Technological Change, Firm Heterogeneity and Wage Inequality*

Guido Matias Cortes[†] Adrian Lerche[‡]

Uta Schönberg[§] Jeanne Tschopp[¶]

April 5, 2021

Abstract

We argue that task-biased technological change does not only differentially affect wages of workers in different tasks, but is also an important driver of the rise in wage inequality within tasks, for workers in different firms. In a framework that allows for heterogeneity between firms in their task composition and in the wage that they pay for workers in a given task, we show that an industry-wide task-biased technological change shock will increase wage inequality between firms within an industry through four channels: selective exit of firms; increased employment concentration in more productive firms; increased wage dispersion between firms for workers of the same task type; and increased dispersion in the task mix that firms employ, due to more sorting of workers in highly rewarded tasks into more productive firms. Using rich administrative matched employer-employee data from Germany spanning two decades, we provide empirical evidence of establishment-level adjustments that are in line with the predictions of the model. We further document that industries with more technological adoption exhibit particularly pronounced adjustment patterns along the dimensions highlighted by the model.

*Cortes and Tschopp are grateful to the Social Sciences and Humanities Research Council of Canada for its financial support. Uta Schönberg is grateful to the European Research Council for its financial support (Grant Number 818992 (FirmIneq)). We thank the Research Data Centre of the German Federal Employment Agency at the Institute for Employment Research for providing data access.

[†]Department of Economics, York University, 4700 Keele Street, Toronto, ON, M3J 1P3, Canada; gmcortes@yorku.ca; +1 (416) 736-2100 (ext 33987).

[‡]Institute for Employment Research (IAB), Regensburger Strasse 104, 90478 Nürnberg, Germany; adrian.lerche@iab.de; +49 911 179 9213.

[§]Department of Economics, University College London, Gower Street, London, WC1E 6BT, United Kingdom; u.schoenberg@ucl.ac.uk, +44 203 549 5351 and Institute for Employment Research (IAB), Regensburger Strasse 104, 90478 Nürnberg, Germany.

[¶]Department of Economics, University of Bern, Schanzeneckstrasse 1, 3001 Bern, BE, Switzerland; jeanne.tschopp@vwi.unibe.ch; +41 31 631 80 81.

1 Introduction

Income inequality has risen dramatically over the past decades in many high-income countries. The development of new technologies has been proposed as a key driver of this pattern. Earlier work on skill-biased technological change (e.g. Katz & Murphy, 1992) argued that technology is complementary to skilled workers. By increasing the (relative) demand for skilled labor, new technologies were argued to increase wage inequality through their impact on the skill premium. A more recent wave of literature has argued that technology is task-biased. Autor et al. (2003) and the subsequent literature on labor market polarization (e.g. Autor et al., 2006; Goos et al., 2014; Jaimovich & Siu, 2020) has discussed how progress in automation and digitization technologies has substituted for workers in occupations that are intensive in routine, easily codifiable tasks. This implies that technology has impacted wage inequality by changing the occupational structure of the economy and the relative wage returns for different tasks.

Intriguingly, however, a recent parallel literature has shown that most of the rise in wage inequality can be traced backed to increasing wage differentials among observationally similar individuals, working in different firms (Card et al., 2013; Barth et al., 2016; Helpman et al., 2017; Song et al., 2019). On the one hand, this literature has argued that individual-level wages have become increasingly dependent on where people work, rather than the type of tasks that they do. On the other hand, this literature has demonstrated that an important proximate reason behind the rise in (between-establishment) wage inequality is the fact that different types of workers have become increasingly segregated across firms, with high-wage workers increasingly clustered in high-wage firms (Card et al., 2013; Song et al., 2019). While this literature has documented many novel empirical facts, it has so far been largely silent on the driving forces behind these findings. In comparison, while the literature on skill- and task-biased technological change has provided deep insights into the evolution of relative wage gaps *across* skill and task groups, it has, by drawing on models of perfectly competitive labor markets with a representative firm, so far abstracted from increasing wage differentials observed *within* groups, across firms.

In this paper, we show that the development of task-biased automation technologies can account not only for increases in between-task inequality, but also for increases in within-task, between-firm inequality, as observed in the data. We show this theoretically using a rich yet tractable heterogeneous firm framework, and empirically verify the predictions of the model using administrative social security data from Germany.

Our data, drawn from the so-called IAB Beschäftigtenhistorik (BEH), cover the universe of private sector workers and establishments in West Germany between 1990 and 2010. We supplement these data with information from the IAB Establishment Panel (IABEP), which provides measures of establishment-level sales and allows us to construct a measure of labor productivity for the establishments covered by the survey.

We begin by verifying the evidence for task-biased technological change, as reflected by the evolution of the aggregate occupational structure of the economy. As has been shown in the literature, there has been a substantial shift in employment away from routine occupations and towards abstract occupations. Importantly, however, we show that essentially all of the increase in wage inequality observed in Germany between 1990 and 2010 is due to growing wage differentials among workers in the same broad task group. The primary importance of within-group changes is also verified when considering much finer 3-digit occupational categories. In line with previous literature, the rise in within-task inequality is primarily due to growing wage differences between, rather than within establishments. More than half of the rise in within-task between-establishments wage inequality is observed across establishments operating within the same 3-digit industry.

These findings motivate us to explore the link between task-biased technological change, and the evolution of between-establishment wage inequality. In order to do this, we set up a model that allows for wage heterogeneity between firms, with a production structure that distinguishes between two tasks, and where firm productivity and task usage are linked.

The theoretical framework that we consider is a version of the model in Helpman et al. (2010). The model introduces search and matching frictions (Diamond, 1982a,b; Mortensen & Pissarides, 1994), as well as match-specific abilities and a screening technology to the heterogeneous firm setting of Melitz (2003). These features generate wage differentials between firms within industries (overall and conditional on worker task) in a rich, yet tractable way. The model predicts that, in equilibrium, more productive firms will find it optimal to employ more workers of both types, but will have a higher abstract employment share. They will also pay higher wages to both types of workers. These predictions are consistent with what we observe in the data.

Our key innovation on the theoretical front is to introduce an aggregate task-biased technological shock in the spirit of Autor et al. (2003) within this type of rich heterogeneous firm framework. The shock increases the relative productivity of abstract workers and, although it is an aggregate shock that is common across all firms, we show that it leads to an increase in between-firm wage inequality, which occurs through four distinct channels.[1] First, task-biased technological change will increase the productivity threshold for production,

[1] It is worth emphasizing that, given our interest in wage inequality, we focus solely on the *relative* impacts of technology across workers. Some recent evidence on the *absolute* impacts of automation technologies are provided, for example, by Acemoglu & Restrepo (2020a,b).

which results in the exit of low-productivity firms. Although this reduces the range of productivities of operating firms, the assumption of Pareto-distributed productivity implies that this change in the composition of firms increases the variance of productivity among operating firms. The second channel is differential employment growth, whereby the most productive, highest-paying firms are predicted to grow more, thus contributing to an increase in worker-weighted measures of between-firm inequality. The third channel is increased sorting: As a reaction to the shock, the most productive firms will disproportionately increase the share of abstract workers that they employ. Since abstract workers earn higher wages, this will exacerbate the differences in average wages between firms. Since more productive firms were already more abstract-intensive at baseline, task-biased technological change will also increase the dispersion in the abstract employment share across firms. Finally, the model also generates endogenous wage changes conditional on task, with more productive firms disproportionately increasing the wage that they pay to workers of each type, thus further contributing to the increase in between-firm wage inequality.

Similar channels are at work in driving the rise in the aggregate abstract employment share in an industry. In particular, the aggregate abstract employment share does not only increase because of within-firm changes, but also because of the selective exit of routine-intensive firms and the differential growth of more productive abstract-intensive firms.

Guided by the model, we return to the BEH and IABEP data and verify its key predictions. In line with the model, we find that in the cross-section and within detailed industries, more productive establishments are larger, employ more abstract workers and pay higher wages not only overall but also to each task type. Interestingly, and consistent with the predictions from the model, these within-industry establishment-level associations between productivity, size, abstract shares and wages have become stronger over our sample period. In line with this evidence, when considering longitudinal changes within establishments, we find that larger establishments tend to pull further away from smaller establishments in the same industry, by experiencing faster productivity growth, by disproportionately employing more abstract workers, and by disproportionately increasing the wages that they pay to the same worker type.

We further show that, in line with the model, establishments within industries have become increasingly heterogeneous in terms of their task mix, rather than converging to more similar technologies of production. This increased dispersion in task mix is driven by increased sorting of abstract workers towards establishments that pay high wage premiums conditional on worker task, which arises endogenously due to task-biased technological change.[2]

[2]See also Wilmers & Aeppli (2021), who show that workers in high-paying occupations in the U.S. are

Guided by the model, we then perform a series of decompositions that allow us to assess the relative importance of the different channels that drive the increase in between-establishment wage inequality. First, we determine the role of entry and exit relative to changes among continuing establishments. We find that changes in establishment composition due to entry and exit have played a non-trivial role in driving the increase in between-establishment wage inequality; changes among continuing establishments, however, are quantitatively much more important. We then analyze the role of sorting along task dimensions. Although we focus on a setting with only two tasks, we find that changes in the sorting of abstract workers towards high wage premium establishments account for 12% of the change in the wage variance among continuing establishments within industries.[3] Finally, we determine the importance of differential employment and wage changes within continuing establishments and find that differential employment growth can account for around half of the rise in the variance of wage premiums among continuing establishments. Thus, an important reason for the rise in between-establishment wage inequality is the fact that establishments that paid higher wages at baseline have expanded in size relative to lower paying establishments – a channel that the literature has so far ignored. Similarly, the selective exit of routine-intensive establishments and the changing reallocation of abstract workers to establishments of different sizes substantially contribute to the aggregate rise in the abstract employment share in an industry.

As a final exercise, we provide direct evidence of the link between task-biased technological change and the establishment-level patterns that we have identified. We do this by exploiting variation across industries in technology adoption, which we measure based on the change in each industry's abstract employment share over our sample period, based on industry-level robot adoption data from the International Federation of Robotics, or based on ICT capital usage data from EUKLEMS. Our key finding is that industries that have adopted more technology have experienced disproportionate increases in between-establishment wage inequality, abstract share heterogeneity, and the sorting of abstract workers to high-wage establishments. This confirms the key role of technology adoption in driving the establishment-level patterns that we have documented.

Our findings make important contributions to several strands of the literature. First, we provide an important innovation to the literature on task-biased technological change (Autor et al., 2003, 2006; Goos et al., 2014; Michaels et al., 2014; Acemoglu & Restrepo, 2020a), by considering the implications of this type of shock within a heterogeneous firm framework that allows for wage differentials between firms for workers in a given task. By

increasingly likely to be employed at high-paying workplaces.

[3] We also show results based on detailed occupational categories, where sorting plays an even larger role.

embedding a task-biased technological shock within a framework that departs from the traditional representative firm frameworks with perfect competition that have been considered in this literature, we are able to show that task-biased technological change not only differently affects wages of workers in different tasks, but can also account for the quantitatively much more important rise in inequality within tasks, for workers in different firms. Our results also paint a richer picture about the individual-level impacts of task-biased technological change, by highlighting that the relative impact across individuals will depend not only on the task that they perform, but also the type of firm that they are matched to. For example, routine workers employed in high-productivity establishments lose out relative to abstract workers in these workplaces, but gain relative to routine workers in low-productivity ones.

Our analysis also provides an important contribution to the literature on the rise in between-firm wage inequality (Card et al., 2013; Song et al., 2019; Barth et al., 2016; Helpman et al., 2017; Criscuolo et al., 2020). While this literature has so far been very successful in highlighting the increasing importance of firms for individual wages, it has been largely silent on the underlying driving forces behind these patterns. We contribute to this literature by considering a tractable theoretical framework that allows us to study the interplay between task-biased technological change and these important workplace-level patterns. We provide novel empirical findings that indicate that task-biased technological change is an important driver of the rise in between-establishment wage inequality in Germany over recent decades. Guided by the model, we also document new findings regarding the role of establishment entry and exit, the role of sorting along task dimensions, and the role of differential employment growth across establishments for wage inequality.

The findings that we present can also be linked to the literature on the rise in concentration and the increased dominance of so-called superstar firms (Autor et al., 2017, 2020; Azar et al., 2020a,b; Bajgar et al., 2019). We show that task-biased technological change leads to the disproportionate growth of the most productive workplaces within an industry. Technological change may therefore be at least partly responsible for the rise in concentration. We also show how the employment shift towards 'superstar' establishments directly contributes to the rise in wage inequality: even in the absence of any wage changes within establishments, the fact that more productive, higher wage workplaces have experienced differential employment growth has led to an increase in measures of worker-weighted between-establishment wage inequality.[4]

Our results are also relevant to the outsourcing literature. This literature has analyzed how workplaces increasingly tend to concentrate on their core tasks, outsourcing secondary

[4]See Webber (2015); Mueller et al. (2017); Rinz (2020); Cortes & Tschopp (2020) for more detailed analyses of the link between rising concentration and rising wage inequality.

tasks to other firms, and leading to increased specialization along task dimensions within workplaces (e.g. Goldschmidt & Schmieder, 2017; Cortes & Salvatori, 2019). In our setting, the re-sorting of workers across workplaces due to task-biased technological change leads to this type of increased task specialization, as reflected by the increased heterogeneity in abstract employment shares across workplaces. Hence, our model generates patterns that are consistent with increased outsourcing, but that arise endogenously as a result of an aggregate task-biased technological change shock (rather than an exogenous decrease in outsourcing costs).

Finally, our paper builds on the major advances made in terms of the analysis of heterogeneous firm frameworks (e.g. Melitz, 2003; Yeaple, 2005; Egger & Kreickemeier, 2009, 2012; Helpman et al., 2017; Trottner, 2019). These models have primarily been used in order to study the impact of trade liberalization, and papers in this literature have provided a very rich set of results regarding the interplay between international trade and various firm-level outcomes, including wages. Here we show that these types of models can also be very useful in terms of understanding the interplay between firm-level heterogeneity and other types of shocks, such as technological change.

2 Data

2.1 Social Security Records (Beschäftigtenhistorik (BEH))

Our main data are drawn from social security records provided by the Institute for Employment Research in Nuremberg (IAB) – the so-called Beschäftigtenhistorik (BEH). While social security records are in principle available for the years 1975 to 2014 (1992 to 2014 for East Germany), we focus on developments after 1990 when overall wage inequality started to increase sharply in Germany (see for example Dustmann et al. (2014); Card et al. (2013)). Due to structural breaks after 2010 in key variables such as occupations (used to classify workers into routine and abstract ones) and workers' full-time status, we end the analysis in 2010. The data source comprises all men and women covered by the social security system – roughly 80% of the German workforce. Not included are civil servants, the self-employed, and military personnel.

Our data source offers some key advantages. A first advantage is its large size, allowing us to accurately capture trends in wage inequality even within detailed industries. Second, our data contain comprehensive and accurate information on a number of worker and establishment characteristics that are not always included in other administrative data sources, such as, for example, workers' occupation, education, employment status and wages (which always refer to a single establishment and are never averaged across establishments) and establishments' industry affiliation. Importantly, unique establishment identifiers allow us not only to decompose overall wage inequality into a within and a between establishment component, but also to study (changes in) establishment heterogeneity within industries more broadly. Establishment identifiers also allow us to paint an accurate picture of entry and exit across industries and time.

From this data source, we select all full- and part-time employment spells that refer to June 30 of each year. We restrict the sample to workers who are currently not in an apprenticeship, are aged between 16 and 65, and are employed in West Germany. We exclude industries in the primary sector and some small industries such as private households and international organizations. We further drop workers with missing occupation, missing employment status, or implausibly low wages below the limit for which social security contributions have to be paid, as well as establishments with missing industry affiliation and establishments employing only part-time workers. These sample restrictions affect less then 1% of all observations.

As is common in administrative data sources, wages are censored at the highest social security limit, affecting on average about 8% of observations. We follow Dustmann et al. (2009) and Card et al. (2013) and impute censored wages, assuming that (log) wages are normally distributed with heterogeneous variances that vary by year, age, education and sex; see Appendix A.1 for details. We deflate wages using 1995 as the base year. Since we do not observe hours worked, we restrict the wage analysis to full-time workers. To compute total employment in establishments and industries (overall or by task), we include part-time workers with a weight of 0.5.

We follow Acemoglu & Autor (2011) to classify individuals into task groups based on their broad occupation codes. For simplicity, we consider two groups: abstract workers (which include all workers in professional, managerial or technical occupations) and all other workers, which we label as "routine" (these include workers in administrative and clerical jobs, production workers, and workers in personal service occupations). Appendix Table A.1 provides details on the mapping of occupation codes to task groups. The two most common abstract occupations are nurses and managers, accounting for more than 15% of all abstract workers. The most common routine occupation is office clerks, comprising 16.9% of all routine workers. A worker's education and the task she is employed in are strongly correlated. Whereas 37.0% of all abstract workers hold a university degree, this is the case for only 6.8% of routine workers.

Our industry classification refers to 3-digit NAICS codes which distinguish between 196 industries. Due to a structural break in the industry classification in the social security data in 1999, we harmonize the industry classification as described in Appendix A.2.

For the main empirical analysis, we aggregate the worker level information to the level of the establishment (by year). We thereby create an establishment panel which records industry affiliation, and tracks entry and exit, size, employment share of abstract workers, and average wages in each establishment over time.

To isolate whether establishments pay higher overall wages because they employ more abstract workers, or because they pay higher wages to the same type of worker, we additionally compute an establishment wage premium as follows: First, we estimate, separately for each year, a regression of individual-level log wages on a task indicator (equal to one for abstract workers) interacted with a full set of 3-digit industry fixed effects, thereby allowing for different abstract wage premiums across industries (and years). We then compute the average residual for each establishment. We refer to this as 'establishment premium (tasks)'

Since workers may substantially differ within broad task types, we compute a second establishment wage premium by estimating, separately for each year, a regression of individual-level log wages on a full set of 3-digit occupation fixed effects (317 occupations) interacted with a full set of 3-digit industry fixed effects, and then calculating the average residual for each establishment. This second establishment wage premium shows whether different establishments pay different wages to workers within the same detailed occupation group, and therefore rules out that the establishment wage premium reflects differences in the occupational structure (at a detailed level) across establishment. We refer to this measure as 'establishment premium (occupations)'.

2.2 The IAB Establishment Panel (IABEP)

Since the social security records drawn from the BEH do not contain information on establishment outcomes such as total sales or labor productivity, we augment the social security records with data from the IAB Establishment Survey (IABEP). The IABEP survey was first administered in 1993 to 4,265 West German establishments, and was extended to East German establishments in 1996. By 2010, the number of surveyed establishments had increased to over 16,000. From this database, we select all West German establishments with at least one full-time employee that participated in the IABEP at least once. Adopting the same sample selection criteria as in the social security records (BEH), we drop establishments with missing industry affiliation as well as establishments in the primary sector and some smaller sectors such as private households and international organizations. Using the unique establishment identifiers, we then merge information from the BEH social security records to the IABEP. We compute an establishment's labor productivity as total sales (obtained from the IABEP), divided by the number of full-time equivalent workers (obtained from the BEH). In the empirical analysis based on the IABEP, we use the weights provided by the survey in order to guarantee representativeness for workers.

2.3 Industry-Level Technology Adoption Measures

We supplement these two main data sources with industry-level data on technology adoption. First, following Graetz & Michaels (2018) and Acemoglu & Restrepo (2020a), we use data on robot usage from the International Federation of Robotics (IFR).[5]

Second, we use data on the adoption of capital related to information and communication technologies (ICT) from the EUKLEMS data set. We use data from the November 2009 release, which uses ISIC revision 3 industry codes which can be matched to the 2-digit industry codes in the BEH social security data. Our measure of ICT assets is based on the real fixed capital stock of computing and communication equipment, and computer software.

3 Motivating Evidence

This section presents some motivating empirical facts related to the evolution of the task structure of employment and wage inequality in Germany between 1990 and 2010. These empirical patterns motivate the setup of our theoretical framework and the subsequent analysis of the impact of technology on (between-establishment) wage inequality.

