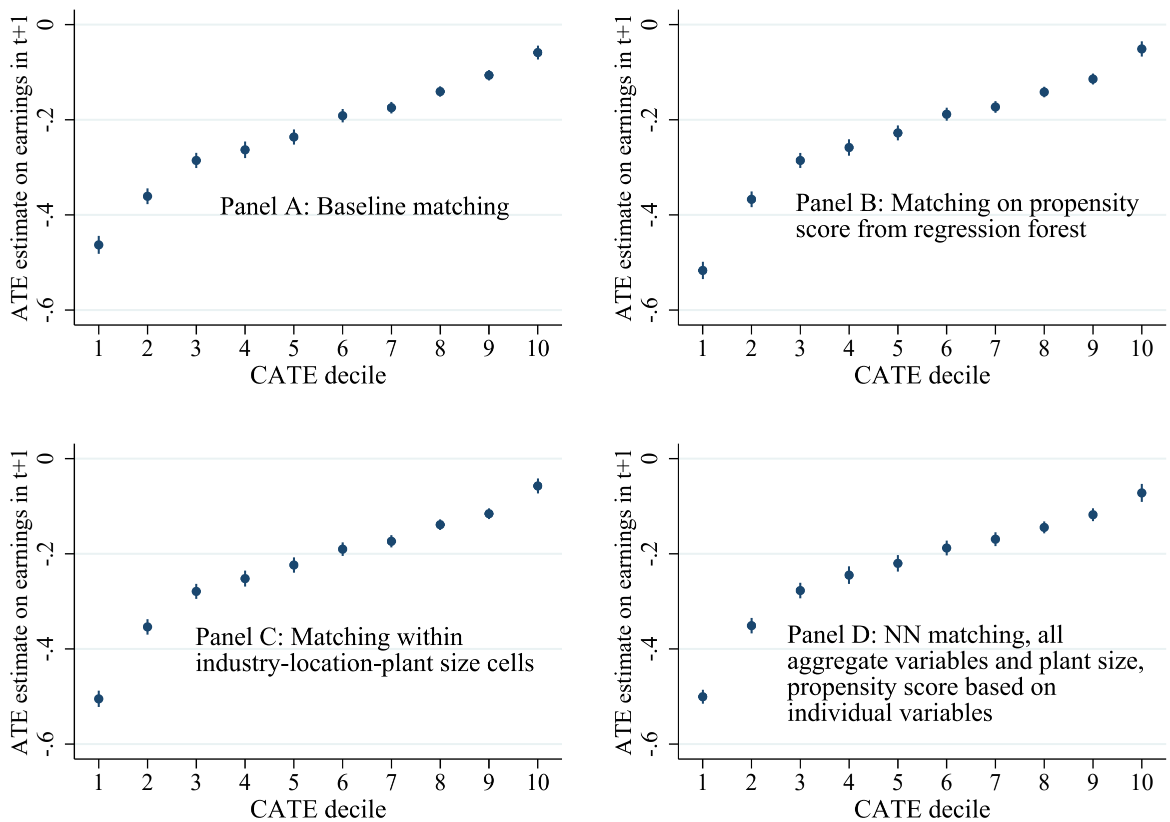

Note: Panel A reports the ATE estimates for earnings in $t + 1$ corresponding to Figure 4B. Panels B–D report robustness exercises when matching the displaced workers to control workers using alternative matching approaches. Treated workers are allocated to deciles according to the baseline CATE:s as in Figure 4. Control workers are placed in the same CATE decile as the treated workers they are matched to. Thus, Panel A is conceptually different from Figure 4B which allocates control workers to deciles based on their own characteristics. Panel B matches on a propensity score estimated using a regression forest instead of the logit estimation used in the baseline analysis. Panel C performs propensity score matching combined with exact matching on industry groups (manufacturing, education/health/public administration and other), a big city indicator and plant size (3 groups). Panel D performs Mahalanobis distance nearest-neighbor matching using the industry and location characteristics as well plant size as separate variables and a propensity score estimated using the other variables. Three control workers matched to each displaced worker in all panels.

Figure B.3: ATE on earnings in t+3 by decile of CATE:s on earnings in t+1

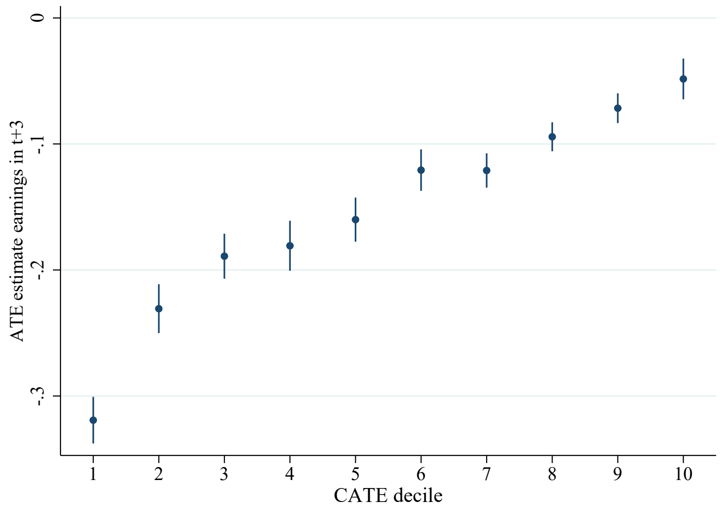

Note: ATE:s (differences between displaced and control workers) in terms of earnings three years after displacement for workers in each decile of the CATE distribution for earnings one year after displacement. 95 percent confidence intervals.