Task Usage. We begin by verifying the substantial shift in the task structure of employment in Germany away from routine occupations and towards abstract occupations – a pattern that has been well documented for many developed countries (see Acemoglu & Autor, 2011; Goos et al., 2009). Figure 1 shows that the aggregate employment share of abstract workers steadily rose from about 18% in 1990 to more than 26% in 2010 – a rise of 44% over two decades. This increase was in part driven by differential industry growth: industries which employ a larger share of abstract workers grew at a faster rate than industries which predominantly employ routine workers. Yet, even when keeping the industry structure constant at 1990 levels (the grey dashed line; see Appendix B.1 for details), the employment share of abstract workers rose substantially by about 27%.

Overall and Between-Task Wage Inequality. Panel A of Figure 2 illustrates the

[5] A robot is defined as an "automatically controlled, re-programmable, and multipurpose machine" and as "fully autonomous machines that do not need a human operator and that can be programmed to perform several manual tasks such as welding, painting, assembling, handling materials, or packaging."

evolution of wage inequality in West Germany between 1990 and 2010. The black line shows that the overall variance of individual log-wages rose sharply from the mid-1990s onwards, from 0.195 in 1995 to 0.276 in 2010, a 41.5% increase (see also Dustmann et al., 2009, 2014).

Panel A further decomposes the variance of log-wages into a within-task and a between-task component as follows:

$$\frac{1}{n}\sum_{i}(\ln w_i - \overline{\ln w})^2 = \underbrace{\frac{1}{n}\sum_{\ell}\sum_{i \in i_\ell}(\ln w_i - \overline{\ln w_\ell})^2}_{\text{Within tasks}} + \underbrace{\frac{1}{n}\sum_{\ell} n_\ell(\overline{\ln w_\ell} - \overline{\ln w})^2}_{\text{Between tasks}} \tag{1}$$

Here, ℓ represent different tasks and $\overline{\ln w_\ell}$ is the average log wage in task ℓ; i_ℓ is the set of individuals in task ℓ, n is the total number of workers, and n_ℓ is the total number of workers in task ℓ. The within-task component captures deviations of individual wages from the mean wage in their task, whereas the between-task component captures deviations of the mean task wage from the aggregate mean wage.

Panel A of Figure 2 shows that, at any given point in time, more than 90% of inequality is within-task. This is not too surprising, given that we consider only two broad task groups. More interestingly, though, the figure shows that essentially all of the increase in inequality is driven by increased wage heterogeneity among workers in the same task category. Although the literature on task-biased technological change (TBTC) focuses on changes in inequality that operate due to changes in the demand for different tasks – and hence changes in task premiums that should impact *between-task* inequality – we see that the between task component in Germany remained stable over this time period. In fact, as shown in Appendix Figure A.1, the abstract wage premium (i.e. the difference between the average log wage of abstract and routine workers) did not rise much in Germany over our sample period.

Panel A of Figure 3 shows an analogous decomposition using detailed 3-digit occupational codes (317 categories). These detailed occupational categories provide us with a much finer proxy for workers' task content. The results support the conclusions drawn from the analysis with two broad task categories. Although the variance between occupations is much higher than between the two task categories, wage differences *within* detailed occupational categories account for the majority of the overall log-wage variance in the cross-section, and account for at least half of the increase in the variance over time. Thus, wage inequality rose sharply *within* task groups, even when we consider very detailed occupational groups.

Wage Inequality Within Tasks: Within and Between Establishments. The finding that essentially all of the increase in inequality is due to increasing wage differences among workers in the same broad task group does not mean that TBTC is not an important driver of wage inequality in Germany. As we show below, in a setting with heterogeneous firms, TBTC can lead to an increase in inequality within tasks, across workers in different workplaces. Panel B of Figure 2 shows that the increase in the within-task component of the wage variance is indeed nearly entirely driven by increasing wage differences across workers in different establishments. In the figure, we decompose the within-task variance into a within-establishment and a between-establishment component as follows:

$$\frac{1}{n}\sum_{\ell}\sum_{i\in i_\ell}(\ln w_i - \overline{\ln w_\ell})^2 = \underbrace{\frac{1}{n}\sum_{f}\sum_{\ell}\sum_{i\in i_{f\ell}}(\ln w_i - \overline{\ln w}_{f\ell})^2}_{\text{Within tasks, within est.}} + \underbrace{\frac{1}{n}\sum_{f}\sum_{\ell} n_{f\ell}(\overline{\ln w}_{f\ell} - \overline{\ln w_\ell})^2}_{\text{Within tasks, between est.}} \tag{2}$$

where f indexes establishments, $i_{f\ell}$ is the set of individuals in task ℓ at establishment f, and $n_{f\ell}$ the total number of workers in task ℓ at establishment f. We find that 95% of the increase in the within-task wage variance between 1990 and 2010 can be attributed to the rise in the variance between establishments; increases in within-establishment wage differentials account for only 5% of the rise. This pattern is in line with the broader evidence in the literature regarding the increasing importance of between-firm wage differentials, documented by e.g. Card et al. (2013); Song et al. (2019); Barth et al. (2016); Helpman et al. (2017).

Using the detailed occupational categories, Panel B of Figure 3 confirms that the majority of the wage differences within detailed occupations are also between, rather than within establishments, and that essentially all of the increase in wage inequality within occupations over time is driven by the between-establishment component.

Wage Inequality Between Establishments Within Tasks: Within and Between Industries. Panel C of Figure 2 further decomposes the wage differentials observed within tasks across workers in different establishments into a component that is due to differences between establishments in the same industry and differences between establishments in different 3-digit industries:

$$\begin{aligned}\frac{1}{n}\sum_{f}\sum_{\ell} n_{f\ell}\left(\overline{\ln w}_{f\ell} - \overline{\ln w_\ell}\right)^2 = \\ \underbrace{\frac{1}{n}\sum_{k}\sum_{f}\sum_{\ell} n_{f\ell}(\overline{\ln w}_{f\ell} - \overline{\ln w}_{k\ell})^2}_{\text{Within tasks, between est., within industries}} + \underbrace{\frac{1}{n}\sum_{k}\sum_{\ell} n_{k\ell}(\overline{\ln w}_{k\ell} - \overline{\ln w_\ell})^2}_{\text{Within tasks, between est., between industries}}\end{aligned} \tag{3}$$

where k indexes industries and $n_{k\ell}$ is the total number of workers in task ℓ in industry k.

While both of the components are important, within-industry differences account for more than half of the within task-between establishment variance in the cross-section, and more than half of its change over time. Using the detailed occupational categories, Panel C of Figure 3 confirms that a significant part of the rise in within-occupation wage inequality between establishments also occurs within, and not between industries.

The within-industry increase in within-task between-establishment wage inequality could in principle be driven by industries with higher within-task between-establishment wage inequality growing at a faster rate than the average industry. To rule out this possibility, in Panel D of Figures 2 and 3 we display the counterfactual within-industry increase in within-task (or within-occupation) between-establishment wage inequality holding the industry structure constant at its 1990 level, which we compute as follows:

$$\sum_{k} \frac{n_{k1990}}{n_{1990}} \sum_{f} \sum_{\ell} \frac{n_{f\ell}}{n_k} (\overline{\ln w}_{f\ell} - \overline{\ln w}_{k\ell})^2 \tag{4}$$

While the counterfactual increase in within-task between-establishment within-industry wage inequality is slightly less pronounced than the actual increase – indicating that industries with above average within task-between establishment wage variances have grown in relative terms – there is a clear increase also in counterfactual inequality between 1990 and 2010, of close to 40% when we consider two broad task groups (Figure 2) and of more than 25% when we consider detailed occupation groups (Figure 3).[6]

To further gauge the relative importance of establishments versus tasks or occupations in wage determination, Panel A of Figure 4 displays the difference between the 90th and 10th and the 80th and 20th percentile in within-industry establishment wage premiums (broad tasks) over time alongside the abstract wage premium, where we average across industries using 1990 industry employment shares. Whereas the 80th-20th and 90th-10th gaps in firm wage premiums amounted to about 30 and 49 log points in 1990, they built up to 39 and 62 log points in 2010, an increase of 9 and 13 log points, respectively. For comparison, the (within-industry) abstract wage premium fluctuates around 30 log points over the same period. Panel B of Figure 4 reveals a similar picture when we focus on establishment wage premiums that hold the establishment's occupation structure (and not only its broad task structure) constant. Whereas the wage gap between the 10% lowest and 10% highest paying occupations increased by about 5 log points between 1990 and 20120, the wage gap between the 10%

[6] Appendix Figure A.2 shows the increase in within-industry between-establishment wage inequality separately for routine and abstract workers. The rise in inequality is noticeable, and of similar magnitude, for both task types.

lowest and highest paying establishments for workers in the same occupation increased by more than 10 log-points over the same period. Therefore, Figure 4 clearly illustrates that wages have become increasingly dependent on where workers work and (in relative terms) less dependent on what workers do.

Establishment Productivity and Task Usage. As a final piece of motivating evidence before outlining the theoretical framework, we explore the link between productivity and task usage at the establishment level. Using sales data from the IAB Establishment Panel, Column (1) of Table 1 explores the relationship between establishments' log productivity (total sales per full-time equivalent worker) and their abstract employment share. This is analyzed by running a regression which includes a set of fully interacted 3-digit industry and year fixed effects so that identification is limited to variation within industry-year cells. Observations are weighted by establishment size and survey weights, to make results representative for workers. Standard errors are clustered at the establishment level. The estimated coefficient shows that there is a positive and statistically significant relationship (at the 5% level) between productivity and abstract employment shares at the establishment level. Column (2) shows that there is a significant positive relationship between establishment size and productivity, such that establishment size can be used as a proxy for productivity. In Column (3) we confirm, using the much larger BEH data, that larger establishments employ a higher share of abstract workers. Hence, there is an empirical link between establishment productivity, size and task usage, which will guide our modelling choices in the following section.

4 Theoretical Framework

In this section we set up a theoretical framework that helps guide our analysis of the link between task-biased technological change and between-firm wage inequality. Motivated by the evidence in the previous section, we set up a model that allows for wage heterogeneity between firms, with a production structure that distinguishes between two tasks, and where firm productivity and task usage are linked. Specifically, we consider a version of the framework developed by Helpman et al. (2010) – a rich, yet tractable model of firm heterogeneity that allows for wage differences across firms within industries. Helpman et al. (2010) extend the Melitz (2003) model by introducing standard Diamond–Mortensen–Pissarides (Diamond, 1982a,b; Mortensen & Pissarides, 1994) search and matching frictions, as well as match-specific ability heterogeneity and a screening technology. We focus on the closed economy version of the extension of the model that allows for two types of labor inputs

(Section 5.2 of their paper) which, in our setting, we think of as two different tasks (abstract and routine).[7]

In Sections 4.1 and 4.2, we briefly outline the key components of the model and the equilibrium conditions, as derived by Helpman et al. (2010).[8] Our key innovation relative to the analysis in Helpman et al. (2010) is in Section 4.3, where we consider the implications of an aggregate routine-replacing technological change shock in the spirit of Autor et al. (2003) and Acemoglu & Autor (2011). We model this as an exogenous aggregate change in the factor-augmenting parameter associated with abstract workers, and study the implications for various workplace-level and industry-level outcomes.

4.1 Overview of the Helpman et al. (2010) Framework

Consumption

Within each sector, consumers demand a continuum of differentiated varieties. The aggregate consumption index is:

$$Q = \left[\int_{j \in J} q(j)^{\beta} dj \right]^{1/\beta}, \tag{5}$$

where j indexes varieties, J is the set of varieties within the sector, $q(j)$ denotes consumption of variety j, and $0 < \beta < 1$. The demand function for variety j is given by:

$$q(j) = A^{1/(1-\beta)} p(j)^{-1/(1-\beta)} \tag{6}$$

where A is a sectoral demand shifter and $p(j)$ is the price of variety j.

Production

As in Melitz (2003), there is a competitive fringe of potential firms that can choose to enter the market by paying an entry cost $f_e > 0$. Once a firm incurs the sunk entry cost, it observes its idiosyncratic value of θ, a parameter that is related to its productivity and its optimal production structure (as discussed below). θ is drawn from a Pareto distribution with scale parameter θ_{min} and shape parameter z, i.e. $G_{\theta}(\theta) = 1 - (\theta_{min}/\theta)^z$ for $\theta \geq \theta_{min} > 0$ and $z > 2$.[9] Once firms observe θ, they decide whether to exit or produce. Production involves

[7] Another framework that generates between-firm wage heterogeneity is the fair wage framework of Egger & Kreickemeier (2009, 2012). That setting requires assumptions about the way in which workers' fairness considerations relate to firm outcomes, which are somewhat ad-hoc.

[8] For full details, we refer the reader to the Helpman et al. (2010) paper.

[9] The assumption that $z > 2$ ensures that the variance of θ is finite.

a fixed cost of $f_d > 0$ units of the numeraire. Since in equilibrium all firms with the same value of θ behave symmetrically, firms can be indexed by θ.

Firms produce using a Constant Elasticity of Substitution (CES) technology with two types of labor inputs: abstract and routine workers (indexed by s and r, respectively). A firm's output depends on its value of θ, as well as its choice of how many workers of each type to hire (h_s and h_r), and the average match-specific ability of these workers ($\overline{a}_s$ and $\overline{a}_r$). Specifically, the production function is given by:

$$y = \left[(\theta\mu_s\overline{a}_s h_s^{\gamma})^{\nu} + (\mu_r\overline{a}_r h_r^{\gamma})^{\nu}\right]^{1/\nu} \tag{7}$$

where $0 < \nu < \beta$, and μ_s and μ_r are aggregate task-augmenting technology parameters.[10] For simplicity, we normalize $\mu_r = 1$. μ_s can therefore be interpreted in relative terms, as the relative aggregate task-bias of technology in favor of abstract tasks. The parameter θ enters into the production function as a firm-specific abstract task-augmenting parameter. Firms that draw higher values of θ will be more productive overall (absolute advantage), but productivity will be particularly high for their abstract workers (comparative advantage). Hence, θ reflects both productivity and the task-bias of production (in favor of abstract workers) of each firm. The model therefore incorporates the link between firm productivity and technological task bias that we documented empirically in Table 1. Given this link, we refer to the parameter θ interchangeably as both firm-specific technology and firm-specific productivity.

Search, Screening and Wage Bargaining

Labor markets are task-specific and there is a fixed aggregate supply of workers of each type (i.e. workers are not mobile across tasks). The firm must pay a search cost of b_ℓ in order to be matched with n_ℓ workers, $\ell = \{s, r\}$.[11] Consistent with the empirical evidence, we assume that abstract workers are relatively scare and hence command a higher search cost, i.e. $b_s > b_r$. Workers of a given task type are ex-ante identical but, upon matching with a firm, draw match-specific abilities from a Pareto distribution with shape parameter k and scale parameter a_{min}: $G_a(a) = 1 - (a_{min}/a)^k$; $a \geq a_{min} > 0$ and $k > 1$.[12] Ability is not observable by the firm or the worker, but a screening technology is available. By paying a screening cost $c\widetilde{a}_\ell^{\delta}/\delta$, firms are able to identify whether a worker's match-specific ability is

[10] The assumption that $\nu < \beta$ ensures that employment and wages of both types of workers are increasing in θ, as discussed below.

[11] b_ℓ is determined endogenously by labor market tightness and is proportional to workers' expected income outside the sector.

[12] This distribution is assumed to be common across both types of workers.

above or below an (endogenously chosen) cutoff $\widetilde{a}_\ell$, where $\ell = \{s, r\}$, $c > 0$, $\delta > 0$.

Wages are determined through Stole & Zwiebel (1996a,b) bargaining, under conditions of symmetric information. Since the screening technology only reveals whether a worker's match-specific ability is above or below $\widetilde{a}_\ell$, but not the specific ability of any individual worker, the expected ability of all hired workers of a given type is the same, and equal to $\overline{a}_\ell$, the expected value of a conditional on being above the threshold $\widetilde{a}_\ell$. Therefore, all workers of a given type within a given firm receive the same wage.

Summary of Firm and Worker Decisions

To summarize, firms first decide whether to produce or not; if they decide to produce, they choose how many routine and abstract workers to match with (n_r and n_s), and they choose the match-specific ability cutoffs that they will screen to ($\widetilde{a}_r$ and $\widetilde{a}_s$). Firms make an offer to all matched workers whose match-specific abilities are revealed to be above these endogenously chosen thresholds $\widetilde{a}_r$ and $\widetilde{a}_s$, and this determines their employment levels h_r and h_s.

Workers decide whether to work in the outside sector (where they can find a job with a certain wage for sure) or whether to search inside the sector under consideration (where the employment probability and wage are uncertain). In equilibrium, workers must be indifferent between the two. Helpman et al. (2010) discuss conditions under which expected worker income ω_ℓ in the outside sector is constant, and we make this assumption in what follows. Workers searching inside the sector are matched with a firm with some probability (which is proportional to labor market tightness in the sector) and, if their match-specific abilities turn out to be above the hiring thresholds set by the firm that they are matched to, they accept the job offer they receive (since this will always be preferable to remaining unemployed). Workers who have decided to search but remain unmatched, or who are matched but turn out to have match-specific abilities that fall below the hiring threshold of the firm that they are matched to, remain unemployed.

4.2 Key Equilibrium Properties

Closed-form solutions can be obtained for the equilibrium values of firm-level employment, wages, revenues, and profits for operating firms. Here we highlight the key properties of interest for our purposes. Derivation details of the key equilibrium relationships are given in Appendix C.1.

Productivity Threshold

As is standard in heterogeneous firm models, the presence of a fixed production cost implies that there is a zero-profit cutoff for productivity, θ_d, such that a firm that draws a productivity below this threshold exits without producing. Appendix C.1.2 shows how this productivity threshold can be pinned down using the Zero-Cutoff Profit condition, which requires the firm at the cutoff θ_d to make zero profits (Equation C.22), along with the Free Entry condition, which states that the expected profits for a potential entrant should equal the fixed entry cost (Equation C.23).

Firm-Level Employment

Firm-level employment for routine workers is given by:

$$h_r(\theta) = h_{dr}\left[1+\varphi(\theta)\right]^{\left(\frac{\beta\Lambda}{\nu\Gamma}-1\right)\left(1-\frac{k}{\delta}\right)} \tag{8}$$

where:

$$\varphi(\theta) = \mu_s^{\nu/\Lambda}\left(\frac{b_s}{b_r}\right)^{-\gamma\nu/\Lambda}\theta^{\nu/\Lambda} \tag{9}$$

Here, $\Lambda \equiv 1-\nu\gamma-\nu(1-\gamma k)/\delta > 0$, $\Gamma \equiv 1-\beta\gamma-\beta(1-\gamma k)/\delta > 0$, and $\Lambda > \Gamma$ due to the assumption that $\nu < \beta$. The derivation of this result is detailed in Appendix C.1.1. The definition of h_{dr}, which is a function of search costs b_r, screening costs c and other model parameters, is also provided in the Appendix (see Equation C.11).

Employment of abstract workers for a firm with productivity level θ is given by:

$$h_s(\theta) = \frac{b_r}{b_s}\varphi(\theta)^{1-k/\delta}h_r(\theta) \tag{10}$$

The firm's abstract share is therefore given by:

$$\frac{h_s(\theta)}{h(\theta)} = \frac{b_r\varphi(\theta)^{1-k/\delta}}{b_s+b_r\varphi(\theta)^{1-k/\delta}}, \tag{11}$$

where $h(\theta) = h_s(\theta) + h_r(\theta)$.

As shown in Appendix C.2, these equilibrium equations imply:

$$\frac{\partial h_r(\theta)}{\partial\theta} > 0, \qquad \frac{\partial h_s(\theta)}{\partial\theta} > 0, \qquad \frac{\partial h_s(\theta)/h(\theta)}{\partial\theta} > 0. \tag{12}$$

The model therefore predicts that more productive firms will employ a larger number of both abstract and routine workers, and, as a consequence, will be larger than less productive firms. More productive firms will also have a higher abstract employment share, implying that abstract workers disproportionately sort towards high-productivity firms – a prediction which is in line with the motivating evidence presented in Table 1.

Firm-Level Wages

Firm-level wages for routine workers are given by:

$$w_r(\theta) = w_{dr}\,[1+\varphi(\theta)]^{\left(\frac{\beta\Lambda}{\nu\Gamma}-1\right)\frac{k}{\delta}} \tag{13}$$

The derivation of this result is also detailed in Appendix C.1.1, with the definition of w_{dr} (which is also a function of search costs b_r, screening costs c and other model parameters) provided in Equation C.17.

Wages for abstract workers are given by:

$$w_s(\theta) = \frac{b_s}{b_r}\varphi(\theta)^{k/\delta} w_r(\theta) \tag{14}$$

We assume that the parameter conditions are such that $w_s(\theta) > w_r(\theta)$, which is consistent with the fact that we observe that abstract workers earn higher wages than routine workers in the data.

As shown in Appendix C.2, these equilibrium wage equations imply:

$$\frac{\partial w_r(\theta)}{\partial \theta} > 0, \qquad \frac{\partial w_s(\theta)}{\partial \theta} > 0 \tag{15}$$

The model therefore generates wage differences between firms, with more productive firms paying higher wages to both types of workers. Intuitively, this arises due to the complementarity between worker abilities and firm productivity, which gives an incentive for more productive firms to screen more intensively and choose a higher ability threshold. In equilibrium, wages are bargained down to the replacement cost of a worker, and given that more productive firms set higher hiring standards, their workers are costlier to replace and hence are paid a higher wage.

Note that both the match-specific heterogeneity and the screening technology are crucial elements (in addition to the search and matching frictions) in order to generate wage differences between firms for workers of a given task type. If workers (within task groups) were homogeneous, firms would have no incentive to screen, and wages would be bargained down to the replacement cost of a worker, which would simply be the search cost b_ℓ which is common across firms for workers of a given type. Firms would be heterogeneous along the size margin, but there would be no wage inequality between firms conditional on task.[13] On the other hand, if screening were not feasible, then the average expected ability of workers across all firms would be common and equal to the average match-specific ability in the population. In this case, the bargaining process would lead to a common wage across firms for all workers and once again there would be no wage inequality.

To summarize, the cross-sectional predictions of the model are that more productive firms are larger and have a higher abstract share, in line with the motivating empirical evidence presented in Table 1. These firms also pay higher wages, both because they hire a larger proportion of abstract workers (who earn higher average wages than routine workers), and because they pay higher wages to their workers (compared to less productive firms) conditional on task type.

4.3 Impacts of Task-Biased Technological Change

We model task-biased technological change (TBTC) as an exogenous increase in μ_s, the aggregate task-augmenting parameter for the abstract labor input in the production function in Equation (7). The literature on task-biased technological change (e.g. Autor et al., 2003; Acemoglu & Autor, 2011) argues that new automation technologies tend to replace labor in performing routine tasks, while complementing labor in abstract (or non-routine cognitive) tasks. Our modeling assumption captures the essence of this idea, by generating an exogenous increase in the relative demand for labor in abstract (relative to routine) tasks. Note that the shock to μ_s is an aggregate shock impacting all firms in the economy; however, as we show below, the impacts of this common shock are very heterogeneous across firms with different productivity levels. It should be noted that we focus here on the *relative* effects of TBTC – that is, the effects of TBTC on wages and employment of routine vs abstract workers in low vs high productivity firms – rather than the *absolute* effects of TBTC on overall wage and employment *levels*.

The key implications of an increase in μ_s are the following:[14]

[13] See for instance Felbermayr et al. (2011).

[14] In what follows, we assume that the search costs b_s and b_r are not affected by technological change. The search costs are proportional to workers' expected income outside the sector (outside option). Helpman et al. (2010) discuss conditions under which the outside options can be assumed to be constant, even when there are shocks with aggregate implications (such as trade opening, in the setting analyzed in their paper).

Prediction 1: *Selection* – Task-biased technological change increases the productivity threshold for production θ_d.

Proof: *See Appendix C.3.*

Implications: By increasing the productivity threshold θ_d, TBTC leads to the exit of firms at the bottom of the productivity distribution. Although this reduces the support of the distribution among operating firms, the variance of productivity among these firms increases. This is due to the fact that the distribution of productivity among operating firms is a truncated Pareto distribution with scale parameter θ_d and shape parameter z, and the variance of this distribution is increasing in the scale parameter θ_d.[15] Intuitively, with a Pareto distribution, the increase in θ_d entails the exit of a mass of relatively homogeneous unproductive firms. This mass is shifted towards the tail of the distribution, leading to an increase in the variance of productivity among the firms that remain in operation. Given the increase in the variance of productivity, wage inequality among operating firms increases even absent any within-firm changes in employment or wages. Moreover, given that the abstract employment share is increasing in θ, the firms that exit are relatively routine-intensive, and hence their exit contributes to the rise in the economy's overall abstract share.

Prediction 2: *Differential Employment Growth* – TBTC strengthens the cross-sectional association between employment and productivity.

Proof: *As shown in Appendix C.3:*

$$\frac{\partial\left(\frac{\partial h_r(\theta)}{\partial\theta}\right)}{\partial\mu_s} > 0 \qquad \text{and} \qquad \frac{\partial\left(\frac{\partial h_s(\theta)}{\partial\theta}\right)}{\partial\mu_s} > 0. \tag{16}$$

Implications: This prediction implies that more productive firms become disproportionately larger in terms of employment relative to less productive firms. This shift in employment of both types of workers towards more productive firms (which pay higher wages) leads to an increase in (worker-weighted) between-firm wage inequality (by task and overall). Moreover, given that more productive firms have a higher abstract employment share, this shift also contributes to the rise in the economy's overall abstract share.

[15] The variance is given by $\frac{z\theta_d^2}{(z-1)^2(z-2)}$. It should be noted that an increase in the productivity threshold θ_d will not increase the variance of productivity among operating firms for all distributions of firm productivity. For example, if firm productivity were uniformly distributed, an increase in the productivity threshold θ_d would lower the variance of firm productivity among operating firms.

Prediction 3: *Increased Sorting* – TBTC increases the abstract employment share within all firms, and strengthens the cross-sectional association between productivity and abstract employment shares.

Proof: *As shown in Appendix C.3:*

$$\frac{\partial\left[h_s(\theta)/h(\theta)\right]}{\partial\mu_s} > 0, \qquad \text{and if } \frac{h_s(\theta)}{h_r(\theta)} < 1, \text{ then } \frac{\partial\left(\frac{\partial h_s(\theta)/h(\theta)}{\partial\theta}\right)}{\partial\mu_s} > 0. \tag{17}$$

Implications: This prediction implies that within-firm changes in task composition also contribute to the increase in the aggregate share of abstract employment (in addition to the two channels highlighted in Predictions 1 and 2). Moreover, as long as firms employ relatively more routine than abstract workers at baseline (which is the empirically relevant case),[16] more productive firms will increase their abstract employment share by more than less productive firms. In consequence, firms within industries will become more heterogeneous in their task input mix as a result of TBTC. The increased sorting of abstract workers towards high-wage firms will also contribute to the increase in between-firm wage inequality.

Prediction 4: *Differential Wage Growth* – TBTC strengthens the cross-sectional association between productivity and wages.

Proof: *As shown in Appendix C.3:*

$$\frac{\partial\left(\frac{\partial w_r(\theta)}{\partial\theta}\right)}{\partial\mu_s} > 0 \qquad \text{and} \qquad \frac{\partial\left(\frac{\partial w_s(\theta)}{\partial\theta}\right)}{\partial\mu_s} > 0. \tag{18}$$

Implications: As a result of TBTC, wages for both types of workers disproportionately increase within more productive firms (relative to less productive firms). This leads to a further increase in wage inequality (overall and by task) across firms.

To summarize, the model unambiguously predicts that task-biased technological change leads to an increase in between-firm wage inequality. This operates through various distinct channels: selection, differential employment growth, sorting, and differential within-firm

[16]More than 80% of workers in our sample are in workplaces that employ more routine than abstract workers.

wage growth – all of which compound each other. Intuitively, aggregate task-biased technological change exacerbates the comparative advantage of firms that are more productive in abstract tasks, which are initially the more productive firms overall. This leads to an increase in their relative size and relative wages. The model also predicts that selection, differential employment growth and sorting all contribute to the rise in the overall share of abstract employment.

5 Empirical Evidence

5.1 Cross-Sectional Relationships

We begin by verifying the empirical relevance of the cross-sectional equilibrium predictions discussed in Section 4.2. The model predicts that more productive firms are larger and employ a higher share of abstract workers. These predictions were verified in Table 1 and motivated our model setup. The model also predicts that more productive firms pay higher average wages, overall and conditional on worker task. We verify these predictions in Table 2. As in Table 1, all regressions include fully interacted 3-digit industry and year fixed effects, and hence exploit only variation across establishments within industries at a given point in time. Observations are weighted by establishment size (and in the case of the IABEP, the survey weights are also used in order to make results representative for workers) and standard errors are clustered at the establishment level. The top panel draws on data from the IABEP linked to BEH social security records, and uses log productivity (revenue per worker) as the key regressor of interest. The bottom panel draws on the full BEH records, and uses log establishment size (which as shown in Table 1 is positively and significantly related to establishment productivity) as the key regressor of interest.

Column (1) in Panel A shows that more productive establishments pay, on average, higher wages. This is partly driven by the higher proportion of abstract workers at these establishments. In Column (2), we regress the establishment's wage premium (two tasks), computed as the establishment's average residual from a year-specific individual log-wage regression that conditions on an abstract indicator interacted with 3-digit industry fixed effects, on (log) establishment productivity. While the coefficient is positive – indicating that more productive establishments pay higher wages conditional on worker task – it is smaller in magnitude than in Column (1), in line with more productive establishments employing more abstract workers. Column (3) confirms that more productive establishments pay higher wages also to workers within the same detailed occupation group; hence, more productive establishments pay higher wages not only because they employ a higher share of workers in higher paying occupations. Columns (4) and (5) further show that more productive establishments employ more workers of both routine and abstract task type, as predicted by the model.

Panel B further corroborates these relationships using establishment size as a key regressor of interest. Larger establishments pay higher wages on average not only overall (Column (1)), but also to workers of the same task type (Column (2)) and to workers within the same detailed occupation (Column (3)).

Note that these relationships occur within 3-digit industries and are thus not accounted for by differences across industries in establishment size, wages and productivity.

5.2 Associations over Time and Longitudinal Changes within Establishments

As discussed in Section 4.3, the model predicts that the cross-sectional associations between firm productivity and firm size, abstract share, and wage should become stronger over time if there is ongoing task-biased technological change – all of which contributes to the rise in wage inequality between firms and the decline in the abstract share within industries. To test whether these relationships have indeed become stronger, we estimate the associations from Tables 1 and 2 separately for each year, controlling for 3-digit industry fixed effects, thereby focusing once again on within-industry associations. Figure 5 plots the coefficients from these yearly regressions, using data from the IABEP and (log) productivity as a key regressor. All of the associations have indeed become substantially stronger over our sample period. For example, as shown in Panel A, while in the early 1990s a 1% increase in the establishment's productivity was associated with an increase in establishment size of about 0.1%, the association had increased to more than 0.4% by 2010. Similarly, as shown in Panel C, the coefficient from the regression of average establishment log wages on log labor productivity tripled from about 0.05 in the early 1990s to 0.15 by 2010. The association between (log) establishment productivity and the establishment wage premium, accounting for the broad task or detailed occupational composition of the establishment, likewise nearly tripled over the time period (Panels D and E).

Figure 6 confirms these findings drawing on the larger BEH social security data and using establishment size as the key regressor of interest. The figure shows that the associations between establishment size and overall wages, as well as between establishment size and establishment wage premiums (though not between establishment size and the abstract share) increased over the sample period. For example, as Panel D shows, while a 1% increase in establishment size was associated with a 0.06% increase in the establishment wage premium (two tasks) in the early 1990s, this association steadily rose to nearly 0.09% by 2010.

Table 3 complements this evidence and shows estimates based on a set of regressions that consider within-establishment changes over (non-overlapping) 5-year windows and their relationships with the establishment's size at baseline, drawing on social security records from the BEH. Establishments that are larger at baseline exhibit a larger increase in labor productivity, a larger increase in the employment share of abstract workers, higher wage growth overall, and a larger increase in their wage premiums, conditional on two broad tasks or on detailed occupation groups. In line with the predictions of the model, these results show that establishments that perform "better" at baseline (in terms of their size) pull away even further from other establishments (in terms of their abstract share and the wages they pay) in their industry. TBTC therefore appears to amplify differences in productivity, task usage and pay across establishments within industries.

5.3 Abstract Share Heterogeneity and Sorting

Another implication of the model discussed in Section 4.3 is that task-biased technological change leads to increased firm heterogeneity within industries in terms of abstract employment shares, with increased sorting of abstract workers towards high-wage firms.

We verify these predictions in Figure 7. Panel A plots the within-industry variance of establishments' abstract employment shares, averaged across industries using either the contemporaneous or the 1990 industry structure (see Equations B.3 and B.4 in Appendix B.2). The figure shows a clear increase in the variance of abstract employment shares across establishments within industries. Thus, rather than having converged towards a more uniform mode of production, establishments have become increasingly heterogeneous in terms of the task input mix that they use, in line with the prediction of the model.

Panel B of Figure 7 shows the within-industry co-variance between establishments' abstract employment shares and their wage premiums (two task groups), averaged once again across industries either using the contemporaneous or the 1990 industry structure (see equations B.5 and B.6 in Appendix B.2). This co-variance also shows a clear positive trend over time: abstract workers increasingly sort into establishments that pay higher wage premiums. This evidence, which our model rationalizes as being driven by task-biased technological change, is consistent with the patterns documented by Card et al. (2013) and Song et al. (2019), which show that high-wage workers increasingly sort into high-wage firms and that high-wage workers are increasingly likely to work with each other.

5.4 Decompositions of Changes in Wage Inequality

The model highlights that TBTC leads to an increased selection of firms that operate in the industry (Prediction 1); to larger employment growth of more productive firms within the industry (Prediction 2); to increased sorting of abstract workers towards more productive firms within the industry (Prediction 3) and to larger wage growth for each worker type in more productive firms within the industry (Prediction 4). Empirical evidence consistent with these predictions has been provided in Figures 5-7 and Table 3. All four channels contribute to the rise in within-industry between-establishment wage inequality documented in Section 3. In this section, we perform various decompositions in order to quantify the relative empirical importance of the different channels.

We begin by assessing the role of changes in the composition of operating establishments (i.e. selection). The model predicts that the composition of operating firms will become more positively selected, and this will lead to an increase in the variance of productivity among operating firms (as firm productivity is drawn from a Pareto distribution), and therefore to more wage inequality.

To quantify the importance of this selection channel, we first classify establishments as continuing, exiting or entering according to their status in periods $t-5$ and t (corresponding to non-overlapping 5-year windows in our data). We can then decompose the change in between-establishment wage inequality in industry k between $t-5$ and t into a "selection effect" due to selective establishment entry and exit and a change in the variance in establishment wages among continuing establishments as follows:

$$\Delta Var_{kt}(\ln w) = \underbrace{Var_{kt} - Var_{kt}^{con} + Var_{kt-5}^{con} - Var_{kt-5}}_{\text{selection}} + \underbrace{\Delta Var_{kt}^{con}}_{\text{continuing establishments}} \tag{19}$$

This decomposition yields two components. The first component captures changes in between-establishment wage inequality due to selective establishment entry and exit. Specifically, the model predicts that in the base period $t-5$ the within-industry variance of establishment wages among continuing establishments (Var_{kt-5}^{con}) exceeds the variance among all establishments (Var_{kt-5}). The second component captures changes in between-establishment wage inequality among continuing establishments.

We compute the two components for each industry over each 5-year windows in our data and average across industries using 1990 employment shares as weights. Panel A of Figure 8 presents the time-averaged value of each component (giving equal weight to all four time intervals). The figure highlights that while selective entry and exit have contributed to the rise in within-industry between-establishment wage inequality over time, this effect is small in magnitude, accounting for only 6% of the total increase. The bulk of the rise in between-establishment wage inequality is therefore driven by continuing establishments.

Next, we further decompose the change in the within-industry between-establishment variance among continuing establishments, ΔVar_{kt}^{con}, to determine the role of worker sorting, increased heterogeneity across establishments in terms of their task usage, and wage premiums conditional on task. The starting point of this decomposition is an individual level regression of log-wages on task indicators (routine vs abstract) interacted with industry fixed effects, estimated separately for each year. The residual from this worker level regression, averaged to the level of the establishment, corresponds to the establishment wage premium (two tasks), which we have analyzed above and denote by $\widetilde{\ln w_f}$ below. Similarly, we can compute the predicted wage from this worker level regression and average the individual-level predicted wages to the level of the establishment; we denote this by $\widehat{\ln w_f}$ below.

We can then decompose the change in the variance among continuing establishments as follows:

$$\Delta Var_{kt}^{con}(\widetilde{\ln w_f} + \widehat{\ln w_f}) = \underbrace{\Delta Var_{kt}^{con}(\widetilde{\ln w_f})}_{\substack{\text{change in variance}\\ \text{estab wage premium}}} + \underbrace{\Delta Var_{kt}^{con}(\widehat{\ln w_f})}_{\substack{\text{change in var predicted}\\ \text{wage given tasks}}} + \underbrace{2\Delta Cov_{kt}^{con}(\widetilde{\ln w_f}, \widehat{\ln w_f})}_{\text{change in co-variance}}. \tag{20}$$

The first component, $\Delta Var_{kt}^{con}(\widetilde{\ln w_f})$, captures changes in the variance of the premiums paid by different establishments within an industry, conditional on their task composition. The second component, $\Delta Var_{kt}^{con}(\widehat{\ln w_f})$, captures changes in the industry-specific abstract premium over time, in addition to changes in heterogeneity across establishments in their task usage. To see this, note that the predicted wage of establishment f in industry k in year t is equal to the average wage of routine workers in the industry and year, $\bar{w}_{kt}^R$, plus the share of abstract workers in the establishment multiplied by the industry-year-specific abstract wage premium, $S_{f(k)t} \cdot \text{AbPrem}_{kt}$. The within-industry variance in predicted establishment wages among continuing establishments thus equals:

$$\begin{aligned} Var_{kt}^{con}(\widehat{\ln w_f}) &= Var_{kt}^{con}(\bar{w}_{kt}^R + S_{f(k)t} \cdot \text{AbPrem}_{kt}) \\ &= \text{AbPrem}_{kt}^2 \cdot Var_{kt}^{con}(S_{f(k)t}) \end{aligned}$$

Hence, changes in the within-industry variance of establishments' predicted wages are driven by both changes in the industry-specific abstract wage premium and by changes in the variance of the abstract employment share in the establishment. As shown in Appendix Figure A.1, the abstract wage premium remained roughly constant over time. Panel A of Figure 7, meanwhile, shows that establishments have become increasingly heterogeneous in terms of the employment share of abstract workers. We therefore expect any change in the within-industry variance of predicted establishment wages to be primarily driven by changes in the variance of establishments' abstract employment shares.

In turn, the third component in Equation (20), $2\ \Delta Cov^{con}_{kt}(\widetilde{\ln w_f}, \widehat{\ln w_f})$, captures the increased sorting of abstract workers into firms paying higher establishment wage premiums. To see this, note that the co-variance between establishments' predicted wages and their wage premiums equals:

$$\begin{aligned} Cov^{con}_{kt}(\widehat{\ln w_f}, \widetilde{\ln w_f}) &= Cov^{con}_{kt}(\bar{w}^R_{kt} + S_{f(k)t} \cdot \text{AbPrem}_{kt}, \widetilde{\ln w_f}) \\ &= \text{AbPrem}_{kt} \cdot Cov^{con}_{kt}(S_{f(k)t}, \widetilde{\ln w_f})) \end{aligned}$$

Since the abstract wage premium remained roughly constant over time, the co-variance primarily captures the increased sorting of abstract workers into establishments that pay higher wage premiums, as documented in Panel B of Figure 7.

The results of this decomposition are presented in Panel B of Figure 8, where we once again average across industries using the 1990 industry structure as weights and give equal weight to each 5-year period. Even though we distinguish between two task groups only, increased sorting of abstract workers to establishments paying higher wage premiums (the third component in Equation 20) can account for about 12% of the overall increase in within-industry between-establishment wage inequality among continuing establishments. While dispersion in task usage has increased over time across establishments within the same industry (see Panel A of Figure 7), its contribution to the overall increase in between-establishment wage inequality is minor (the second component in Equation 20). Not surprisingly, given that we only distinguish between two task groups, changes in establishment wage premiums account for the majority (86.1%) of the change in the variance of wages among continuing establishments.