Figure B.4: Effects of displacement on employment over time

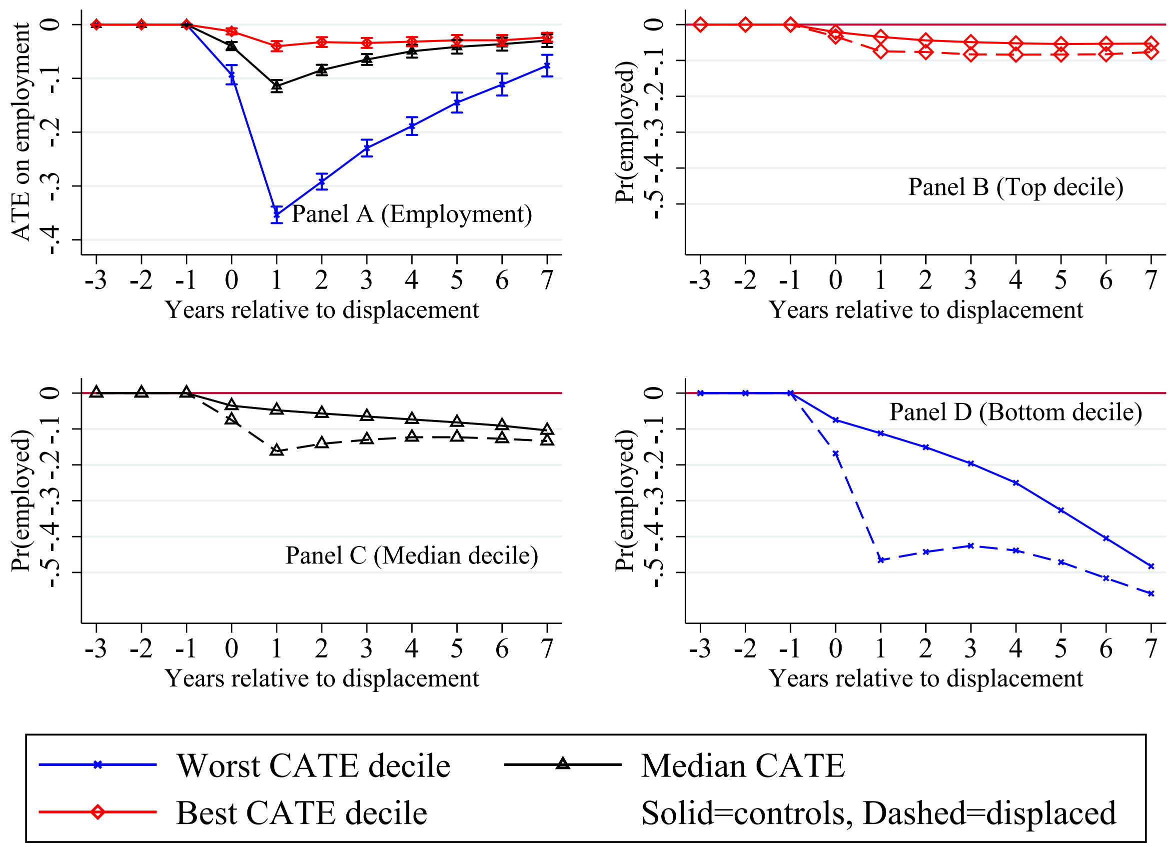

Note: The figure replicates figure 5 in the paper, but with employment as the outcome instead of earnings. It shows statistics for three decile groups of the CATE distribution. The "median" group straddles the median (i.e. it contains the $10th$ and $11th$ ventile). Panel A shows the ATE over time, similar to figure 1, but separately for the decile groups. Point estimates and 95 percent confidence interval with standard errors clustered at the establishment in $t = -1$. Panels B–D show the underlying employment trajectories for displaced and matched controls within each decile group for the top, bottom and median deciles respectively. Only workers observed in each of the periods $t = -3$ to $t = 7$ are included (this entails excluding the years 2011-2014). Sample restrictions ensure that employment is equal to 1 for all until $t = -1$.

Figure B.5: Differences in characteristics across CATE quartiles and deciles

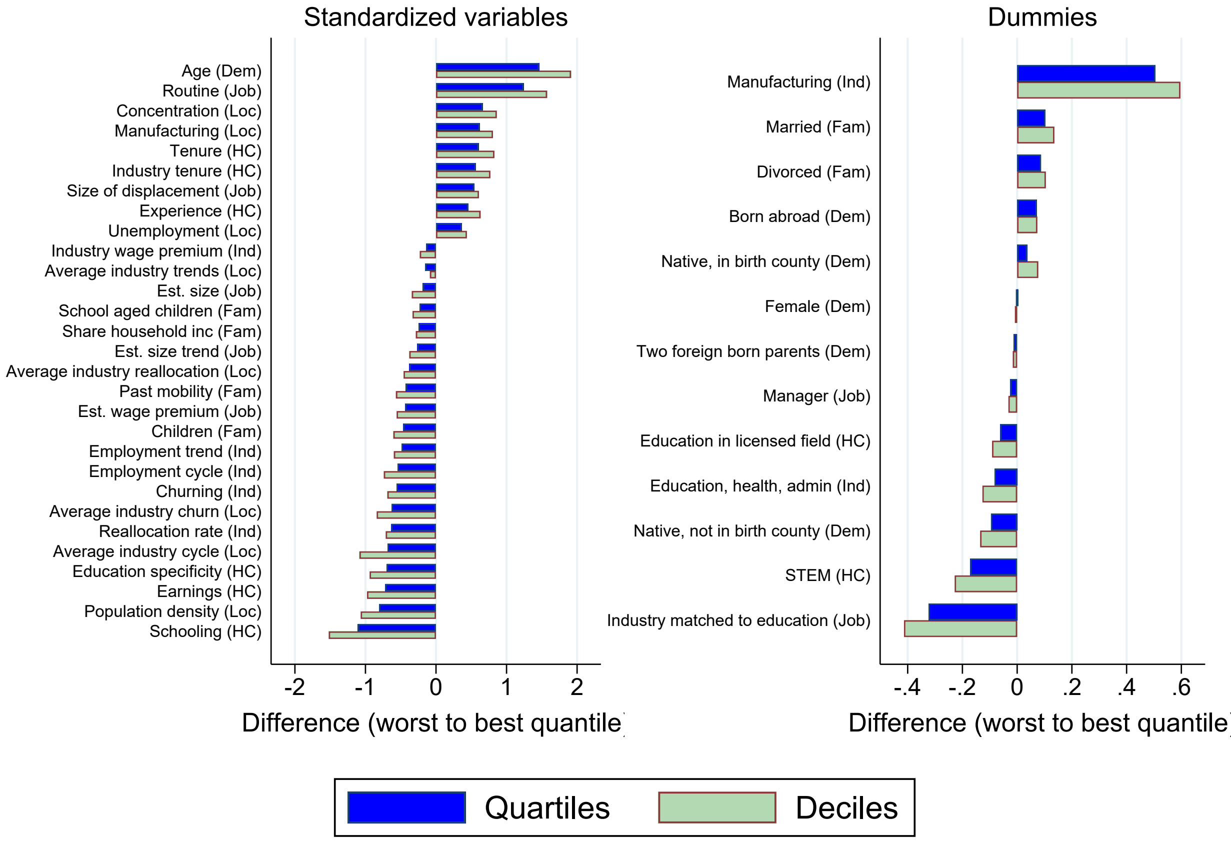

Note: The figure shows differences in characteristics between individuals in the lowest and highest quartiles and deciles of CATE:s, using the training data set. CATE:s estimated using 5-fold estimation and ranking done within each fold. The left-hand panel contains standardized (mean 0 and standard deviation 1) continuous variables and the right-hand panel contains dummy variables. Blue bars indicate quartiles and green bars indicate deciles.