In Panel A of Appendix Figure A.3, we repeat the exercise distinguishing between 317 occupations, rather than two tasks. As expected, increased dispersion in the occupational structure across establishments (the second component in Equation 20) and increased sorting of workers in high-paying occupations into establishments paying high establishment premiums (the third component in Equation 20) become quantitatively more important, accounting for about 20% and 25% of the overall increase in the within-industry wage variance among continuing establishments.

In a final step, our goal is to gauge the importance of differential employment growth for the increase in between-establishment wage inequality. The model highlights that TBTC leads to larger employment growth of more productive firms (Prediction 2). In line with this prediction, we have documented that the association between establishment productivity and size has become stronger over time (Panel A of Figures 5 and 6). If establishments that were initially more productive and hence paid higher wages grow at a faster rate, then this will result in an increase in wage inequality across establishments, even if the wage premiums that establishments pay remain unchanged.

To assess the importance of this channel, we focus here on the change in the variance of the establishment wage premiums, $\Delta Var^{con}_{kt}(\widetilde{\ln w})$, and decompose this into the role that is due to differential employment growth, and a residual component that is attributable to differential wage growth within establishments. This final decomposition is given by:

$$\Delta Var^{con}_{kt}(\widetilde{\ln w}) = \underbrace{\sum_{f \in con} \left(\omega^{con}_{f(k)t} - \omega^{con}_{f(k)t-5}\right)\left(\widetilde{\ln w}_{f(k)t-5} - \overline{\widetilde{\ln w}}_{kt-5}\right)^2}_{\text{differential employment growth}} + \underbrace{\sum_{f \in con} \omega^{con}_{f(k)t}\left[\left(\widetilde{\ln w}_{f(k)t} - \overline{\widetilde{\ln w}}_{kt}\right)^2 - \left(\widetilde{\ln w}_{f(k)t-5} - \overline{\widetilde{\ln w}}_{kt-5}\right)^2\right]}_{\text{residual (differential wage growth)}}. \quad (21)$$

Here, $\omega^{con}_{f(k)t}$ is the employment share of establishment f among continuing establishments of the same industry; its wage premium is $\widetilde{\ln w}_{f(k)t}$, and the average wage premium in industry k in period t among continuing establishments is $\overline{\widetilde{\ln w}}_{kt}$. The first component of the decomposition therefore shows the changes in the variance that arise solely from changes in establishment size, holding establishment wage premiums at $t-5$ levels.

Panel C of Figure 8 shows the results from this decomposition. The figure shows that differential employment growth accounts for around half of the increase in the variance of log wage premiums among continuing establishments. The residual component, however, is also important, indicating that differential changes in wage premiums within establishments are also a major driver of the change in wage inequality among continuing establishments.[17]

Overall, we can conclude that all four channels highlighted by the model have contributed to the rise in wage inequality between establishments within industries. Changes in the composition of operating establishments, sorting of abstract workers to establishments pay-

[17] Panel B of Appendix Figure A.3 presents the analogous results based on the establishment wage premiums that control for detailed occupations, rather than two tasks.

ing higher wages to the same worker type, differential employment growth, and differential within-establishment wage growth are all quantitatively important, with the latter two being of primary importance.

5.5 Decompositions of Changes in the Abstract Employment Share

We can perform a similar decomposition analysis to assess the quantitative importance of the different channels in accounting for the rise in the aggregate share of abstract employment within industries. First, the within-industry abstract employment share may increase because of selective establishment entry and exit. Specifically, the model highlights that TBTC drives the least productive firms out of the market. As these firms employ relatively fewer abstract workers, selective firm exit will lead to an increase in the abstract share in the industry.

To determine the importance of this channel, we first decompose the within-industry change in the abstract share, $\sum_k \omega_{k1990} \Delta S_{kt}$, into changes among continuing establishments and changes due to entering and exiting establishments as follows:

$$\begin{aligned}\sum_k \omega_{k1990} \Delta S_{kt} \quad = \quad & \underbrace{\sum_k \omega_{k1990} \Delta S_{kt}^{con}}_{\text{Change continuing establishments}} \\ & \underbrace{+ \sum_k \omega_{k1990} \left[\frac{E_{kt}^{exit}}{E_{kt}} (S_{kt-5}^{con} - S_{kt-5}^{exit}) + \frac{E_{kt}^{entry}}{E_{kt}} (S_{kt}^{entry} - S_{kt}^{con}) \right]}_{\text{Change establishment entry and exit}},\end{aligned} \tag{22}$$

where $\frac{E_{kt}^{I}}{E_{kt}}$ is the employment-weighted share of establishments in the respective industry and year that are either entering or exiting between $t-5$ and t.

Panel A of Figure 9 shows the results from this decomposition, where we once again consider four 5-year windows and average across time periods giving each period an equal weight. While selective establishment entry and exit contribute to the rise in the within-industry abstract share, the effect is quantitatively small, accounting for 11.3% of the overall change. The rise in the abstract employment share within industries is therefore primarily due to changes among continuing establishments.

In a next step, we follow an approach analogous to Autor et al. (2020) and decompose the within-industry rise in the abstract employment share among continuing establishments into a within-establishment and a reallocation component as follows:

$$\Delta S_{kt}^{con} = \underbrace{\sum_{f\in con} \Delta S_{f(k)t}}_{\text{within establishments}} \underbrace{+\Delta \sum_{f\in con} (\omega_{f(k)t} - \bar{\omega}_{kt}^{con})(S_{f(k)t} - S_{kt}^{con})}_{\text{reallocation}} \tag{23}$$

where $\omega_{f(k)t}$ and $S_{f(k)t}$ denote the relative size and abstract employment share of establishment f at time t, while $\bar{\omega}_{kt}^{con}$ and S_{kt}^{con} denote the industry averages among continuing establishments. The first component in Equation (23) captures changes in the abstract employment share within continuing establishments, where each establishment is given the same weight, irrespective of their size. The second component corresponds to the change in the co-variance between establishment size and the establishment's abstract share between the two time periods and thus captures changes over time in the allocation of abstract workers to establishments of different sizes.

Panel B of Figure 9 shows the results from this decomposition. The figure highlights that less than a quarter of the within-industry rise in the abstract share occurred within establishments. The reallocation of abstract workers to establishments of different sizes is therefore the primary channel through which abstract employment shares adjust within an industry.

Yet, Panel C of Figure 9 suggests that worker reallocation is in large part driven by larger establishments experiencing a larger increase in their abstract employment share compared to smaller establishments in their industry (as documented in Table 3). The panel contrasts the unweighted within-establishment change in the abstract share where all establishments in the industry are given the same weight (i.e., $\sum_{f\in con} \Delta S_{f(k)t}$) with the weighted within-establishment change where establishments that were larger at baseline are given a larger weight (i.e., $\sum_{f\in con} \omega_{f(k)t-5} \Delta S_{f(k)t}$). The weighted within-establishment increase in the abstract share is more than five times larger than the unweighted increase – implying that larger establishments increased their abstract employment share by more than smaller establishments.

5.6 Technology Adoption: Industry-Level Analysis

As a final support of the implications of the model, we exploit variation at the industry level and make use of more direct measures of technology adoption. Specifically, we analyze whether industries with more technology adoption have experienced larger increases in between-establishment wage inequality, in the variance of abstract employment shares across establishments, and in the sorting of abstract workers to high-wage establishments.

We first consider industry-level variation in the overall change in the employment share of abstract workers between 1990 and 2010. Since our focus is on task-biased technological change, we can think of industries that experience larger increases in their abstract employment shares as being more exposed to this type of technological innovation. For simplicity, we divide industries into two groups, based on whether they experience above-median or below-median increases in the abstract employment share over the entire period.[18]

Panel A of Figure 10 shows the evolution over time of the variance of establishment wages for these two groups of industries. In line with our prediction, we find that industries that experience larger increases in their abstract employment share also experience a stronger increase in wage inequality between establishments. As shown in Panel B, these industries also show a stronger increase in the between-establishment variance of the employment share of abstract workers. Hence, abstract-intensive establishments increase their abstract employment share by more than routine-intensive establishments particularly in industries experiencing a larger overall increase in abstract employment. Panels C and D further highlight that the variance in establishment wage premiums, adjusting for the task and occupation composition in the establishment, increased more in industries that experienced a larger overall increase in the abstract employment share. Finally, Panel E provides evidence of increased sorting of abstract workers to high wage premium establishments in industries characterized by larger increases in their abstract employment shares.

Figure 11 exploits a more direct measure of technology exposure based on the change in robots per worker within industries over the 1993-2010 time period, using data from the International Federation of Robotics. Once again we divide industries into two groups, according to whether they experience above or below median changes during this time period.

Panel A first confirms that we can think of robot adoption as task-biased technological change: Industries with above-median robot adoption experience a much larger increase in their abstract employment shares. Panels B to D further show that these industries also experience larger increases in the variance of average establishment wages and establishment wage premiums that adjust for the task and occupation structure in the establishment. Importantly, not only are establishments becoming increasingly heterogeneous in terms of their task mix particularly in industries with above-median robot adoption (Panel E), but the sorting of abstract workers into high-wage establishments is also particularly pronounced in these industries (Panel F).

Finally, Figure 12 shows that we obtain broadly consistent results if we use a measure of technology adoption based on the industry's change in ICT capital stock per worker between 1991 and 2007 from the EUKLEMS data. Industries with more technology adoption tend to experience larger increases in their abstract share (Panel A), larger increases in between-establishment inequality in average establishment wages and wage premiums (Panels B to

[18]The median is computed based on the employment distribution across industries in 1990.

D), larger increases in the dispersion of establishments' task usage (Panel E), and more sorting of abstract workers towards high wage establishments.

6 Conclusions

In this paper, we show that task-biased technological change is an important driver of the rise in between-establishment wage inequality in Germany between 1990 and 2010. While a large literature has considered the role of task-biased technological change for wage inequality, it has focused on its implications in terms of wage differentials *between* workers in different task groups. Empirically, however, the increase in wage inequality is primarily driven by increased wage differentials *within* task groups, across establishments. By embedding a task-biased technological change shock within a rich, yet tractable heterogeneous firm framework, we show that this type of shock will lead to heterogeneous responses at the firm level that will generate a rise in between-workplace wage inequality along the lines observed in the data. The rise in inequality occurs due to endogenous changes in establishment composition, worker sorting, establishment size and establishment wages paid to the same worker type. Using rich administrative social security data from Germany, we provide evidence consistent with the model and quantify the importance of the different channels that it highlights. We also provide evidence that the key workplace-level patterns that we identify as being driven by the technological change shock are indeed more pronounced within industries that have experienced stronger rates of technology adoption.

Our results highlight the importance of moving beyond the traditional representative firm setting with competitive labor markets when considering the impact of aggregate shocks such as technological change. While the literature has generally thought about the individual-level impacts of task-biased technological change as being related to the tasks that individuals perform, our findings indicate that the type of firm that individuals are matched to is also crucial: Routine workers employed in low-productivity firms lose out not only relative to abstract workers in these firms, but also relative to routine workers in high-productivity ones.

The literature on technological change has long thought about increases in educational attainment as being a useful tool in order to offset the rise in inequality (Tinbergen, 1974, 1975; Acemoglu & Autor, 2011). In a homogeneous firm setting with competitive labor markets, a rise in the supply of abstract workers will offset the rise in the abstract wage premium and thus the rise in (between-task) inequality that arises due to task-biased technological change. When moving away from this competitive, homogeneous firm framework, however, the idea that there is a simple 'race' between technology and the supply of skills becomes less clear. Even if an increase in educational attainment could dampen the rise in the abstract wage premium (as is indeed observed in the German case over our sample period), this may not be enough in order to avoid a rise in wage inequality across firms among workers in the same type of occupation. Understanding what type of policies can mitigate the increase in inequality in this more realistic environment with heterogeneous firms and various market frictions remains an important avenue for future work.

References

Acemoglu, D. & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4, 1043–1171.

Acemoglu, D. & Restrepo, P. (2020a). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.

Acemoglu, D. & Restrepo, P. (2020b). Unpacking skill bias: Automation and new tasks. *Working Paper*.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2017). Concentrating on the fall of the labor share. *American Economic Review: Papers & Proceedings*, 107(5), 180–185.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2), 645–709.

Autor, D. H., Katz, L. F., & Kearney, M. S. (2006). The polarization of the US labor market. *The American Economic Review*, 96(2), 189–194.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.

Azar, J., Marinescu, I., Steinbaum, M., & Taska, B. (2020a). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, 66, 101886.

Azar, J., Marinescu, I. E., & Steinbaum, M. (2020b). Labor market concentration. *Journal of Human Resources*, Forthcoming.

Bajgar, M., Berlingieri, G., Calligaris, S., Criscuolo, C., & Timmis, J. (2019). Industry concentration in Europe and North America. *OECD Productivity Working Papers, No. 18*.

Barth, E., Bryson, A., Davis, J. C., & Freeman, R. (2016). It's where you work: Increases in the dispersion of earnings across establishments and individuals in the United States. *Journal of Labor Economics*, 34(S2), S67–S97.

Card, D., Heining, J., & Kline, P. (2013). Workplace heterogeneity and the rise of West German wage inequality. *The Quarterly Journal of Economics*, 128(3), 967–1015.

Cortes, G. M. & Salvatori, A. (2019). Delving into the demand side: Changes in workplace specialization and job polarization. *Labour Economics*, 57, 164–176.

Cortes, G. M. & Tschopp, J. (2020). Rising concentration and wage inequality. *IZA Discussion Paper No. 13557.*

Criscuolo, C., Hijzen, A., Schwellnus, C., Barth, E., Chen, W.-H., Fabling, R., Fialho, P., Grabska, K., Kambayashi, R., Leidecker, T., Skans, O. N., Riom, C., Roth, D., Stadler, B., Upward, R., & Zwysen, W. (2020). Workforce composition, productivity and pay: The role of firms in wage inequality. *OECD Social, Employment and Migration Working Papers No. 241.*

Diamond, P. A. (1982a). Aggregate demand management in search equilibrium. *Journal of Political Economy*, 90(5), 881–894.

Diamond, P. A. (1982b). Wage determination and efficiency in search equilibrium. *The Review of Economic Studies*, 49(2), 217–227.

Dustmann, C., Fitzenberger, B., Schönberg, U., & Spitz-Oener, A. (2014). From sick man of Europe to economic superstar: Germany's resurgent economy. *Journal of Economic Perspectives*, 28(1), 167–88.

Dustmann, C., Ludsteck, J., & Schönberg, U. (2009). Revisiting the German wage structure. *Quarterly Journal of Economics*, 124(2), 843–881.

Egger, H. & Kreickemeier, U. (2009). Firm heterogeneity and the labor market effects of trade liberalization. *International Economic Review*, 50(1), 187–216.

Egger, H. & Kreickemeier, U. (2012). Fairness, trade, and inequality. *Journal of International Economics*, 86(2), 184–196.

Felbermayr, G., Prat, J., & Schmerer, H.-J. (2011). Globalization and labor market outcomes: Wage bargaining, search frictions, and firm heterogeneity. *Journal of Economic Theory*, 146(1), 39–73.

Goldschmidt, D. & Schmieder, J. F. (2017). The rise of domestic outsourcing and the evolution of the German wage structure. *The Quarterly Journal of Economics*, 132(3), 1165–1217.

Goos, M., Manning, A., & Salomons, A. (2009). Job polarization in Europe. *American Economic Review*, 99(2), 58–63.

Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8), 2509–2526.

Graetz, G. & Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, 100(5), 753–768.

Helpman, E., Itskhoki, O., Muendler, M.-A., & Redding, S. J. (2017). Trade and inequality: From theory to estimation. *Review of Economic Studies*, 84(1), 357–405.

Helpman, E., Itskhoki, O., & Redding, S. (2010). Inequality and unemployment in a global economy. *Econometrica*, 78(4), 1239–1283.

Jaimovich, N. & Siu, H. E. (2020). Job polarization and jobless recoveries. *The Review of Economics and Statistics*, 102(1), 129–147.

Katz, L. F. & Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1), 35.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.

Michaels, G., Natraj, A., & Van Reenen, J. (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics*, 96(1), 60–77.

Mortensen, D. T. & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3), 397–415.

Mueller, H. M., Ouimet, P. P., & Simintzi, E. (2017). Wage inequality and firm growth. *American Economic Review: Papers & Proceedings*, 107(5), 379–83.

Rinz, K. (2020). Labor market concentration, earnings, and inequality. *Journal of Human Resources*, Forthcoming.

Song, J., Price, D. J., Guvenen, F., Bloom, N., & von Wachter, T. (2019). Firming up inequality. *The Quarterly Journal of Economics*, 134(1), 1–50.

Stole, L. A. & Zwiebel, J. (1996a). Intra-firm bargaining under non-binding contracts. *The Review of Economic Studies*, 63(3), 375–410.

Stole, L. A. & Zwiebel, J. (1996b). Organizational design and technology choice under intrafirm bargaining. *The American Economic Review*, (pp. 195–222).

Tinbergen, J. (1974). Substitution of graduate by other labor. *Kyklos*, 27(2), 217–226.

Tinbergen, J. (1975). *Income Differences: Recent Research*. North-Holland Publishing Company.

Trottner, F. (2019). Who gains from scale? Trade and wage inequality within and between firms. *Working Paper*.

Webber, D. A. (2015). Firm market power and the earnings distribution. *Labour Economics*, 35, 123 – 134.

Wilmers, N. & Aeppli, C. (2021). Consolidated advantage: New organizational dynamics of wage inequality. *Washington Center for Equitable Growth Working Paper*.

Yeaple, S. R. (2005). A simple model of firm heterogeneity, international trade, and wages. *Journal of international Economics*, 65(1), 1–20.

Figure 1: Abstract Employment Share

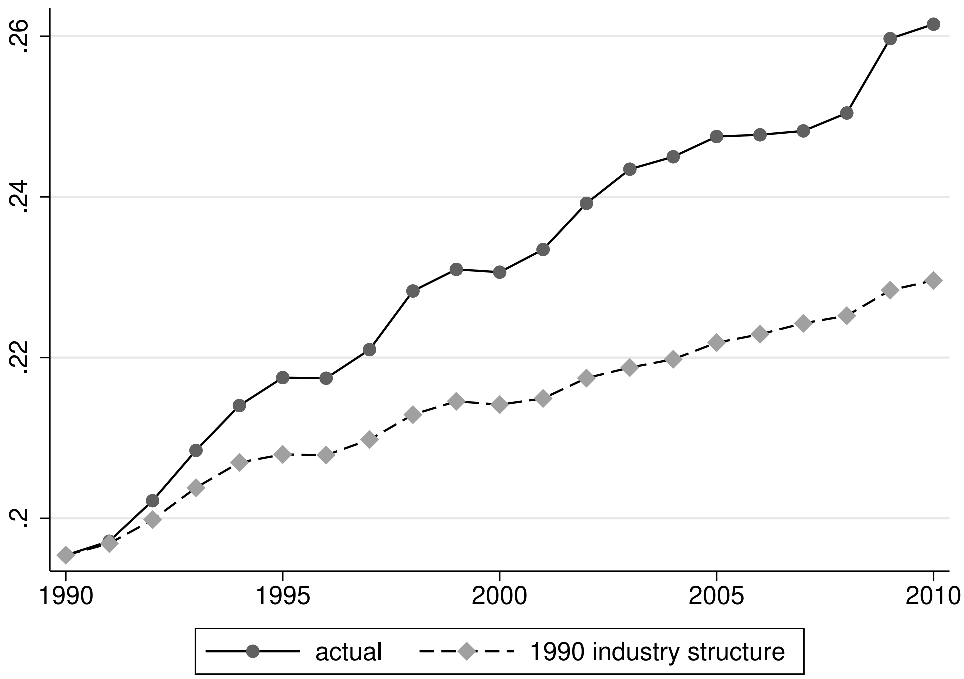

Note: The figure shows the evolution of the share of aggregate employment in abstract occupations in Germany between 1990 and 2010 based on data from the BEH social security records. The solid line uses the contemporaneous industry structure in each year, while the dashed line fixes the industry composition using 1990 employment shares and therefore captures only changes in the task composition of employment within industries (see Appendix B.1). The mapping of detailed occupation codes to broad task categories is detailed in Appendix Table A.1.

Figure 2: Evolution of the Variance of Log Wages

Note: Panel A displays the evolution of the variance of individual log wages and decomposes it into within- and between-task components (based on two broad task groups); see Equation (1). Panel B decomposes the within-task component into within- and between-establishment components; see Equation (2). Panel C decomposes the within-task between-establishment component into within- and between-industry components; see Equation (3). Panel D plots the actual and counterfactual within-task between-establishment within-industry wage variance, where the within-task between-establishment wage variance in an industry are averaged across industries using the actual and 1990 industry employment structure; see Equation (4).

Figure 3: Evolution of the Variance of Log Wages (and Decomposition Based on 3-Digit Occupations)

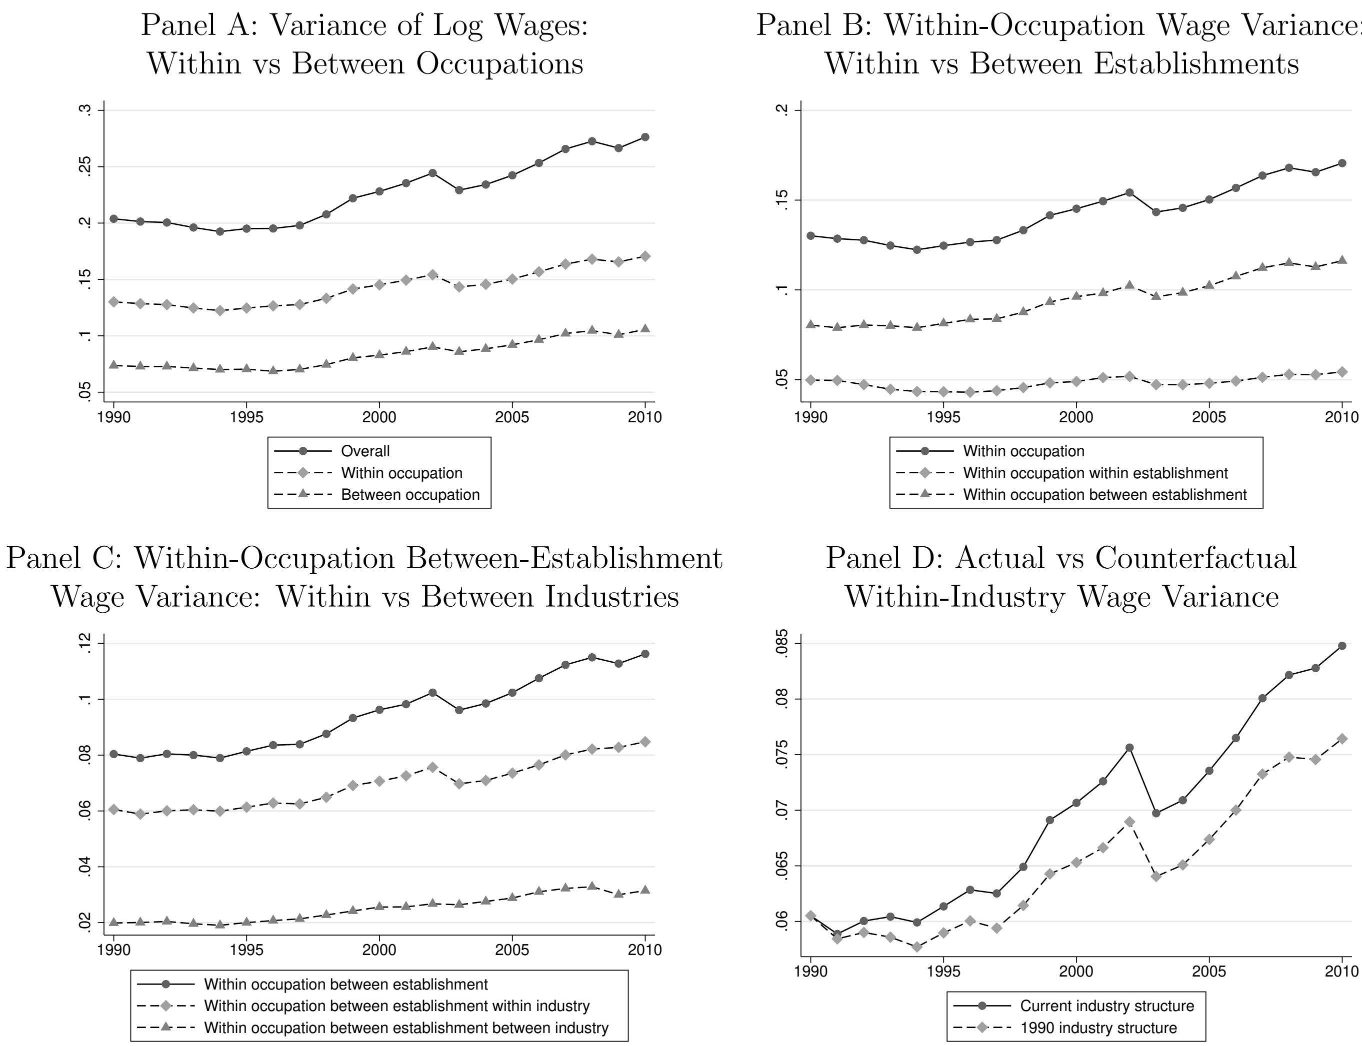

Note: Panel A displays the evolution of the variance of log wages and decomposes it into within- and between-occupation components (based on 317 occupation groups); see Equation (1). Panel B decomposes the within-occupation component into within- and between-establishment components; see Equation (2). Panel C decomposes the within-occupation between-establishment component into within- and between-industry components; see Equation (3). Panel D plots the actual and counterfactual within-occupation between-establishment within-industry wage variance, where the within-occupation between-establishment wage variance in an industry are averaged across industries using the actual and 1990 industry employment structure; see Equation (4).

Figure 4: Establishment vs Task/Occupation Premiums

Panel A: Establishment Premium Gaps (Two Tasks) and Abstract Wage Premium

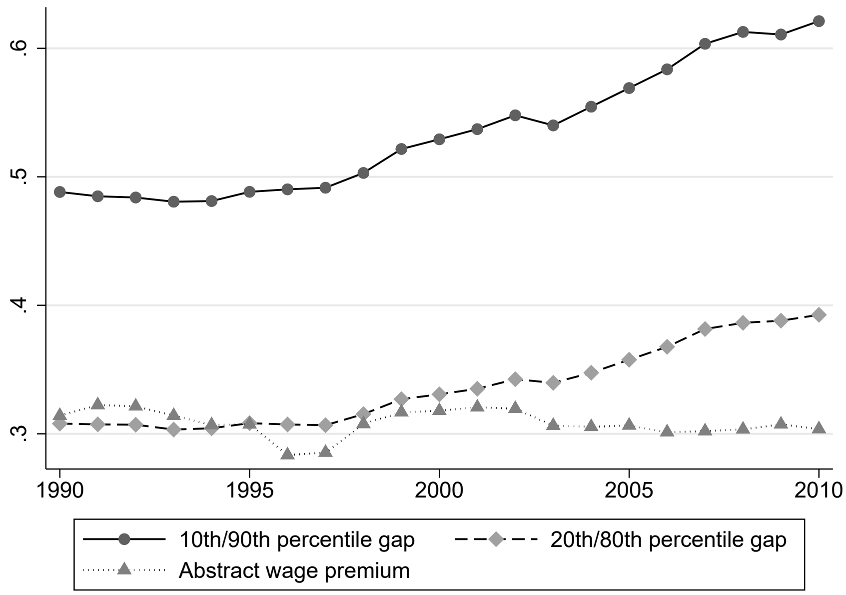

Panel B: Establishment Premium Gaps (Occupations) and Occupation Wage Premiums

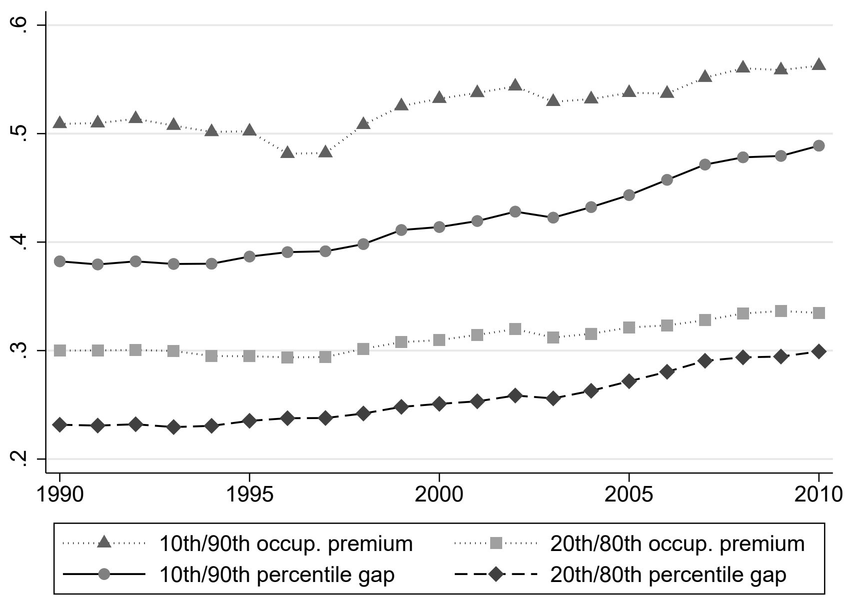

Note: Panel A shows the evolution of the gap between the 90th and 10th and the 80th and 20th percentile in within-industry establishment wage premiums (broad tasks), alongside the abstract wage premium. Panel B displays the evolution of the gap between the 90th and 10th and the 80th and 20th percentile in within-industry establishment wage premiums (317 occupations). It further shows the gap between the 90th and 10th and the 80th and 20th percentile in within-industry occupational wages over time. The within-industry values are averaged across industries using 1990 industry shares.

Figure 5: Year-by-Year Associations between Establishment Productivity and Other Establishment Characteristics

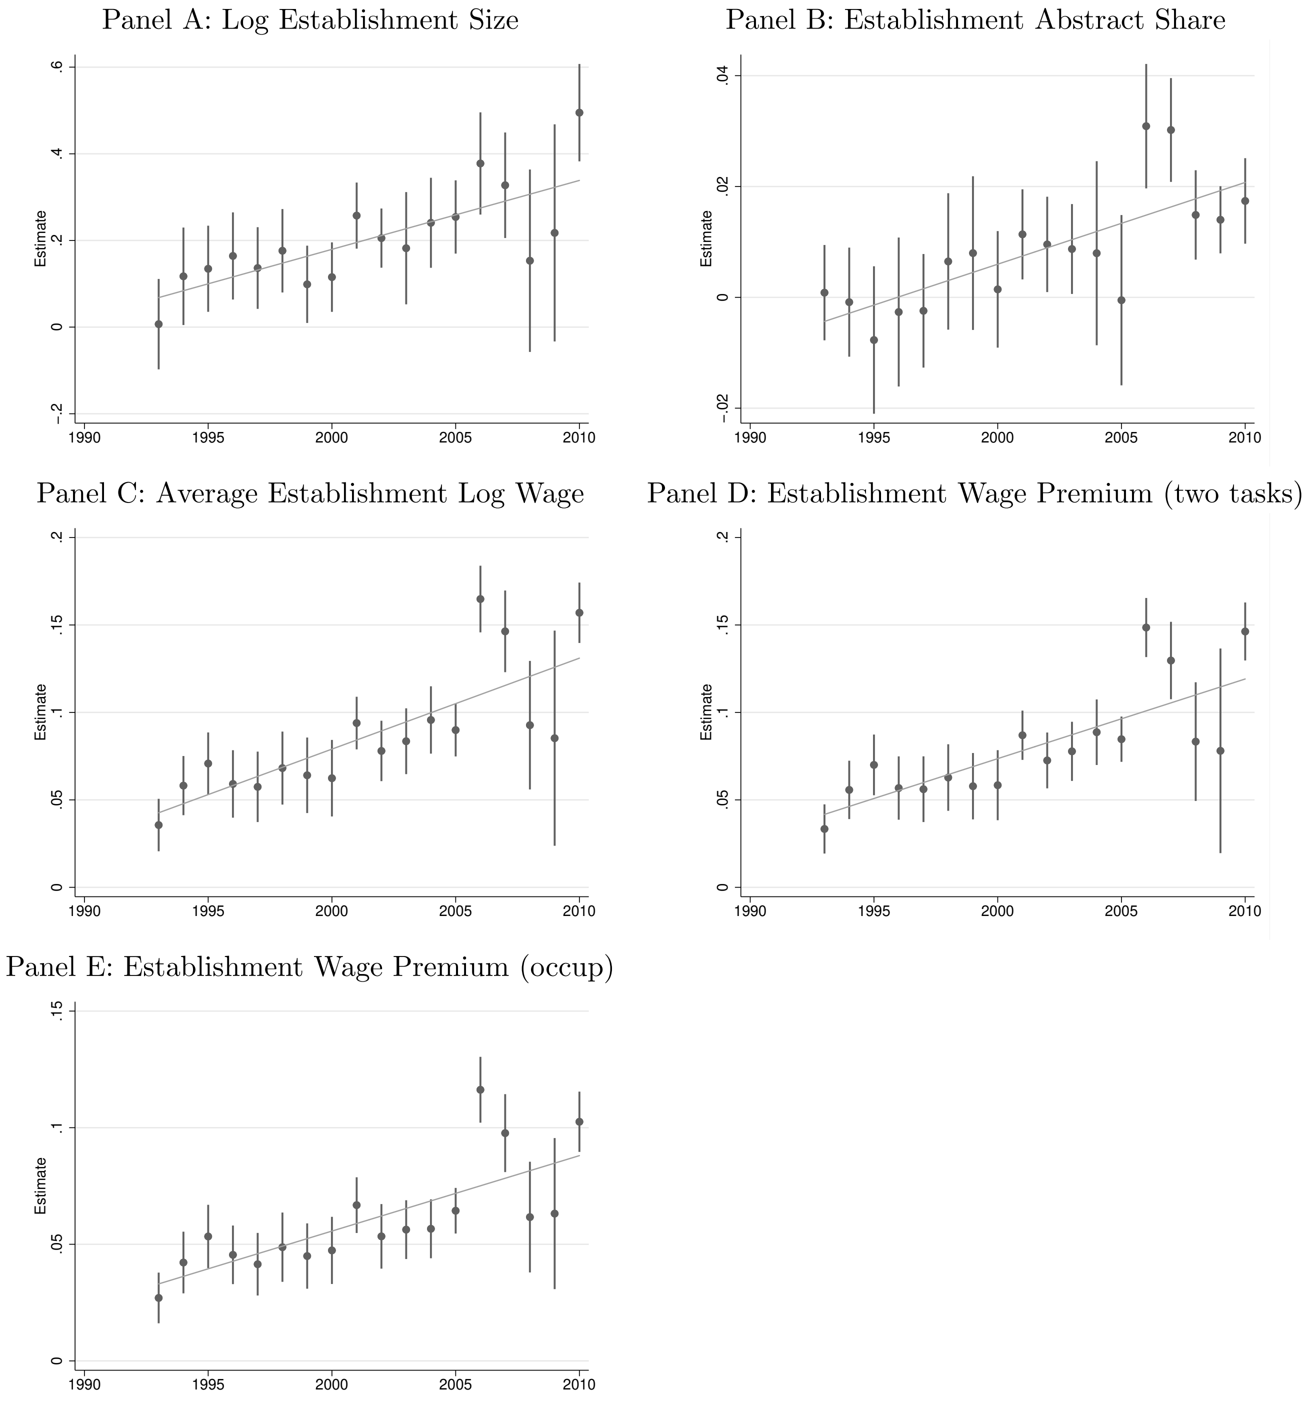

Note: The figure shows estimated coefficients and 95% confidence intervals from regressions of the outcome that appears in the title of each panel on log establishment productivity and a full set of 3-digit industry fixed effects, with the regressions estimated separately for each year. Results are based on establishments in the IABEP and observations are weighted by establishment size and survey weight. Establishment wage premiums adjust for the task (two broad tasks) and occupation (317 occupations) composition of the establishment; see text for details.

Figure 6: Year-by-Year Associations between Establishment Size and Other Establishment Characteristics

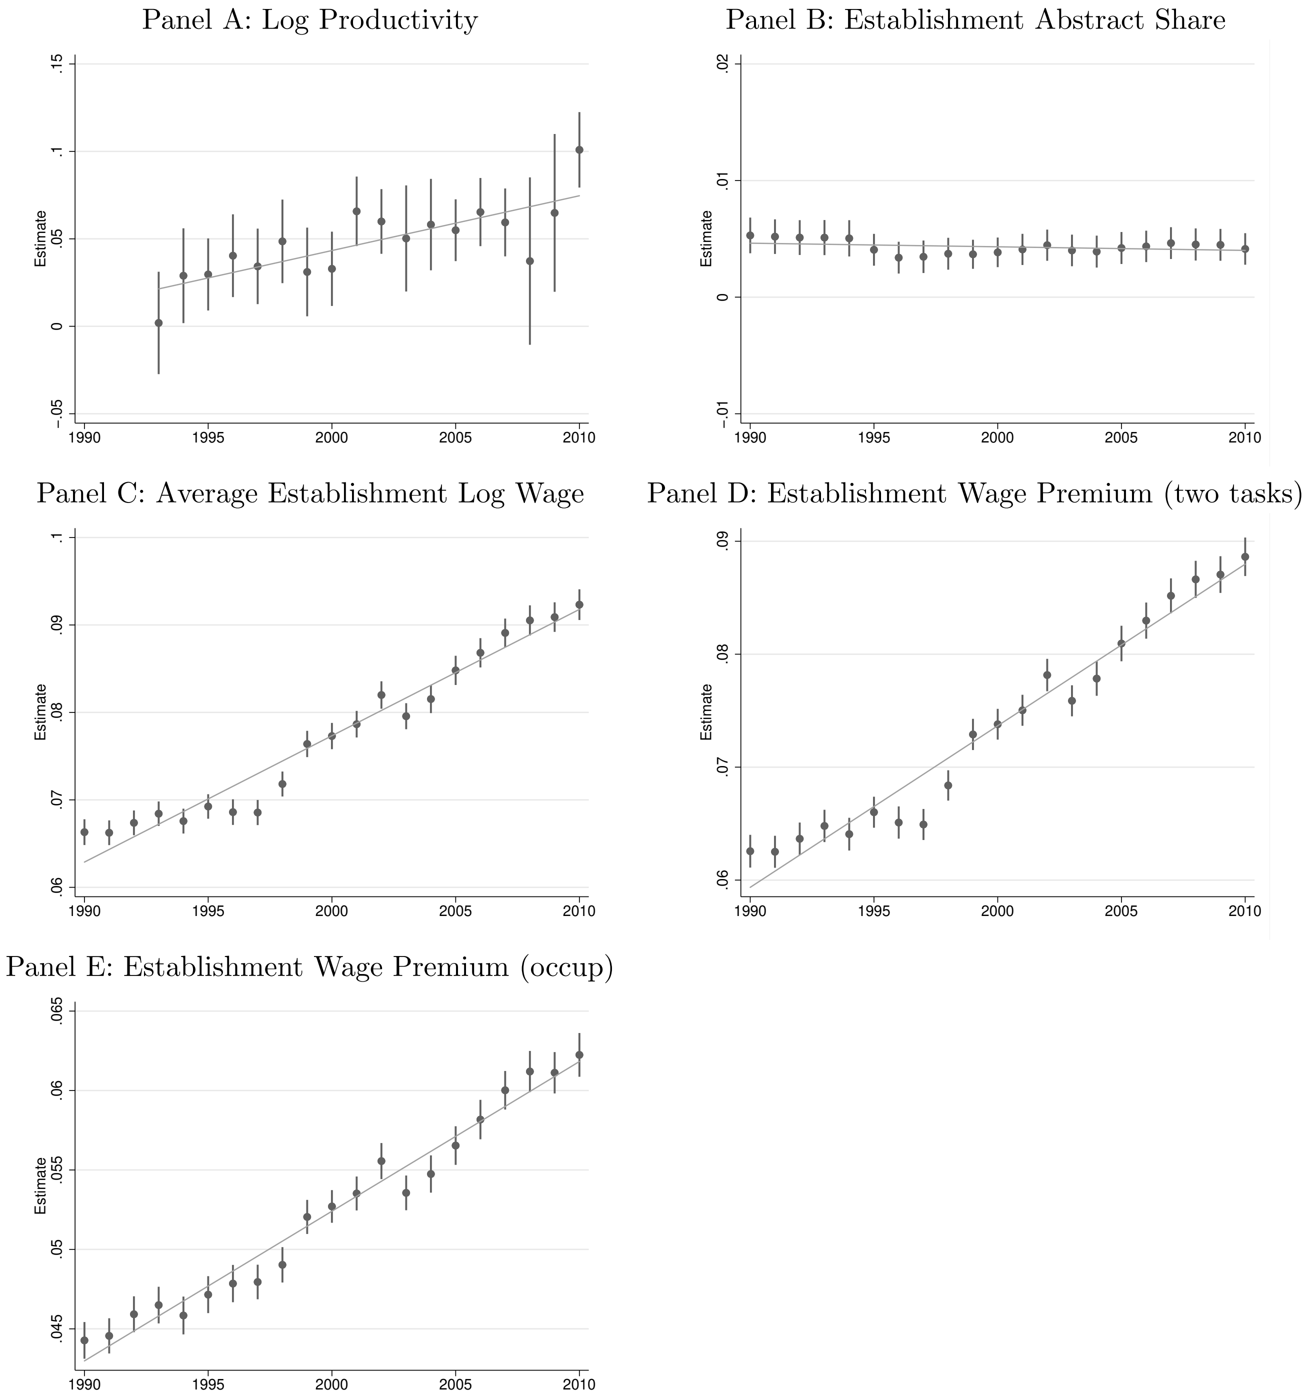

Note: The figure shows estimated coefficients and 95% confidence intervals from regressions of the outcome that appears in the title of each panel on log establishment size and a full set of 3-digit industry fixed effects, with the regressions estimated separately for each year. Results are based on establishments in the BEH and observations are weighted by establishment size (except Panel A which uses establishments in the IABEP and weights observations based on establishment size and survey weights). Establishment wage premiums adjust for the task (two broad tasks) and occupation (317 occupations) composition of the establishment; see text for details.