Figure B.6: CATE across combinations of age and schooling

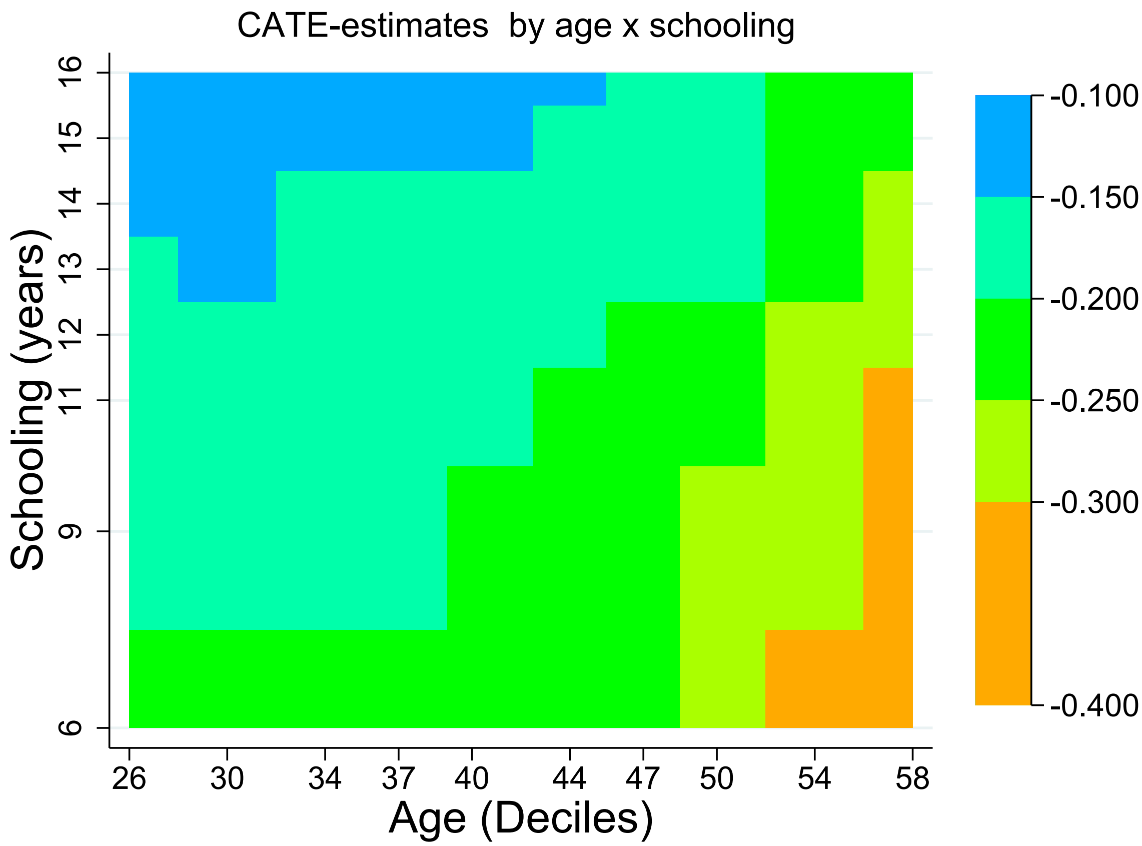

Note: The figure replicates Figure 7A (and B) in the paper, but using CATE instead of ATE (or AIPW). It divides training set workers into cells by age and schooling. Schooling has been aggregated to 8 groups by pooling the few with 10 years of schooling together with those with 9 years of schooling, and by letting the top group include all with 16 or more years of schooling. Age is defined in deciles among the displaced and the x-axis shows the median in each age group. Colors indicate the size of point estimates. It shows estimated CATE:s by combinations of age and schooling.

Figure B.7: Heterogeneity within extreme combinations of age and schooling

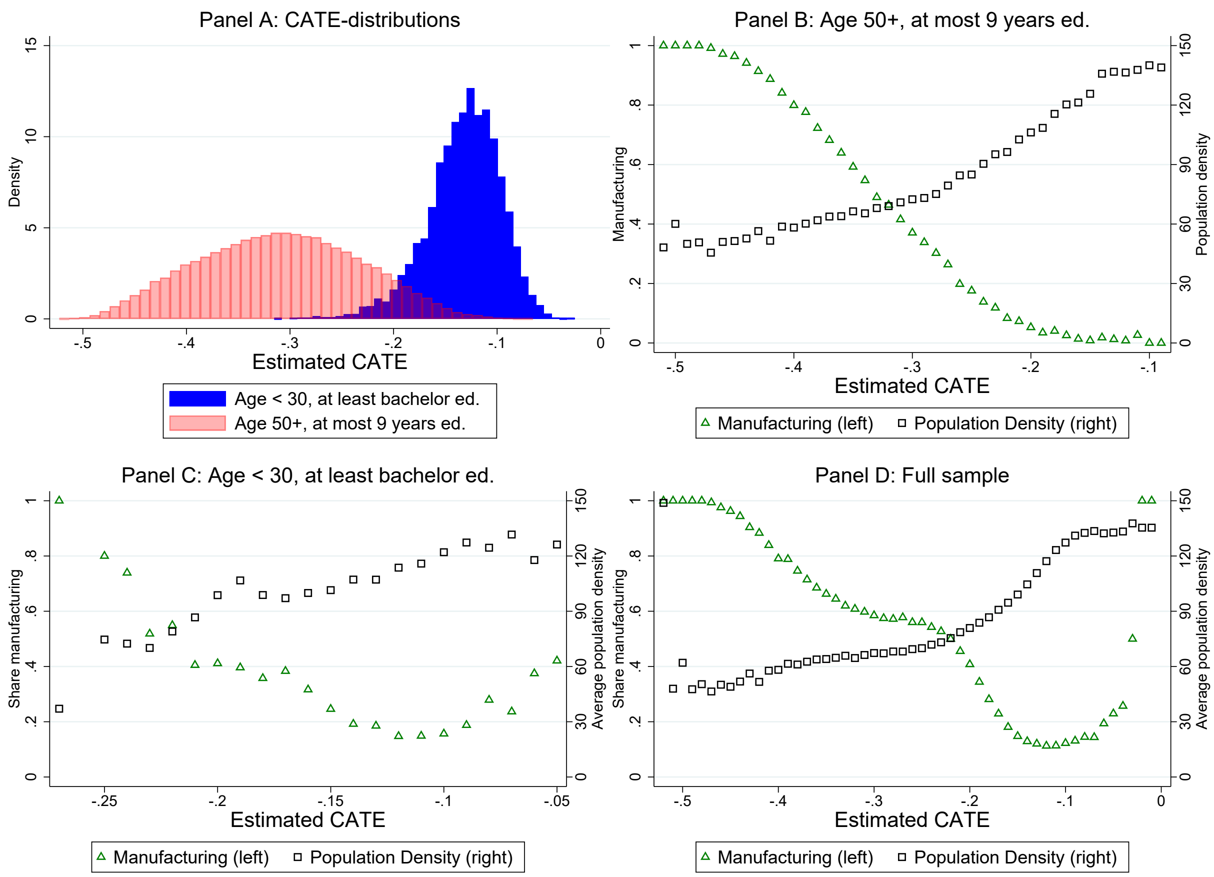

Note: Training set workers who are (a) younger than 30 and have at least 15 years of education or (b) older than 50 and have at most 10 years of education. Panel A: Histograms of 5-folds GRF CATEs within these groups. Panels B-C: Share of manufacturing workers and average population density within each CATE cell for the two groups. Panel D repeats this for the full sample. Cells defined as CATE bins of one percentage point. Bins containing fewer than ten workers are dropped.

Figure B.8: Robustness analyses: Displacement effects across and within establishments using AIPW