Figure 7: Abstract Share Heterogeneity and Sorting (Within Industries)

Panel A: Variance of Abstract Employment Shares

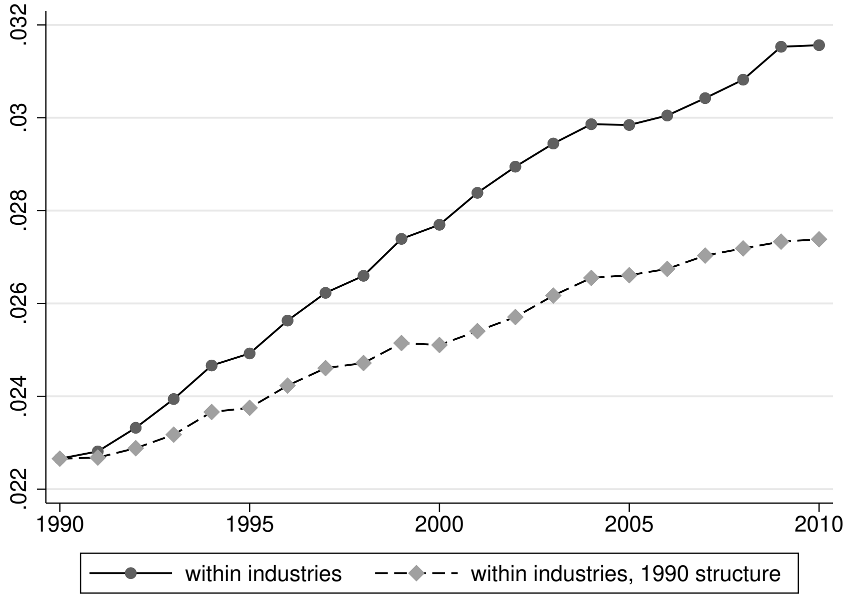

Panel B: Co-variance between Abstract Share and Establishment Wage Premium

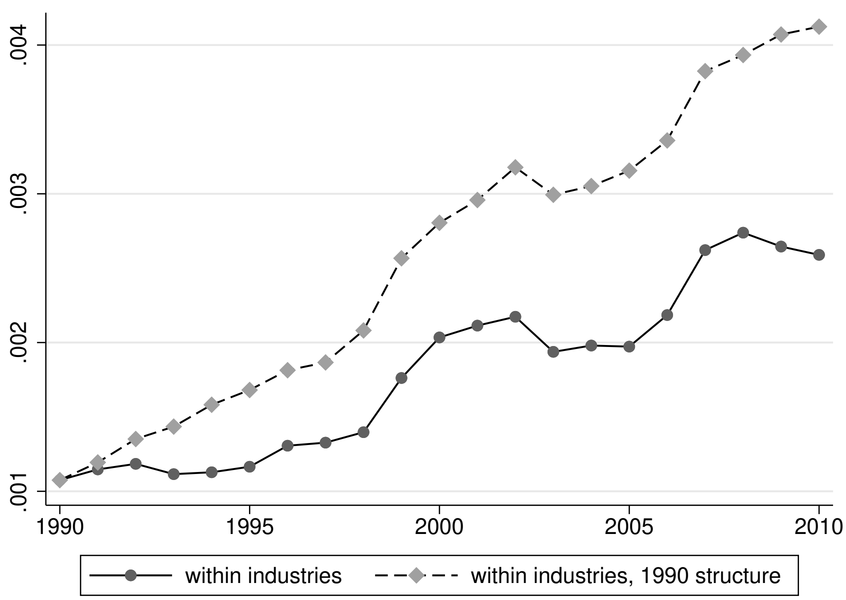

Note: Panel A shows the evolution of the variance of the abstract employment share across establishments within industries, averaging across industries using either observed industry employment shares in each year (solid line) or constant 1990 industry employment shares in all periods (dashed line). Panel B shows the co-variance between establishments' abstract employment shares and their wage premium (two tasks). Establishment wage premiums are computed as the establishment average of the residual of an individual log wage regression, estimated separately for each year, that controls for task type interacted with 3-digit industry fixed effects.

Figure 8: Decomposition of Changes in the Within-Industry Variance of Log Wages between Establishments

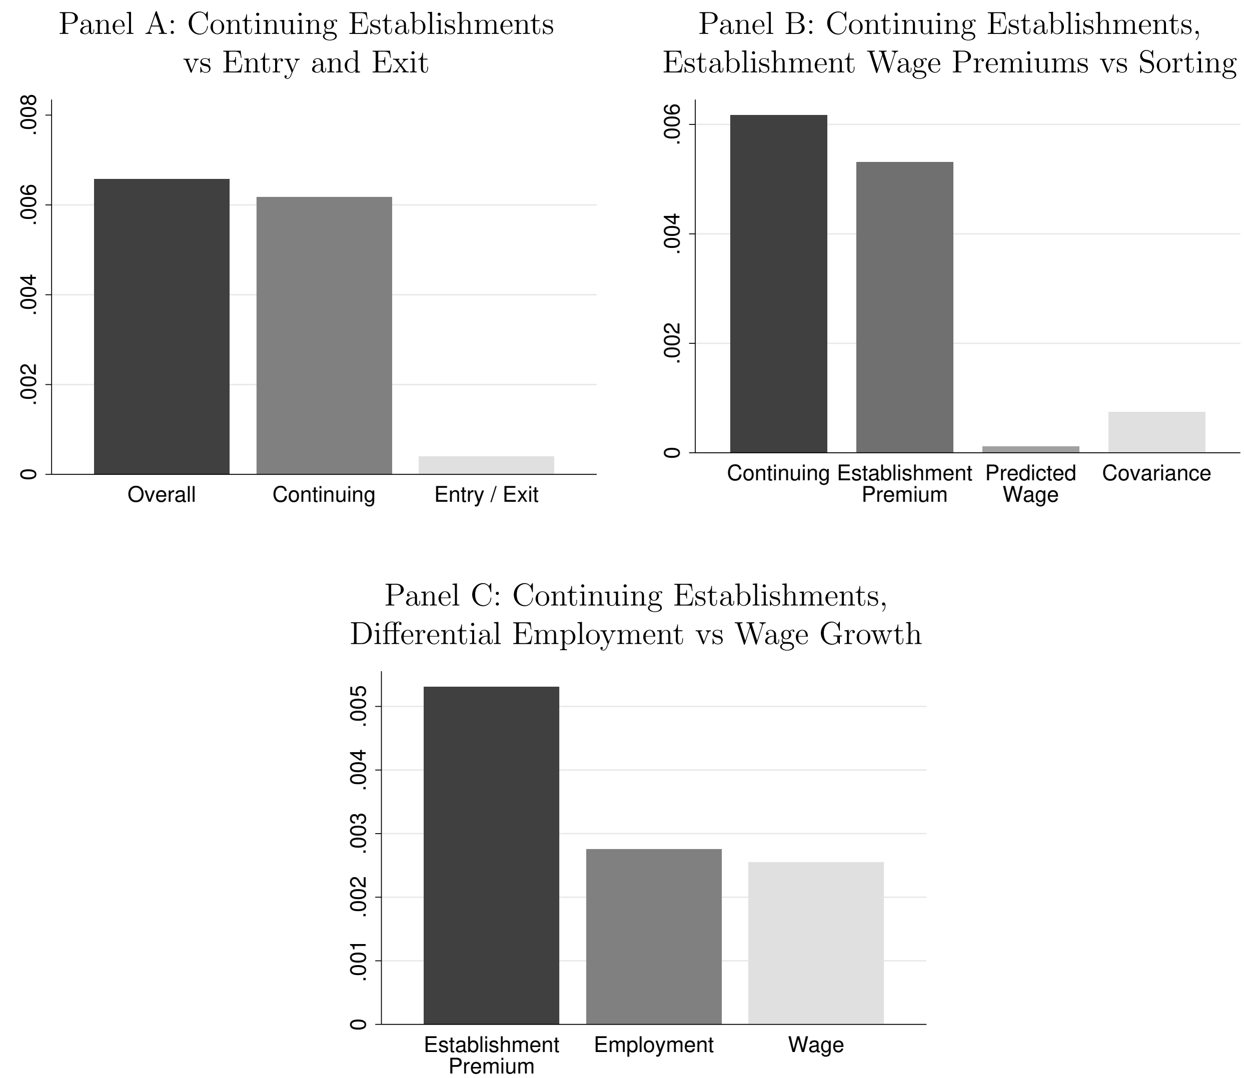

Note: Panel A decomposes changes in the within-industry, between-establishment variance of log wages into two components: selection due to establishment entry and exit, and changes among continuing establishments; see Equation (19). Panel B decomposes changes in the within-industry variance of log wages among continuing establishments into components related to establishment premiums, dispersion in task usage (predicted wage) and sorting (covariance); see Equation (20). Panel C decomposes changes in the within-industry variance of establishment wage premiums among continuing establishments into a differential employment and wage growth component; see Equation (21). The decompositions are performed over non-overlapping 5-year windows within industries and then averaged across industries and time using 1990 employment shares as weights and giving equal weight to each time period.

Figure 9: Decomposition of Within-Industry Changes in the Abstract Employment Share

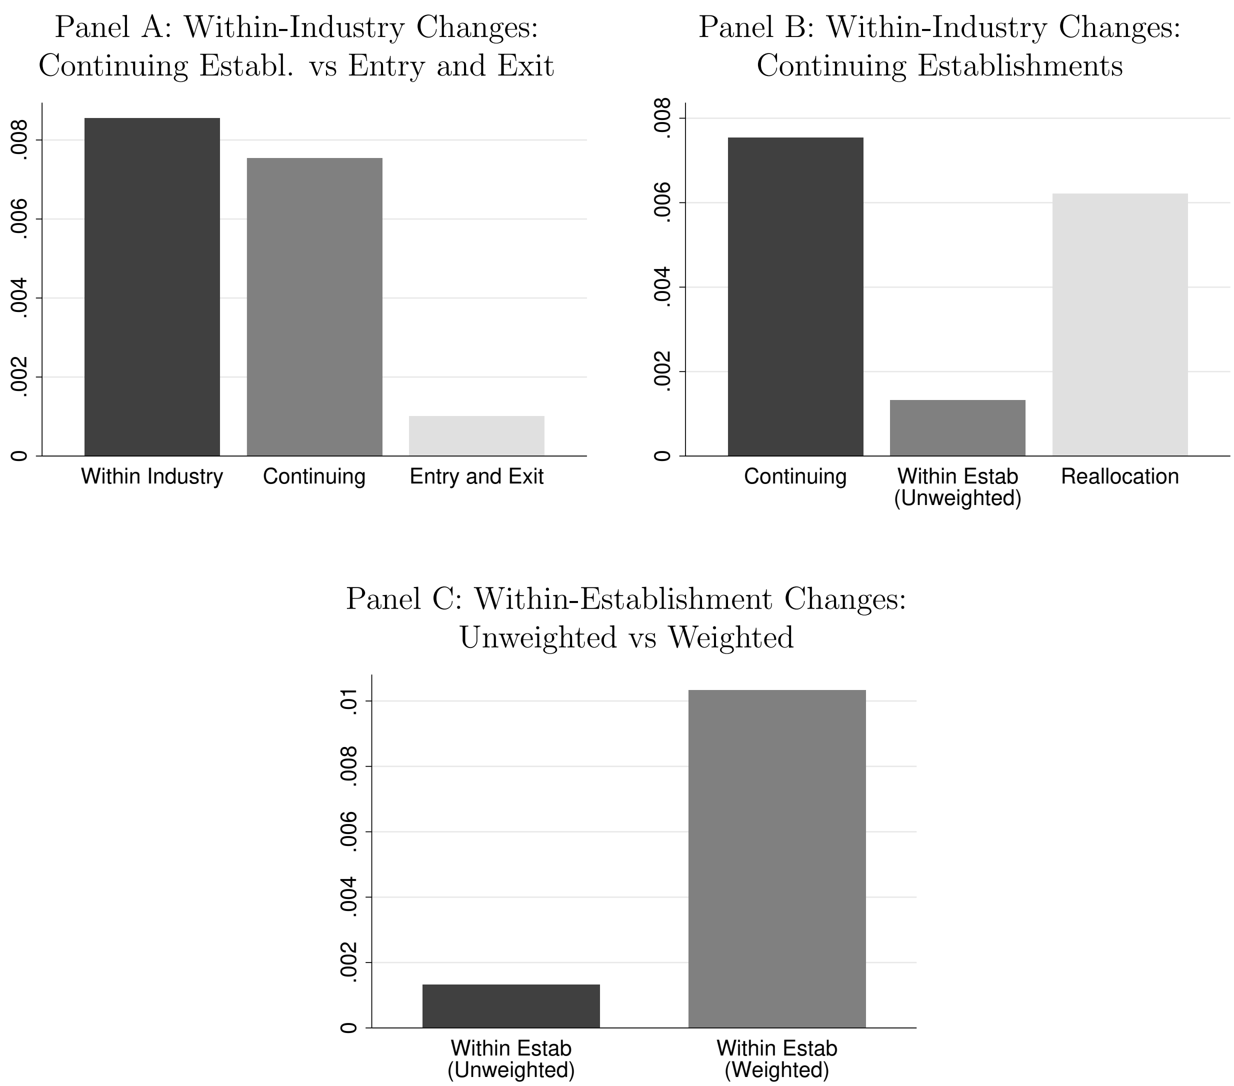

Note: The figure decomposes within-industry changes in the abstract employment share. Panel A decomposes the overall change in the industry into changes among continuing establishments and changes stemming from selective establishment entry and exit; see Equation (22). Panel B decomposes the within-industry change among continuing establishments into a within-establishment and a reallocation component; see Equation (23). Panel C contrasts the unweighted within-establishment change in the abstract employment share where all establishments are given the same weight with the weighted change where larger establishments at baseline are given a larger weight. The decompositions are performed over non-overlapping 5-year windows within industries and then averaged across industries and time using 1990 employment shares as weights and giving equal weight to each time period.

Figure 10: Industries with Below vs Above Median Increases in Abstract Employment Shares

Panel A: Change in Variance of Establishment Wages

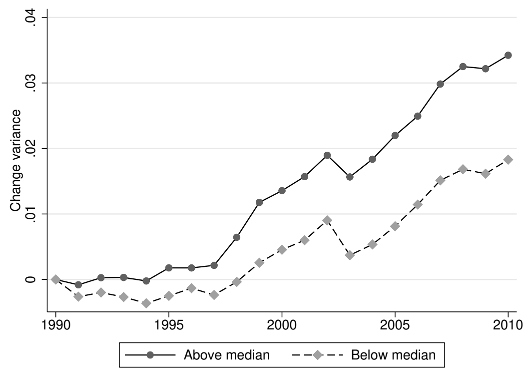

Panel B: Change in Variance of Establishment Abstract Shares

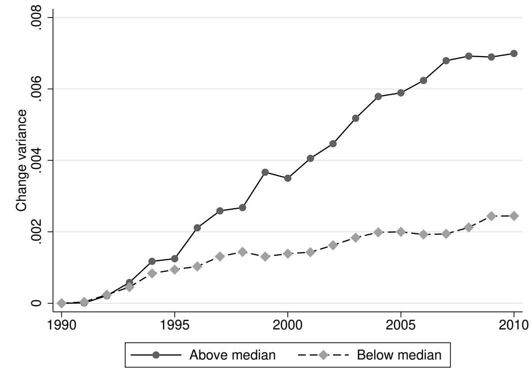

Panel C: Change in Variance of Establishment Wage Premium (tasks)

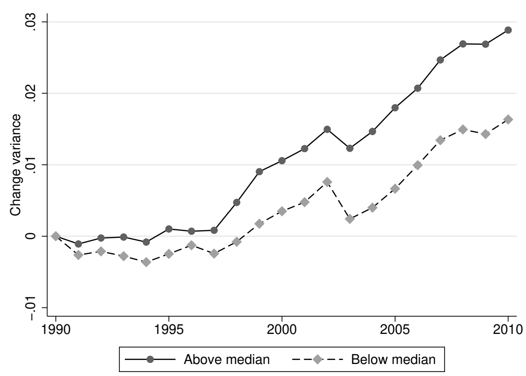

Panel D: Change in Variance of Establishment Wage Premium (occup)

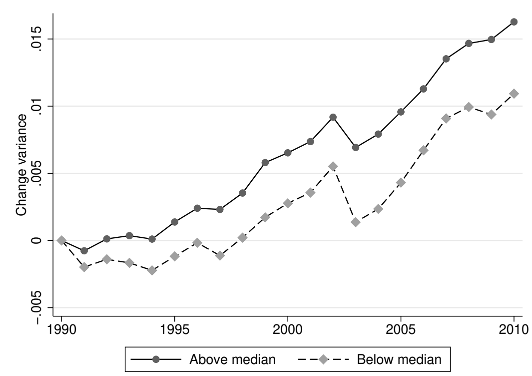

Panel E: Change in Co-variance of Establishment Wage Premium and Abstract Share

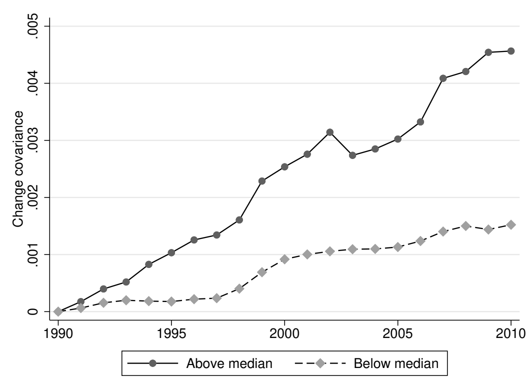

Note: The figures contrast the evolution of the variance of average establishment wages (Panel A), the variance of establishments' abstract employment shares (Panel B), the variance of establishments' wage premiums adjusting for their broad task and detailed occupation structure (Panels C and D), and the co-variance between establishments' abstract employment shares and wage premiums (Panel E) for two types of industries: industries with below median and above median changes in the industry-level abstract employment share between 1990 and 2010. We average across industries using the 1990 industry employment structure as weights.