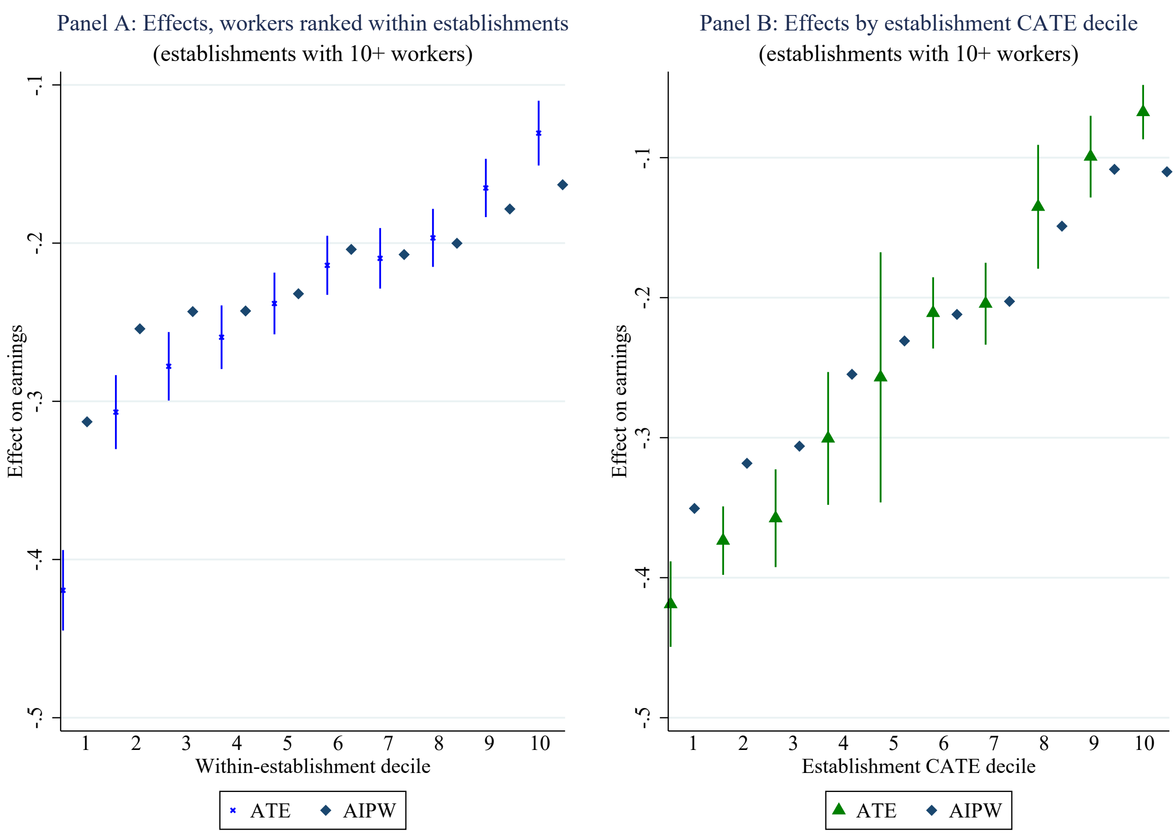

Note: The figure shows robustness analyses using both ATE:s and average AIPW scores in the respective CATE deciles. Panel A shows ATE/AIPW estimates when displaced workers are ranked based on their within-establishment CATE. Panel B shows ATE/AIPW estimates when displaced workers are ranked based on the CATE of their co-workers (defined as the leave-out mean for the workers at the establishment and then averaging over the individuals in the CATE decile.) This analysis excludes establishments with fewer than 10 displaced workers. Controls are allocated to the same decile sample as the treated workers they were matched to.

Figure B.9: Robustness analyses: Displacement effects across and within establishments for additional outcomes

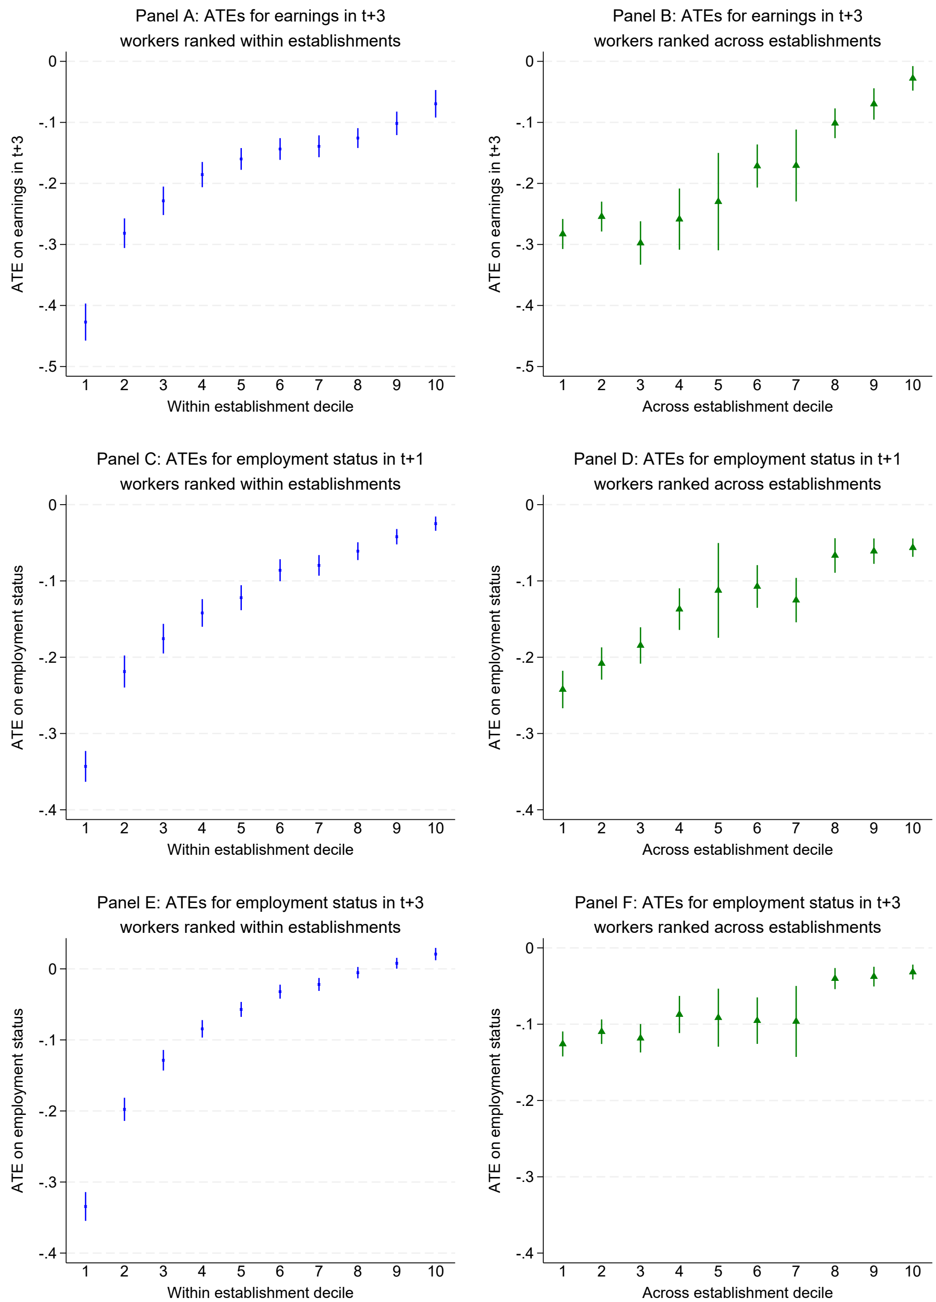

Note: The figure shows robustness analyses using employment status in $t + 1$, earnings in $t + 3$ and employment status in $t + 3$ as outcomes. Panels A, C and E show ATE estimates when displaced workers are ranked based on their within-establishment CATE. Panels B, D and F show ATE estimates when displaced workers are ranked based on the CATE of their co-workers (defined as the leave-out mean for the workers at the establishment and then averaging over the individuals in the CATE decile.) Panels A and B exclude establishments with fewer than 10 displaced workers. Controls are allocated to the same decile sample as the treated workers they were matched to.

Figure B.10: Robustness analyses: Displacement effects across and within establishments by AKM plant estimates

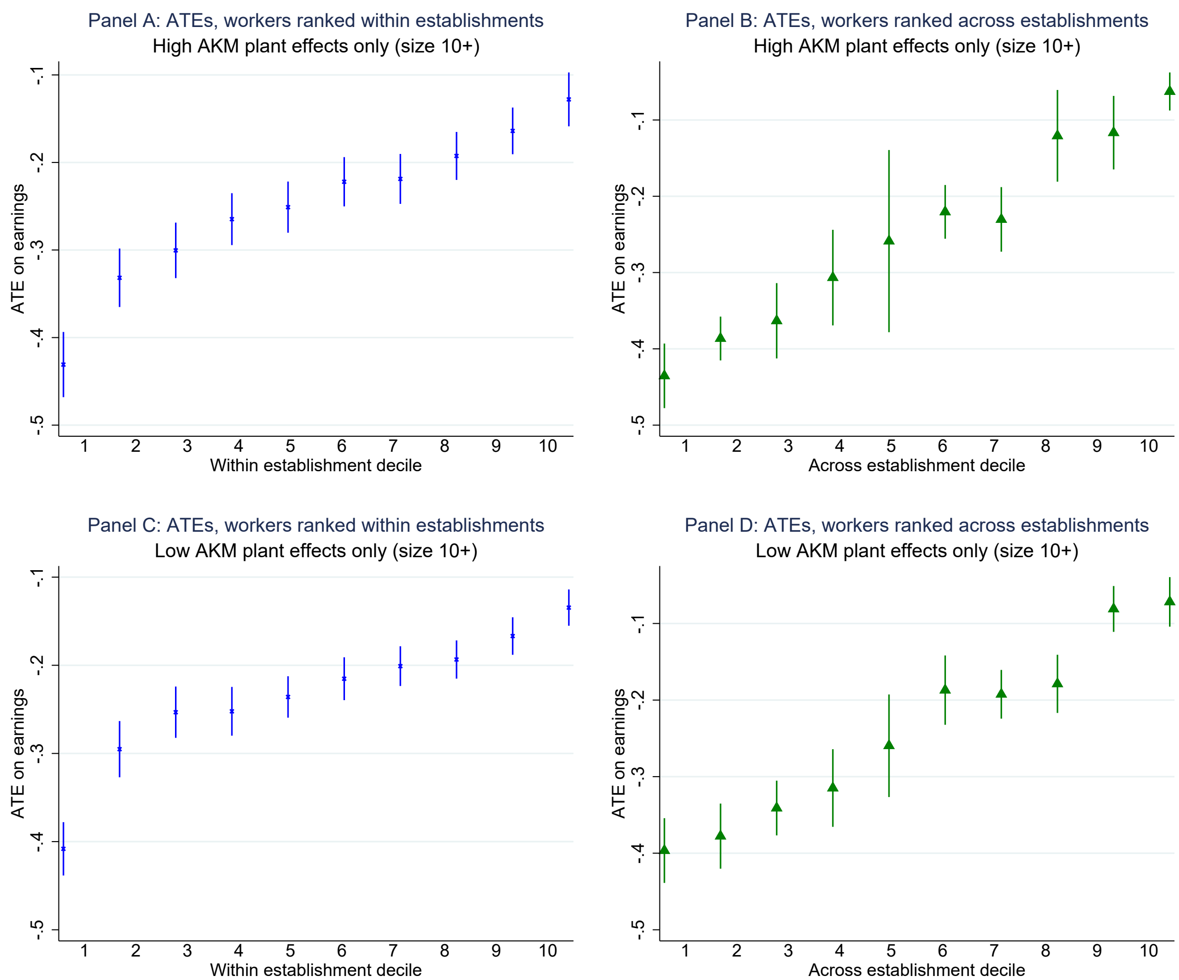

Note: Panel A (high AKM-plant effects) and C (low AKM-plant effects) shows ATE estimates when displaced workers are ranked based on their within-establishment CATE. Panel B (high AKM-plant effects) and D (low AKM-plant effects) shows ATE estimates when displaced workers are ranked based on the CATE of their co-workers (defined as the leave-out mean for the workers at the establishment and then averaging over the individuals in the CATE decile.) All panels exclude establishments with fewer than 10 displaced workers. Controls are allocated to the same decile sample as the treated workers they were matched to. AKM-estimates (Abowd et al., 1999) are derived from estimates on rolling 5-year samples ending in the year prior to the event ($t = -1$). Models have fixed establishment (”plants”) effects, fixed person effects, year dummies, and third order polynomial in age (in deviation from 45), separately for three education groups (less than high school, high school, more than high school). The outcome is log earnings conditional on employment. Sample is split into high vs. low plant-effects samples by the median plant effect within each displacement cohort.

Figure B.11: CATE across and within establishments

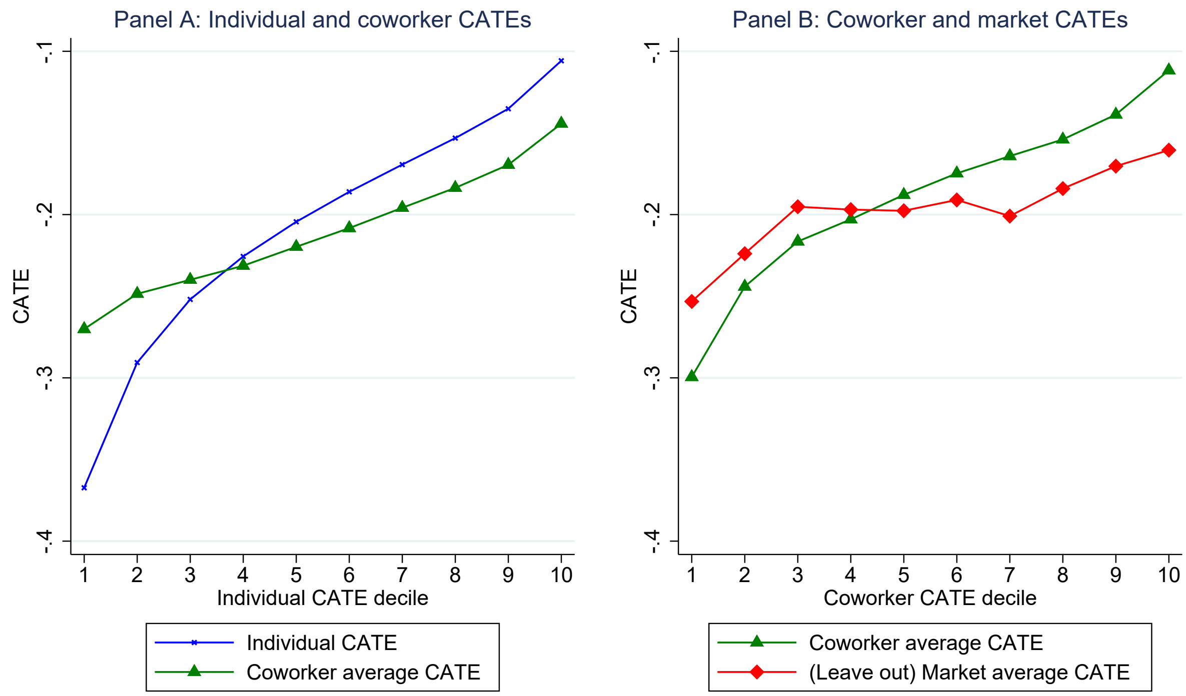

Note: Panel A shows statistics for displaced workers in the main (training) set by CATE decile (5-fold estimation and ranking done within each fold). It reports average individual CATE and the average CATE for the coworkers of the individuals in each CATE decile taken as a leave-out mean across the workers at the establishment and then averaging over the individuals in the CATE decile. Panel C shows average coworker CATE and the average CATE for other closures in the same industry, location and year by coworker CATE decile. Establishments with less than 10 displaced workers are discarded in Panel A. Workers from markets with only one event are discarded in Panel B.