Figure 11: Industries with Below vs Above Median Robot Adoption

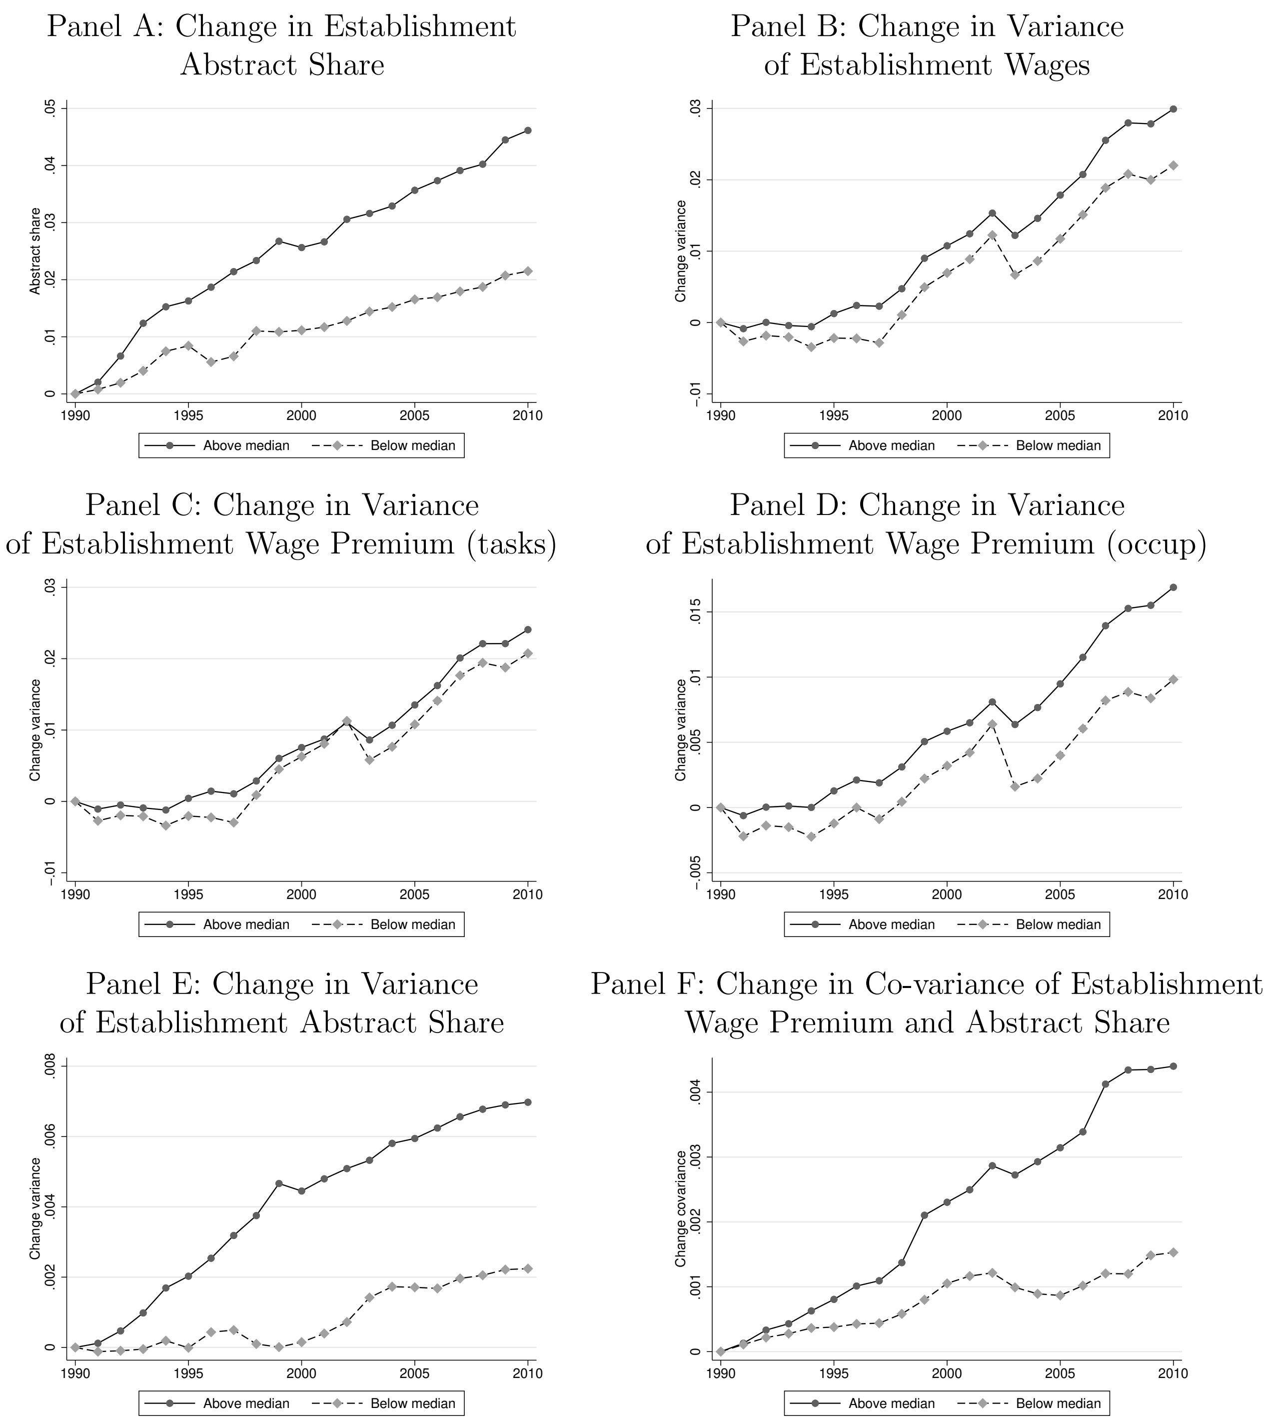

Note: The figures contrast the evolution of the increase in the abstract employment share (Panel A), the variance of average establishment wages (Panel B), the variance of establishments' wage premiums adjusting for their broad task and detailed occupation structure (Panels C and D), the variance of establishments' abstract employment shares (Panel E), and the co-variance between establishments' abstract employment shares and wage premiums (Panel F) for two types of industries: industries with below median and above median robot adoption between 1993 and 2010. We average across industries using the 1990 industry employment structure as weights.

Figure 12: Industries with Below and Above Median ICT Capital Adoption

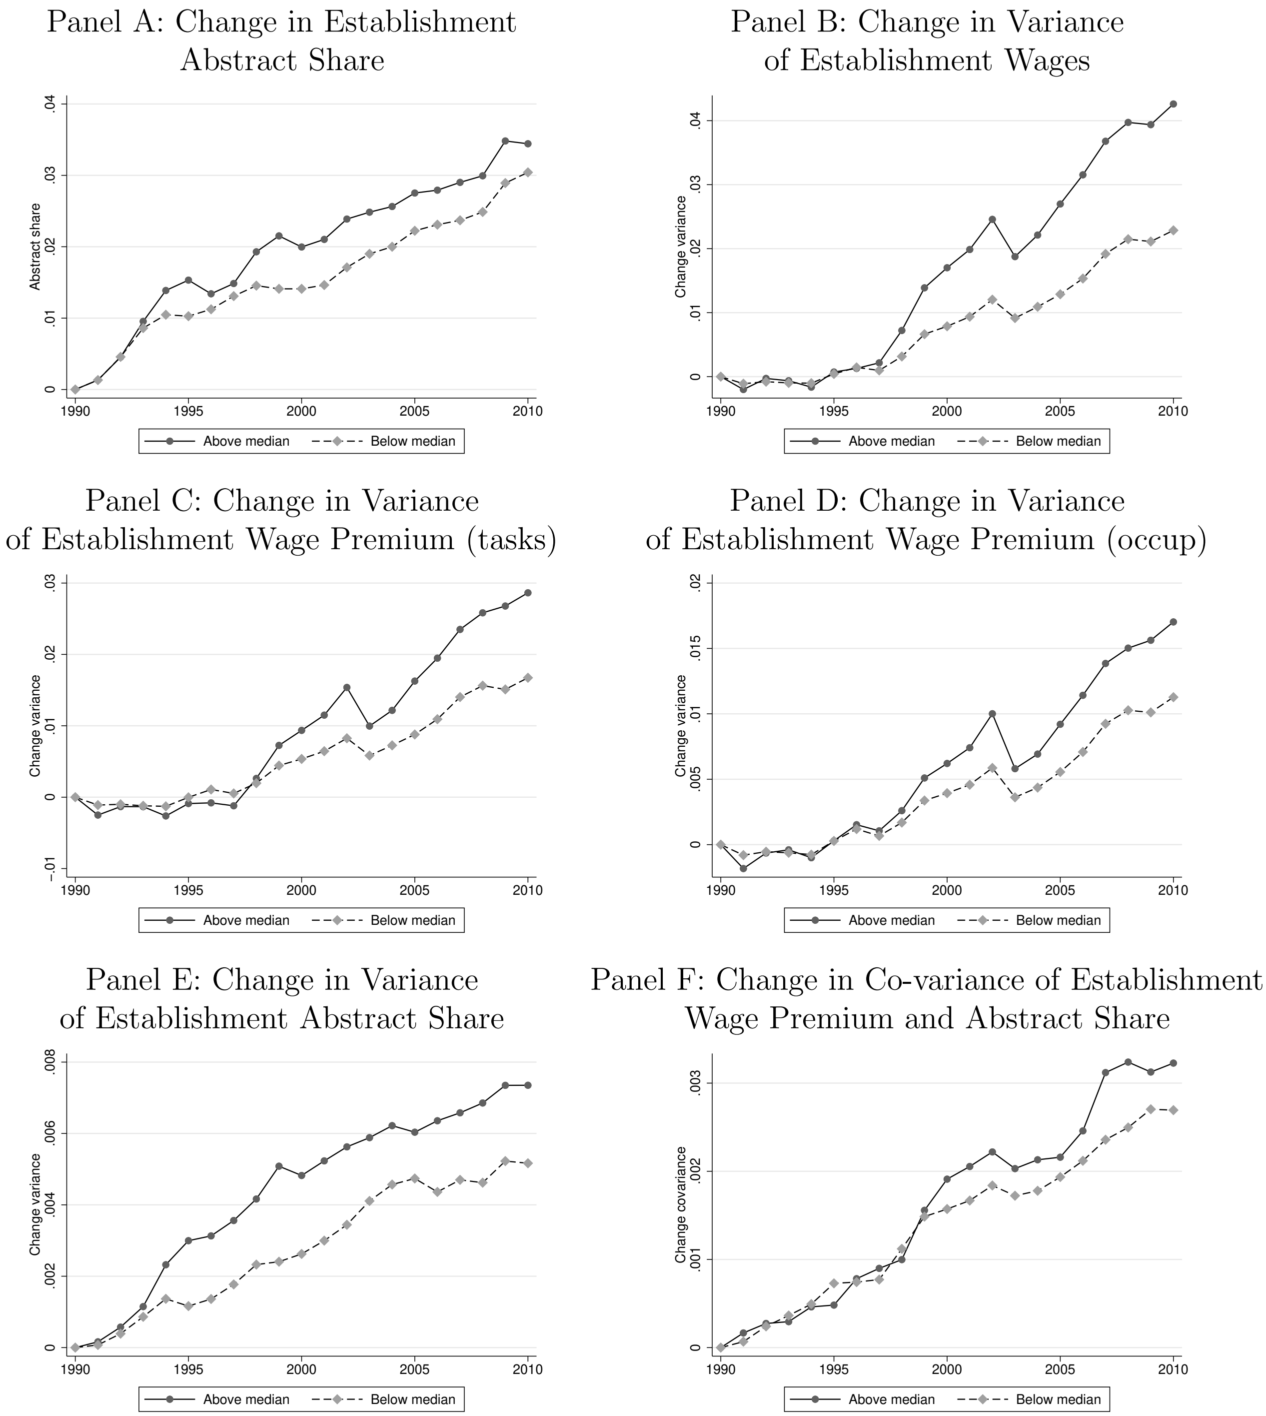

Note: The figures contrast the evolution of the increase in the abstract employment share (Panel A), the variance of average establishment wages (Panel B), the variance of establishments' wage premiums adjusting for their broad task and detailed occupation structure (Panels C and D), the variance of establishments' abstract employment shares (Panel E), and the co-variance between establishments' abstract employment shares and wage premiums (Panel F) for two types of industries: industries with below median and above median increases in ICT capital between 1993 and 2010. We average across industries using the 1990 industry employment structure as weights.

Table 1: Cross-Sectional Relationships Between Establishment Productivity and Abstract Share (Within Industry-Year)

	Abstract Share		Productivity	Abstract Share
	(1)		(2)	(3)
Log Productivity	0.0065^{**}	Log Estab Size	0.047^{***}	0.0043^{***}
(Rev. p. Worker)	(0.003)	(# of employees)	(0.006)	(0.001)
N	86,883		86,883	26,814,744

Note: All regressions include a set of fully interacted 3-digit industry and year fixed effects. Observations are weighted by establishment size and survey weights in Columns (1) and (2) and by establishment size in Column (3). Standard errors are clustered at the establishment level. Productivity is measured as total sales per full-time equivalent worker using sales data from the IABEP. Columns (1) and (2) are based on establishments observed in the IABEP; Column (3) is based on establishments observed in the full BEH data; see Sections 2.1 and 2.2 for details on sample restrictions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Cross-Sectional Relationships between Wages, Productivity and Employment (Within-Industry and Year)

Panel A: Relationship with (Log) Productivity (IABEP)

	Avg. Log wage	Establishment Premium (tasks)	Establishment Premium (occup)	Log Abstract Workers	Log Routine Workers
	(1)	(2)	(3)	(4)	(5)
Log Productivity	0.080***	0.075***	0.056***	0.21***	0.21***
(Rev. p. Worker)	(0.0050)	(0.0047)	(0.0035)	(0.033)	(0.024)
N	86,883	86,883	86,883	58,089	82,722

Panel B: Relationship with Establishment Size (BEH)

	Avg. Log wage	Establishment Premium (tasks)	Establishment Premium (occup)
	(1)	(2)	(3)
Log Estab Size	0.077***	0.073***	0.052***
(# of employees)	(0.00064)	(0.00061)	(0.00049)
N	26,814,744	26,814,744	26,814,744

Note: All regressions include a set of fully interacted 3-digit industry and year fixed effects. Productivity is measured as total sales per full-time equivalent worker using sales data from the IABEP. Establishment wage premiums adjust for the task (two broad tasks) and occupation (317 occupations) composition of the establishment, and are computed as the establishment average of the residual of an individual wage regression, estimated separately for each year, that controls for task type (occupations) interacted with 3-digit industry fixed effects. Panel A is based on establishments observed in the IABEP and observations are weighted by establishment size and survey weight. Panel B is based on establishments observed in the full BEH data and observations are weighted by establishment size. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Baseline Establishment Size and Within-Establishment Changes (Within Industries)

	Δ Estab Productivity	Δ Abstract Share	Δ Avg. Log Wage	Δ Estab Premium (Tasks)	Δ Estab Premium (Occup)
	(1)	(2)	(3)	(4)	(5)
Estab size at baseline	0.032***	0.0030***	0.0081***	0.0069***	0.0047***
	(0.0096)	(0.00025)	(0.00034)	(0.00028)	(0.00023)
N	5,460	3,452,385	3,452,385	3,452,385	3,452,385

Note: The table shows estimated coefficients from regressions of within-establishment changes in the outcome variable shown in each column of the table on baseline establishment size. Regressions include a set of fully interacted 3-digit industry and year fixed effects. Within-establishment changes are taken over non-overlapping 5-year windows. With the exception of Column (1), results are based on establishments in the full BEH. Observations are weighted by establishment size, and standard errors are clustered at the establishment level. Column (1) uses establishments in the IABEP. Productivity is measured as total sales per full-time equivalent worker using sales data from the IABEP, and establishments are weighted based on total employment and survey weights. Establishment wage premiums adjust for the task (two broad tasks) and occupation (317 occupations) composition of the establishment, and are computed as the establishment average of the residual of an individual log wage regression, estimated separately for each year, that controls for task type (occupations) interacted with 3-digit industry fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Online Appendix for:

Technological Change, Firm Heterogeneity and Wage Inequality

Guido Matias Cortes (York University)
Adrian Lerche (Institute for Employment Research)
Uta Schönberg (University College London and Institute for Employment Research)
Jeanne Tschopp (University of Bern)

Appendix A Data

Appendix A.1 Imputation of Censored Wages

To impute top-coded wages, we first define age-education cells based on five age groups (with 10-year intervals) and three education groups (no post-secondary education, vocational degree, college or university degree). Within each of these cells, following Dustmann et al. (2009) and Card et al. (2013), we estimate Tobit wage equations separately by year while controlling for age; firm size (quadratic, and a dummy for firm size greater than 10); occupation dummies; the focal worker's mean wage and mean censoring indicator (each computed over time but excluding observations from the current time period); and the firm's mean wage, mean censoring indicator, mean years of schooling, and mean university degree indicator (each computed at the current time period by excluding the focal worker observations). For workers observed in only one time period, the mean wage and mean censoring indicator are set to sample means, and a dummy variable is included. A wage observation censored at value c is then imputed by the value $\mathrm{X}\hat{\beta} + \hat{\sigma}\Phi^{-1}[\mathrm{k} + \mathrm{u}(1 - \mathrm{k})]$, where Φ is the standard normal CDF, u is drawn from a uniform distribution, $\mathrm{k} = \Phi[(\mathrm{c} - \mathrm{X}\hat{\beta})/\hat{\sigma})]$ and $\hat{\beta}$ and $\hat{\sigma}$ are estimates for the coefficients and standard deviation of the error term from the Tobit regression.

Appendix A.2 Harmonization of Industry Codes

In 1999, the industry classification in the BEH social security data changed in order to make the industry classification compatible with international NAICS codes. In 1999, both the old and new industry code are included for all establishments. We use this information to compute the most common NAICS industry code for each old industry code. For establishments that still exist in 1999, we assign their 1999 NAICS industry code for all earlier years.

For establishments that have exited by 1999, we assign the NAICS industry code that, in 1999, is the most common given the establishment's old industry code.

Appendix B Decomposition Details

Appendix B.1 Evolution of Abstract Share Holding Industry Structure Constant (Figure 1)

The actual abstract share in year t (the black line in Figure 1) is a weighted average of the abstract share in each industry k:

$$S_t = \sum_k \omega_{kt} S_{kt} \tag{B.1}$$

where ω_{kt} denotes the employment share of industry k at time t.

The counterfactual abstract share in year t holding the industry structure constant at its 1990 level (the grey line in Figure 1) can then be computed as:

$$S_t^{1990} = \sum_k \omega_{k1990} S_{kt} \tag{B.2}$$

Appendix B.2 Abstract Share and Sorting within Industries (Figure 7)

The within-industry variance in establishments' abstract employment share in year t is computed as:

$$Var_t(S_{f(k)t}) = \sum_k \omega_{kt} \sum_f \omega_{f(k)t}(S_{f(k)t} - S_{kt})^2 \tag{B.3}$$

where ω_{kt} denotes the employment share of industry k in year t, $\omega_{f(k)t}$ denotes the employment share of establishment f in industry k at time t, and $S_{f(k)t}$ and S_{kt} the abstract employment share in establishment f and industry k, respectively.

The counterfactual within-industry variance in establishments' abstract employment share, holding the industry structure constant at its 1990 level, equals:

$$Var_{1990}(S_{f(k)t}) = \sum_k \omega_{k1990} \sum_f \omega_{f(k)t}(S_{f(k)t} - S_{kt})^2 \tag{B.4}$$

Let $\widetilde{lnw}_{f(k)t}$ denote establishment f's wage premium, computed as the residual from an individual-level wage regression, estimated separately for each year, of log wages on an indicator variables for task usage (two tasks) interacted with a full set of industry fixed effects, averaged to the level of the establishment. The within-industry co-variance between establishments' abstract employment shares and their wage premiums is computed as follows:

$$Cov_t(S_{f(k)t}, \widetilde{lnw}_{f(k)t}) = \sum_k \omega_{kt} \sum_f \omega_{f(k)t}(S_{f(k)t} - S_{kt})(\widetilde{lnw}_{f(k)t} - \overline{\widetilde{lnw}}_{kt}) \tag{B.5}$$

The counterfactual within-industry co-variance, holding the industry structure constant at its 1990 level, equals:

$$Cov_{1990}(S_{f(k)t}, \widetilde{lnw}_{f(k)t}) = \sum_k \omega_{k1990} \sum_f \omega_{f(k)t}(S_{f(k)t} - S_{kt})(\widetilde{lnw}_{f(k)t} - \overline{\widetilde{lnw}}_{kt}) \tag{B.6}$$

Appendix C Model

This section contains details of the model and of derivations that were omitted in the main text. The presentation is not necessarily self-contained but rather complementary with Section 4 of the paper. We refer to Section 5.4 of the technical appendix of Helpman et al. (2010) for more details on the model with a CES production function and two types of workers.