Figure B.12: Heterogeneity within and across establishments, interactions

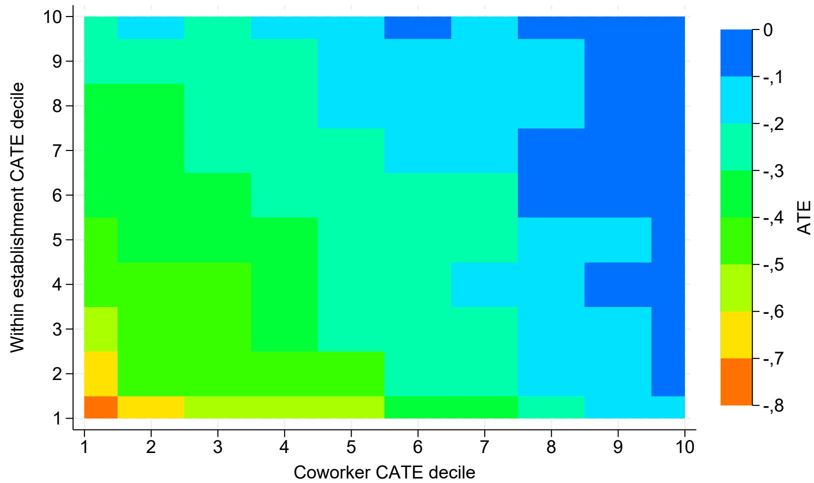

Note: The figure divides training set workers by combinations of coworker 5-fold CATE decile and within-establishment decile. Coworker CATE decile is determined by the ranking of the leave-out mean CATE at the establishment. For the within-establishment CATE, individuals are ranked by their CATE:s within the establishment. The figure restricts to establishments with at least 10 workers. Colors indicate the size of point estimates of the CATE within each cell.

Figure B.13: Insurance across the CATE distribution

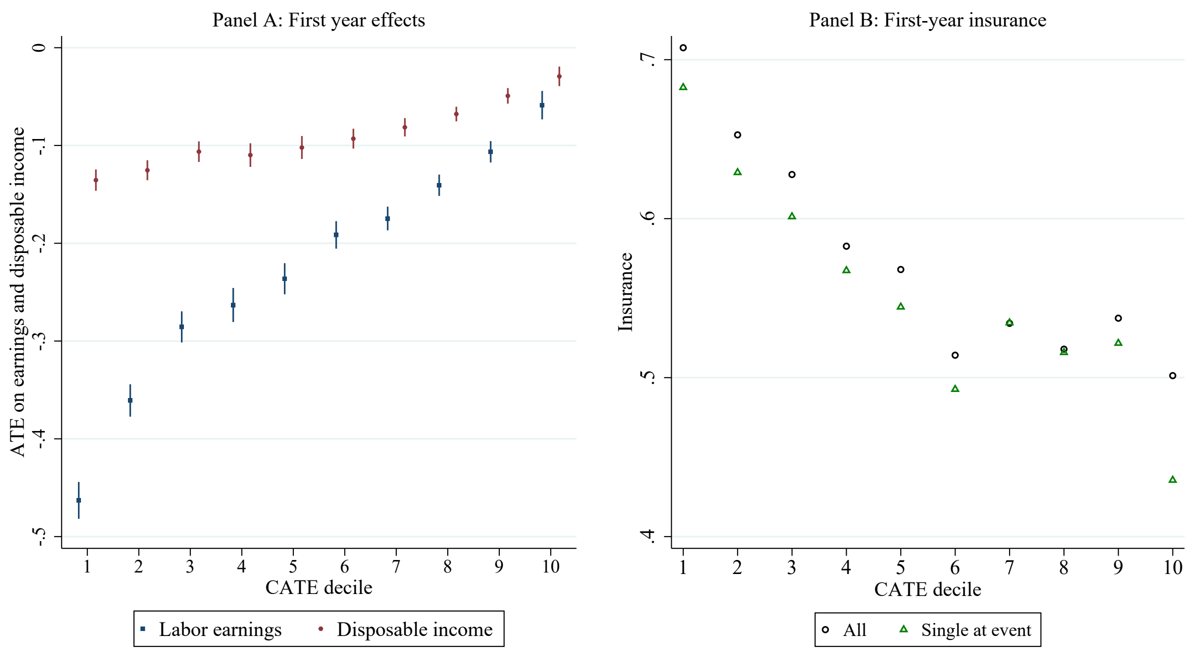

Note: The figure shows outcomes when dividing the training data set into deciles based on CATE:s estimated with 5-fold estimation. Panel A shows ATE estimates on labor earnings (as in figure 4) and corresponding ATE estimates for individualized disposable income. Point estimates and 95 percent confidence interval with standard errors clustered at the establishment level. Panel B shows the implied degree of insurance, defined as the difference between the ATE on disposable income the ATE on labor earnings, divided by the ATE on labor earnings. It provides separate estimates for the full population of displaced and for singles without children (in the year before displacement).

Figure B.14: Differences in characteristics across CATE quartiles for disposable income

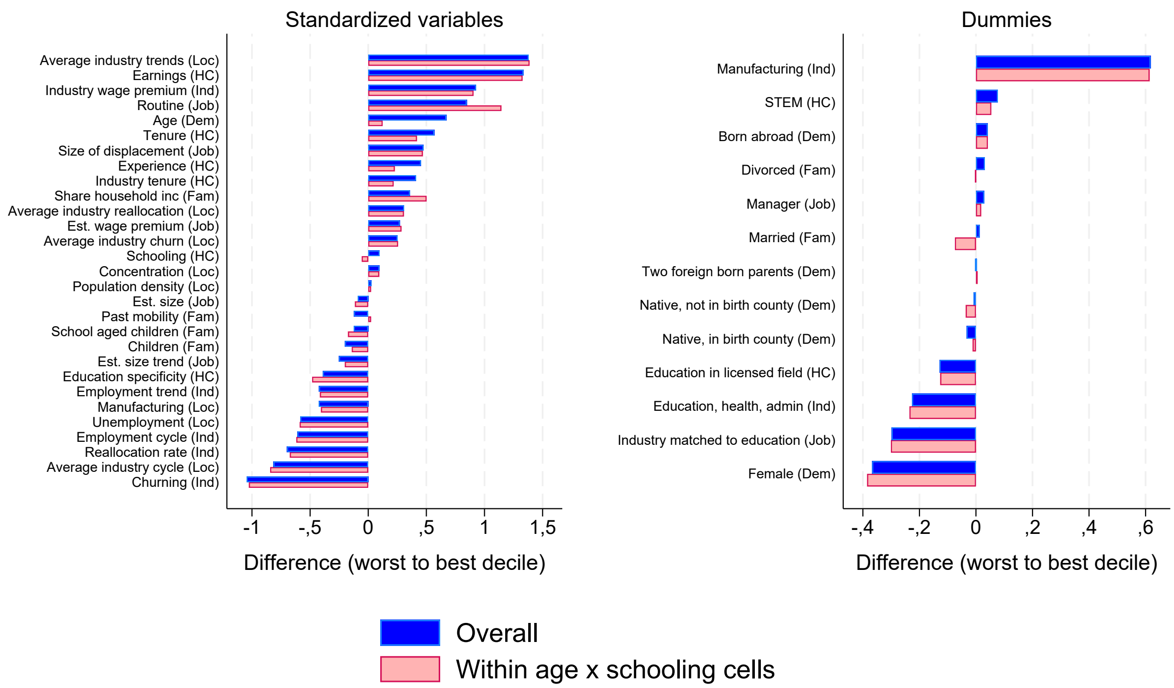

Note: The figure divides the sample by the CATE estimates for disposable income. It shows differences in characteristics between individuals in the lowest and highest quartiles and deciles of CATE:s, using the training data set. CATE:s estimated using 5-fold estimation and ranking done within each fold. The left-hand panel contains standardized (mean 0 and standard deviation 1) continuous variables and the right-hand panel contains dummy variables. Blue bars are for the overall quartiles and red bars cover the highest and lowest quartiles of CATE within each combination of 8 schooling and 10 age categories

C Online Appendix: Partial Effects of Market Conditions

This appendix describes the most important industry and location variables in more detail. Isolating the importance of market-level factors is not trivial since workers are likely to be sorted across regions and industries. For that reason we use the GRF to estimate partial dependence functions, holding the individual characteristics at their empirical levels and sequentially rotating across all observed sets of market characteristics within the same year. This way, we capture the predicted role of aggregate conditions across the full distribution of displaced workers, making full use of the GRF's non-linear nature. Because market-level factors are correlated, we split locations and industries into groups, and then characterize effects and attributes for combinations of these groups.

Formally, we divide the vector of characteristics X into parts related to the location (X^L), the industry (X^S), or the displaced worker and lost job (X^I). Worker i was displaced from location $l(i)$ and industry $s(i)$, with $X^L_{l(i)}$ and $X^S_{s(i)}$ being the corresponding location and industry characteristics. Then we compute:

$$\tau^L_l = \frac{1}{N}\sum_{i=1}^{N} \text{CATE}(X^L = X^L_l, X^S_{s(i)}, X^I_i), \quad \tau^S_s = \frac{1}{N}\sum_{i=1}^{N} \text{CATE}(X^L_{l(i)}, X^S = X^S_s, X^I_i), \tag{4}$$

where N is the number workers in the sample. τ^L_l and τ^S_s are the average CATEs if all N workers experienced the conditions in location l/industry s. We use these predictions to rank industries and locations by τ^L and τ^S respectively.

We classify workers using a similar strategy. Here, we hold worker characteristics fixed and combine them with each observed combination of location and industry conditions. Then for each worker i we compute:

$$\tau^I_i = \frac{1}{N^L \times N^S}\sum_{l=1}^{N^L}\sum_{s=1}^{N^S} \text{CATE}(X^L_l, X^S_s, X^I_i), \tag{5}$$

where N^L and N^S are the number of locations and industries respectively. τ^I_i only reflects worker-level characteristics since we average over the same set of location and industry conditions for all workers (by year).

In the end, we compute how ATEs, and characteristics, vary with (combinations of) quartiles of τ^L, τ^S and τ^I. This strategy allows us to reduce the dimensionality of the heterogeneity and study how different types of heterogeneity interact.

C.1 ATE Estimates by Partial Market and Worker Quartiles

Table C.1 shows estimated ATEs for different cuts of the data, using combinations of τ_l^L, τ_s^S and τ_i^I. Panel A shows estimates for the lowest (Column 1) and highest (Column 2) quartile of worker-level effects (τ_i^I). Column 3 shows the inter-quartile differences, with standard errors in Column 4. The top row shows that workers in the lowest quartile of τ_i^I experience 26 pp larger losses than the top quartile.

In the rows below, we zoom in on workers exposed to different combinations of location and industry conditions. We see significant differences in ATEs across worker types within each set of market conditions (ranging from 0.30 to 0.16). Market conditions are particularly important predictors for workers with poor individual characteristics.[43] The ATEs vary nearly as much when comparing across market (industry and location) conditions for the lowest worker quartile (25 pp difference, comparing across rows) as when comparing across worker quartiles within the worst market (industry and location) quartile setting (30 pp).

In Panels B and C, the columns instead represent location and industry (defined at the 3-digit level) quartiles, respectively. Here, the different rows represent worker types.[44] Panel B shows that the location matters for the size of the displacement effect even when holding the worker quartile (τ_i^I) fixed (interquartile differences are 6 to 12 pp depending on type). Workers also suffer significantly larger effects if displaced from "bad" industries (Panel C), in particular if they belong to the bottom worker quartile. In this group, estimated losses are 19 pp larger if displaced in a low-quartile industry as compared to a high-quartile industry.

Market conditions should matter less (or not at all) if workers are mobile. We therefore estimate how displacement events affect mobility across locations and industries. Panel D shows that workers are more likely to change location if displaced under worse conditions. The estimated effects on location mobility are small, partly because location mobility overall is very low, but the results show that this is true even for workers displaced in very bad locations. The results also indicate that job loss leads to more mobility across locations among workers who are more resilient.

Panel E shows that workers are more likely to move away from their (1-digit) industry if they are displaced from an industry with large displacement effects. This analysis is performed on the endogenous subsample of workers who find new employment, which

[43]See Appendix Figure B.12 for similar results on the importance of establishment heterogeneity for low ranked workers.

[44]Standard errors for the cross-quartile differences are clustered at location (Panel B) and industry (Panel C) since the number of locations and industries is small relative to the sample.

warrants some caution. With this caveat in mind, the results suggest that cross-industry mobility due to job loss is more common among workers in industries where effects are larger; 33 (23) percent move because of displacement in industries in the worst (best) industry quartile. The magnitudes here are much more substantial than for location mobility thorughout, which is natural since the old job disappears, but (other) ties to the location remain. Even though we find that industry mobility responds to displacement during distressful conditions, the fact that industry characteristics in general correlate so strongly with displacement effects suggests that the degree of industry mobility is insufficient to offset the negative effects of being displaced in a bad industry.

Table C.1: Earnings and mobility effects, by partial market and worker characteristics

	(1) Worst Quartile	(2) Best Quartile	(3) Interquartile Difference	(4) Standard Error of Difference
Earnings Estimates				
Panel A: Worker Quartiles, earnings effects				
All Market types	-0.370	-0.114	-0.256	0.008
Worst Market quartiles	-0.502	-0.200	-0.302	0.090
Median Markets	-0.361	-0.149	-0.211	0.026
Best Market quartiles	-0.254	-0.093	-0.162	0.019
Panel B: Location Quartiles, earnings effects				
All worker types	-0.298	-0.179	-0.119	0.016
Worst worker quartile	-0.424	-0.306	-0.118	0.012
Median workers	-0.262	-0.179	-0.083	0.017
Best worker quartile	-0.158	-0.100	-0.059	0.029
Panel C: Industry Quartiles, earnings effects				
All worker types	-0.330	-0.180	-0.150	0.020
Worst worker quartile	-0.463	-0.274	-0.188	0.022
Median workers	-0.286	-0.183	-0.102	0.020
Best worker quartile	-0.152	-0.102	-0.049	0.037
Mobility Estimates				
Panel D: Location quartiles, effects on location mobility				
All worker types	0.017	0.005	0.012	0.002
Worst worker quartile	0.011	0.011	0.000	0.002
Median workers	0.020	0.003	0.017	0.004
Best worker quartile	0.022	0.003	0.019	0.006
Panel E: Industry quartiles, effects on industry mobility				
All worker types	0.332	0.225	0.107	0.027
Worst worker quartile	0.411	0.254	0.158	0.029
Median workers	0.309	0.242	0.067	0.026
Best worker quartile	0.239	0.183	0.056	0.047

Note: Panels A–C display displacement effects on earnings on year after displacements (calculated as displaced-control differences, ”ATEs”). Panel A shows earnings estimates for the worst quartile of workers (Column 1) and the best quartile (Column 2), using the worker resiliency measure described above. Results for workers in all markets (interaction between industry and location), for workers in the worst (best) market quartiles, and for median markets (workers in the 3d or 4th market quintiles). Column 3–4 show the difference between Columns (1) and (2), and the standard error for the difference (clustered at the level of the pre-displacement establishment). Panel B shows earnings estimates for the best and the worst quartiles of locations using the partial-effects procedure described above. Panel C shows similar earnings measures when dividing by the best and worst industries. Both Panel B and C report estimates for all workers and when dividing workers by worker resiliency using the worker resiliency measures described above. Median workers are those in the 3d or 4th worker quintiles. Panel D shows displacement effects on location mobility (mobility defined as moving to another local labor market between one year before and three years after the displacement), and Panel E displacement effects on industry mobility (mobility defined as switching 1-digit industry between one year before and three years after the displacement). Panels D–E show results for the full sample of workers and when dividing the sample by worker resilience as in Panels B–C. Standard errors in Panel B and D clustered at the location level, and standard errors in Panel C and E clustered at the industry level. All estimates use the main data set. Rankings of workers and reshuffling done using the k-folds procedure.

C.2 Characteristics of Markets with Large Partial Effects

Table C.2 describes the best and worst locations and industries (top and bottom quartiles of τ_l^L and τ_s^S). Panel A shows location characteristics (X^L) for the "good" and "bad" locations as well as cross-quartile differences in X^L. Standard errors for differences are clustered at the location level. Locations with bottom-quartile partial effects (τ_l^L) are in particular characterized by much lower population density. These locations also have high unemployment rates and a more concentrated industry structure dominated by declining industries and manufacturing jobs.

Panel B shows similar results for industry characteristics. Industries with large predicted earnings losses for the average worker (τ_s^S) are exclusively found in manufacturing, while industries with small predicted losses are in non-manufacturing sectors. Industries with very negative τ_s^S-estimates also have higher wage premia, are less dynamic (lower churning and reallocation rates), and experience declining employment trends over both the short and the long run. All of these attributes are typical of manufacturing, but they are also related to effect heterogeneity if we only compare different manufacturing industries to each other, or if we only compare non-manufacturing industries to each other (see Table C.3).

Table C.2: Location and industry characteristics for the worst and best locations/industries

	(1) Worst Quartile	(2) Best Quartile	(3) Interquartile Difference	(4) Standard Error of Difference
Panel A: Location characteristics, by location quartiles				
Population density	20.410	147.913	-127.503	3.026
Unemployment rate	0.105	0.062	0.043	0.003
Industry concentration (HHI)	0.040	0.025	0.015	0.001
Share manufacturing	0.230	0.106	0.124	0.016
Average industry trend	0.058	0.113	-0.055	0.005
Average industry cycle	0.005	0.011	-0.006	0.001
Average industry churn	0.203	0.268	-0.065	0.004
Average industry reallocation	0.144	0.161	-0.017	0.002
Panel B: Industry characteristics, by industry quartiles				
Employment trend	-0.150	0.215	-0.366	0.061
Employment cycle	-0.033	0.024	-0.057	0.006
Industry wage premium	0.067	-0.049	0.116	0.027
Churning	0.177	0.253	-0.076	0.021
Reallocation rate	0.116	0.164	-0.048	0.008
Manufacturing	1.000	0.000	1.000	0.000
Education, health, admin	0.000	0.240	-0.240	0.113

Note: Panel A displays average location characteristics for the best and the worst quartiles of locations, and Panel B displays average industry characteristics for the best and the worst quartiles of industries. Both panels use the partial-effects procedure described above. The location and industry characteristics are described in Section 3.2. Standard errors in Panel A clustered at the location level, and standard errors in Panel B clustered at the industry level. All estimates use the main data set.

Table C.3: Characteristics for the worst and best manufacturing and non-manufacturing industries

	Manufacturing			Non-manufacturing		
	Worst Quartile (1)	Best Quartile (2)	Interquartile Difference (3)	Worst Quartile (4)	Best Quartile (5)	Interquartile Difference (6)
Employment trend	-0.236	0.080	0.316	0.004	0.242	0.238
Employment cycle	-0.036	-0.007	0.030	-0.008	0.022	0.030
Industry wage	0.085	0.036	-0.049	0.090	-0.033	-0.123
Churning	0.179	0.160	-0.019	0.212	0.245	0.033
Reallocation rate	0.113	0.117	0.004	0.138	0.155	0.018
Manufacturing	1.000	1.000	0.000	0.000	0.000	0.000
Education, health, admin	0.000	0.000	0.000	0.004	0.375	0.370

Note: The table displays average industry characteristics for the best and the worst quartiles of manufacturing industries (Columns 1–3) and non-manufacturing industries (Columns 4–6). Both panels use the re-shuffling procedure described above. The industry characteristics are described in Section 3.2. Standard errors clustered at the industry level. All estimates use the main data set.

Table C.4 describes the characteristics of the top and bottom partial worker-level quartiles (τ_i^I) and compare these to differences across raw CATE quartiles.[45] Overall, the results show relatively small differences between the partial and overall quartiles. This suggests that a very small share of the differences in treatment effects associated with worker characteristics arises because of correlations with market factors. The table reaffirms the strong relationships with age, schooling, and specific human capital. The most striking difference is related to gender. Females are highly over-represented in the lowest *partial* quartile (τ_i^I), despite not being in the lowest *overall* CATE quartile. The (statistical) reason is that females are underrepresented in manufacturing where effects tend to be larger. There are also fewer immigrants in the best partial quartile than in the best overall CATE quartile, reflecting both an under-representation of immigrants in manufacturing and an over-representation of immigrants in metropolitan areas where effects tend to be more muted. Similarly, an over-representation of STEM graduates in the most resilient partial quartile reflects that many engineers work in manufacturing, but still suffer from relatively modest effects.

[45] In contrast to Figure 6, we use raw non-standardized numbers in this table.

Table C.4: Worker characteristics for resilient and non-resilient workers

	Workers ranked by *partial* CATE			Workers ranked by CATE		
	(1) Worst Worker Quartile	(2) Best Worker Quartile	(3) Difference [1]-[2]	(4) Worst Worker Quartile	(5) Best Worker Quartile	(6) Difference [1]-[2]
Demographics						
Age	52.448	35.286	17.162	50.728	35.900	14.827
Female	0.403	0.314	0.089	0.366	0.366	0.000
Born abroad	0.168	0.057	0.111	0.163	0.091	0.072
Two foreign born parents	0.024	0.033	-0.009	0.025	0.037	-0.012
Native, not in birth country	0.215	0.307	-0.092	0.201	0.297	-0.096
Family						
Married	0.637	0.539	0.097	0.629	0.526	0.103
Divorced	0.162	0.045	0.116	0.146	0.058	0.088
Children	0.317	0.916	-0.599	0.394	0.864	-0.470
School-aged children	0.235	0.492	-0.257	0.287	0.475	-0.188
Share houshold inc.	0.716	0.798	-0.081	0.724	0.788	-0.064
Past mobility	0.089	0.360	-0.271	0.101	0.344	-0.244
Human capital						
Schooling (years of education)	10.073	12.986	-2.913	10.094	12.704	-2.611
Experience	9.381	8.578	0.803	9.341	8.510	0.831
Tenure	6.676	5.220	1.456	6.763	5.172	1.590
Industry tenure	7.926	6.456	1.470	7.946	6.445	1.501
Earnings (rank)	0.364	0.646	-0.282	0.377	0.587	-0.210
STEM education	0.031	0.254	-0.224	0.032	0.203	-0.172
Education in licensed field	0.058	0.088	-0.031	0.044	0.106	-0.062
Education specificity	0.428	0.528	-0.100	0.420	0.528	-0.108

Note: The table shows average worker characteristics differ for the best and worst quartiles of workers. Columns 1–3 rank workers by CATE and Columns 4–6 by the partial worker resiliency measure described above. Columns 2–3 and 4–5 display averages within the respective quartiles. Column (3) shows differences between Columns (1) and (2). The worker characteristics are described in Section 3.2. All sample statistics are for the main data set. Rankings of workers and reshuffling done using the k-folds procedure.