For the derivations below, it is useful to note that $\varphi(\theta)$ and ϕ_ℓ, where $\ell \in \{s, r\}$ are defined as follows:

$$\varphi \equiv \frac{\mu_s^\nu \left(\theta \overline{a}_s h_s^\gamma\right)^\nu}{\left(\overline{a}_r h_r^\gamma\right)^\nu}, \qquad \phi_s \equiv \frac{\varphi}{1+\varphi}, \qquad \phi_r \equiv \frac{1}{1+\varphi}. \tag{C.1}$$

Appendix C.1 Derivations of the Key Equilibrium Relationships

This section derives the equilibrium relationships for the variables which play a crucial role when examining the impact of task-biased technological change on wage inequality.

Appendix C.1.1 Firm-level Equilibrium Variables

Below we use the following first-order conditions from the profit maximization problem to derive firm-level equilibrium revenues, employment and wages by tasks:

$$\frac{\beta\gamma}{1+\beta\gamma}\phi_\ell r(\theta) = b_\ell n_\ell(\theta) \tag{C.2}$$

$$\frac{\beta(1-\gamma k)}{1+\beta\gamma}\phi_\ell r(\theta) = c\widetilde{a}_\ell(\theta)^\delta \tag{C.3}$$

Revenues As Helpman et al. (2010) mention in Appendix 5.4 footnote 1, revenues can be expressed as:

$$r(\theta) = \kappa_y^\beta A\left[1+\varphi(\theta)\right]^{\beta/\nu}\left[\widetilde{a}_r(\theta)^{1-k\gamma}n_r(\theta)^\gamma\right]^\beta, \tag{C.4}$$

where $\kappa_y \equiv \frac{ka_{min}^{\gamma k}}{k-1}$. Using the first-order conditions along with equation (C.4) and the definition of ϕ_r, one obtains the revenue equation:

$$r(\theta) = \kappa_r\left[1+\varphi(\theta)\right]^{\frac{\beta\Lambda}{\nu\Gamma}}, \tag{C.5}$$

where κ_r is equivalent to:

$$\kappa_r \equiv A^{1/\Gamma}\left[\kappa_y\left(\frac{\beta}{1+\beta\gamma}\right)^{\frac{1-k\gamma}{\delta}+\gamma}\left(\frac{1-\gamma k}{c}\right)^{\frac{1-k\gamma}{\delta}}\left(\frac{\gamma}{b_r}\right)^\gamma\right]^{\beta/\Gamma}. \tag{C.6}$$

Employment by task and abstract employment share To obtain firm-level employment, note that from equation (C.2):

$$\begin{aligned} n_r(\theta) &= \frac{\beta\gamma}{1+\beta\gamma}\left[1+\varphi(\theta)\right]^{-1}b_r^{-1}r(\theta) \\ &= \left(\frac{\beta\gamma}{1+\beta\gamma}\right)b_r^{-1}\kappa_r\left[1+\varphi(\theta)\right]^{\frac{\beta-\nu}{\nu\Gamma}}, \end{aligned} \tag{C.7}$$

where $\frac{\beta-\nu}{\nu\Gamma} = \frac{\beta\Lambda}{\nu\Gamma} - 1 > 0$, and from equation (C.3):

$$\begin{aligned} \widetilde{a}_r(\theta) &= \left\{\frac{\beta(1-\gamma k)}{1+\beta\gamma}\left[1+\varphi(\theta)\right]^{-1}c^{-1}r(\theta)\right\}^{1/\delta} \\ &= \left[\frac{\beta(1-\gamma k)}{1+\beta\gamma}\right]^{1/\delta}c^{-1/\delta}\kappa_r^{1/\delta}\left[1+\varphi(\theta)\right]^{\frac{\beta-\nu}{\delta\nu\Gamma}}. \end{aligned} \tag{C.8}$$

Using expression $h_\ell(\theta) = n_\ell(\theta)\left(\frac{a_{min}}{\widetilde{a}_\ell(\theta)}\right)^k$, along with (C.7) and (C.8), we have that:

$$\begin{aligned} h_r(\theta) &= n_r(\theta)\left(\frac{a_{min}}{\widetilde{a}_r(\theta)}\right)^k \\ &= \left(\frac{\beta\kappa_r}{1+\beta\gamma}\right)^{1-k/\delta}\left(\frac{c}{1-\gamma k}\right)^{k/\delta} b_r^{-1}a_{min}^k\,[1+\varphi(\theta)]^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)} && \text{(C.9)} \\ &= h_{dr}\,[1+\varphi(\theta)]^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)}, && \text{(C.10)} \end{aligned}$$

where:

$$h_{dr} \equiv \left(\frac{\beta\kappa_r}{1+\beta\gamma}\right)^{1-k/\delta}\left(\frac{c}{1-\gamma k}\right)^{k/\delta} b_r^{-1}a_{min}^k. \tag{C.11}$$

Proceeding in a similar way for firm-level employment of abstract workers, we obtain:

$$h_s(\theta) = \frac{b_r}{b_s}\varphi(\theta)^{1-k/\delta}h_r(\theta), \tag{C.12}$$

and it follows that the firm's employment share of abstract workers is given by:

$$\frac{h_s(\theta)}{h(\theta)} = \frac{b_r\varphi(\theta)^{1-k/\delta}}{b_s + b_r\varphi(\theta)^{1-k/\delta}}, \tag{C.13}$$

where $h(\theta) = h_s(\theta) + h_r(\theta)$.

Wages by task To derive equilibrium firm-level wages by task, it is useful to note that the solution of the Stole and Zwiebel bargaining game takes the following form:

$$w_\ell = \frac{\beta\gamma}{1+\beta\gamma}\frac{\phi_\ell r}{h_\ell} \tag{C.14}$$

Using (C.14) along with (C.4) and (C.10), we have that firm wages of routine workers are given by:

$$\begin{aligned} w_r(\theta) &= \frac{\beta\gamma}{1+\beta\gamma}\phi_r(\theta)\frac{r(\theta)}{h_r(\theta)} \\ &= \left(\frac{\beta\gamma}{1+\beta\gamma}\right)\left(\frac{\kappa_r}{h_{dr}}\right)[1+\varphi(\theta)]^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}} && \text{(C.15)} \\ &= w_{dr}\,[1+\varphi(\theta)]^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}}, && \text{(C.16)} \end{aligned}$$

where:

$$w_{dr} \equiv \left(\frac{\beta\gamma}{1+\beta\gamma}\right)\left(\frac{\kappa_r}{h_{dr}}\right). \tag{C.17}$$

Proceeding in a similar way for firm-level wages of abstract workers, we obtain:

$$w_s(\theta) = \frac{b_s}{b_r}\varphi(\theta)^{k/\delta} w_r(\theta). \tag{C.18}$$

Finally, combining the definition of $\varphi(\theta)$ together with the first-order conditions of the profit maximization problem, we obtain:

$$\varphi(\theta) = \mu_s^{\nu/\Lambda}\left(\frac{b_s}{b_r}\right)^{-\gamma\nu/\Lambda}\theta^{\nu/\Lambda}. \tag{C.19}$$

Appendix C.1.2 Determination of the Productivity Threshold

As is standard in Melitz-type heterogeneous firm models, the productivity threshold for production, θ_d, is pined down by both the Zero-Cutoff Profit (ZCP) and the Free Entry (FE) conditions.

The ZCP condition, which requires that the firm at the cutoff θ_d makes zero profits, implies:[19]

$$\frac{\Gamma}{1+\beta\gamma} r(\theta_d) = f_d. \tag{C.20}$$

Moreover, given equation (C.5), relative revenues across two firms with productivities θ_1 and θ_2 can be written as:

$$\frac{r(\theta_1)}{r(\theta_2)} = \left[\frac{1+\varphi(\theta_1)}{1+\varphi(\theta_2)}\right]^{\frac{\beta\Lambda}{\nu\Gamma}}. \tag{C.21}$$

Combining equation (C.21) along with the ZCP condition (C.20) we obtain:

$$r(\theta) = f_d\left(\frac{\Gamma}{1+\beta\gamma}\right)^{-1}\left[\frac{1+\varphi(\theta)}{1+\varphi(\theta_d)}\right]^{\frac{\beta\Lambda}{\nu\Gamma}}. \tag{C.22}$$

The FE condition states that the expected profits for a potential entrant should equal the fixed entry cost:

$$\int_{\theta_d}^{\infty} \pi(\theta) dG(\theta) = f_e. \tag{C.23}$$

[19] This is obtained by noting that profits can be written as:

$$\pi(\theta) = \frac{\Gamma}{1+\beta\gamma} r(\theta) - f_d.$$

Therefore, combining equations (C.22) and (C.23) implies:

$$f_d \int_{\theta_d}^{\infty} \left(\left[\frac{1+\varphi(\theta)}{1+\varphi(\theta_d)} \right]^{\frac{\beta\Lambda}{\nu\Gamma}} - 1 \right) dG(\theta) = f_e. \tag{C.24}$$

Equation (C.24) pins down the equilibrium threshold θ_d as a function of the parameters of the model and the search costs b_s and b_r.

Appendix C.2 The Relationship between Firm-specific Equilibrium Outcomes and Productivity

This section presents the proofs for the results in Equations (12) and (15) on the relationship between firm-level employment and wages and firm productivity.

First, note that:

$$\frac{\partial\varphi(\theta)}{\partial\theta} = \frac{\nu}{\Lambda}\mu_s^{\frac{\nu}{\Lambda}} \left(\frac{b_s}{b_r}\right)^{-\frac{\gamma\nu}{\Lambda}} \theta^{\frac{\nu}{\Lambda}-1} \quad > \quad 0, \tag{C.25}$$

and recall that $\beta > \nu$, $\Lambda > \Gamma$ and $\delta > k$ such that $\frac{\beta-\nu}{\nu\Gamma} > 0$ and $1 - \frac{k}{\delta} > 0$.

Prediction 1: More productive firms are larger, employing more of both types of workers.
Proof: Taking the derivative of equations (C.10) and (C.12), we obtain:

$$\begin{aligned} \frac{\partial h_r(\theta)}{\partial\theta} &= h_{dr}\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)[1+\varphi(\theta)]^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)-1} \cdot \frac{\partial\varphi(\theta)}{\partial\theta} && > \quad 0 \\ \frac{\partial h_s(\theta)}{\partial\theta} &= \frac{b_r}{b_s}\left[\left(1-\frac{k}{\delta}\right)\varphi(\theta)^{-\frac{k}{\delta}} \cdot \frac{\partial\varphi(\theta)}{\partial\theta} \cdot h_r(\theta) + \varphi(\theta)^{1-\frac{k}{\delta}} \cdot \frac{\partial h_r(\theta)}{\partial\theta}\right] && > \quad 0 \end{aligned} \tag{C.26}$$

Prediction 2: More productive firms have a higher employment share of abstract workers.
Proof: Taking the derivative of equation (C.13), we have that:

$$\frac{\partial}{\partial\theta}\left[\frac{h_s(\theta)}{h(\theta)}\right] = \frac{b_s b_r\left(1-\frac{k}{\delta}\right)\varphi(\theta)^{-\frac{k}{\delta}} \cdot \frac{\partial\varphi(\theta)}{\partial\theta}}{\left[b_s + b_r\varphi(\theta)^{1-\frac{k}{\delta}}\right]^2} \quad > \quad 0 \tag{C.27}$$

Prediction 3: More productive firms pay higher wages to both types of workers.

Proof: Taking the derivative of equations (C.16) and (C.18), we obtain:

$$\begin{aligned}
\frac{\partial w_r(\theta)}{\partial\theta} &= w_{dr}\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}\left[1+\varphi(\theta)\right]^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}-1}\cdot\frac{\partial\varphi(\theta)}{\partial\theta} && > \ 0 \\
\frac{\partial w_s(\theta)}{\partial\theta} &= \frac{b_s}{b_r}\left[\frac{k}{\delta}\varphi(\theta)^{\frac{k}{\delta}-1}\cdot\frac{\partial\varphi(\theta)}{\partial\theta}\cdot w_r(\theta)+\varphi(\theta)^{\frac{k}{\delta}}\cdot\frac{\partial w_r(\theta)}{\partial\theta}\right] && > \ 0
\end{aligned} \tag{C.28}$$

This result, combined with the prediction that more productive firms employ a higher share of abstract workers, unambiguously implies that firm average wages are increasing in firm productivity.

Appendix C.3 Impact of Task-Biased Technological Change

We model task-biased technological change (TBTC) as an increase in the parameter μ_s, i.e. as a factor-augmenting shock favoring abstract workers.

In order to evaluate how this task shock affects firms differentially across the productivity distribution we examine the second-order derivative of firm outcome variables, with respect to both the common abstract-augmenting technology parameter μ_s and firm productivity. To this end note that:

$$\frac{\partial\varphi(\theta)}{\partial\mu_s}=\frac{\nu}{\Lambda}\mu_s^{-1}\varphi(\theta) \ > \ 0 \ , \qquad \frac{\partial^2\varphi(\theta)}{\partial\mu_s\partial\theta}=\left(\frac{\nu}{\Lambda}\right)^2\mu_s^{-1}\theta^{-1}\varphi(\theta) \ > \ 0 \ , \tag{C.29}$$

and

$$\frac{\partial\varphi(\theta)}{\partial\mu_s}\frac{\partial\varphi(\theta)}{\partial\theta}=\varphi(\theta)\cdot\frac{\partial^2\varphi(\theta)}{\partial\mu_s\partial\theta}>0 \ . \tag{C.30}$$

Prediction 1: *Selection* – TBTC increases the productivity threshold for production θ_d.
Proof: We prove Prediction 1 by contradiction. Consider equation (C.24), which pins down the equilibrium threshold as a function of parameters of the model:

$$f_d\int_{\theta_d}^{\infty}\left(\left[\frac{1+\varphi(\theta)}{1+\varphi(\theta_d)}\right]^{\frac{\beta\Lambda}{\nu\Gamma}}-1\right)dG(\theta)=f_e \tag{C.31}$$

Suppose first that TBTC has no effect on θ_d. Holding θ_d fixed, the increase in $[1+\varphi(\theta)]/[1+\varphi(\theta_d)]$ induced by the increase in μ_s would imply an increase in the term in the square brackets for all relevant values of θ evaluated in the integral. Hence, with a fixed θ_d the LHS of equation (C.31) would increase while the RHS would remain fixed. This implies that θ_d cannot remain constant if μ_s increases.

Suppose now that θ_d falls as a reaction to the increase in μ_s. This would lead to a further increase in the value of the term in the square brackets for all relevant values of θ (as there would now be a larger gap between θ and θ_d). At the same time, a fall of θ_d would increase the range of values of θ that are integrated over. Hence, a decrease in θ_d would unambiguously increase the LHS of equation (C.31) while the RHS would remain fixed. This implies that θ_d cannot decrease either.

This proves that the only change in θ_d consistent with condition (C.31) is an increase in θ_d when μ_s increases. Therefore:

$$\frac{\partial \theta_d}{\partial \mu_s} \quad > \quad 0 \tag{C.32}$$

leading to the exit of the least productive firms.

Recall that productivity among operating firms follows a Pareto distribution with scale parameter θ_d. Thus, even if the range of productivity of operating firms decreases, TBTC increases the variance of productivity among operating firms. Given the dependence of firm wages on productivity, this results in an increase in the variance of firm wages and therefore wage inequality. In addition, since low-productivity firms have a lower share of abstract workers, their exit contributes to the increase in the economy's overall abstract share.

Prediction 2: *Differential Employment Growth* – TBTC strengthens the cross-sectional association between employment and productivity.

Proof: Taking the first- and second-order derivatives of (C.10), we obtain:

$$\begin{aligned}
\frac{\partial h_r(\theta)}{\partial \mu_s} &= \left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right) h_r(\theta)\,[1+\varphi(\theta)]^{-1}\,\frac{\nu}{\Lambda}\mu_s^{-1}\varphi(\theta) \quad > \quad 0 \\
\frac{\partial^2 h_r(\theta)}{\partial \mu_s \partial \theta} &= \left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right) h_r(\theta)\,[1+\varphi(\theta)]^{-2} \cdot \frac{\partial^2 \varphi(\theta)}{\partial \mu_s \partial \theta}\left[1+\varphi(\theta)\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)\right] \quad > \quad 0
\end{aligned}$$

Hence, TBTC increases routine employment for all firms, but more so for more productive firms. Similarly, taking the derivatives of (C.12):

$$\begin{aligned}
\frac{\partial h_s(\theta)}{\partial \mu_s} &= \frac{\nu}{\Lambda}\mu_s^{-1}\left(1-\frac{k}{\delta}\right)\left[1+\left(\frac{\beta-\nu}{\nu\Gamma}\right)\cdot\frac{\varphi(\theta)}{1+\varphi(\theta)}\right] h_s(\theta) \quad > \quad 0 \\
\frac{\partial^2 h_s(\theta)}{\partial \mu_s \partial \theta} &= \frac{\nu}{\Lambda}\mu_s^{-1}\left(1-\frac{k}{\delta}\right)\left\{\frac{\beta-\nu}{\nu\Gamma}\,[1+\varphi(\theta)]^{-2}\,\frac{\partial \varphi(\theta)}{\partial \theta} h_s(\theta) + \left[1+\left(\frac{\beta-\nu}{\nu\Gamma}\right)\cdot\frac{\varphi(\theta)}{1+\varphi(\theta)}\right]\frac{\partial h_s(\theta)}{\partial \theta}\right\} > 0
\end{aligned}$$

Hence, TBTC also increases abstract employment for all firms, but more so for more productive firms.

Prediction 3: *Increased Sorting* – TBTC increases the abstract employment share within all firms, and strengthens the cross-sectional association between productivity and abstract employment shares.
Proof: Taking the first-order derivative of (C.13) we get:

$$\frac{\partial}{\partial \mu_s}\left[\frac{h_s(\theta)}{h(\theta)}\right] = b_s\left(1-\frac{k}{\delta}\right)\frac{\nu}{\Lambda}\mu_s^{-1}\cdot\frac{1}{b_s+b_r\varphi(\theta)^{1-\frac{k}{\delta}}}\cdot\frac{h_s(\theta)}{h(\theta)} \quad > \quad 0$$

Hence, TBTC increases the share of abstract workers for all firms. Taking the second-order derivative of firm abstract employment share yields:

$$\frac{\partial^2}{\partial \mu_s \partial \theta}\left[\frac{h_s(\theta)}{h(\theta)}\right] = \frac{1}{\left[b_s+b_r\varphi(\theta)^{1-\frac{k}{\delta}}\right]^3}\left(1-\frac{k}{\delta}\right)\varphi(\theta)^{-\frac{k}{\delta}}\frac{\partial\varphi(\theta)}{\partial\theta}\left[b_s-b_r\varphi(\theta)^{1-\frac{k}{\delta}}\right]$$

Given that the ratio of abstract to routine workers is $h_s(\theta)/h_r(\theta) = \frac{b_r}{b_s}\varphi(\theta)^{1-\frac{k}{\delta}}$, the term $\left[b_s-b_r\varphi(\theta)^{1-\frac{k}{\delta}}\right]$ is positive if $h_s(\theta)/h_r(\theta) < 1$ and negative if $h_s(\theta)/h_r(\theta) > 1$. Therefore,

$$\frac{\partial^2}{\partial \mu_s \partial \theta}\left[\frac{h_s(\theta)}{h(\theta)}\right] \quad > \quad 0 \quad \text{if} \quad \frac{h_s(\theta)}{h_r(\theta)} < 1, \tag{C.33}$$

implying that the increase in the share of abstract workers will be larger for more productive firms if the number of abstract workers outweighs the number of routine workers at baseline. Starting from a situation with relatively more routine workers, TBTC will therefore increase sorting of abstract workers into firms with relatively higher productivity levels.

Prediction 4: *Differential Wage Growth* – TBTC strengthens the cross-sectional association between productivity and wages.
Proof: Taking the first- and second-order derivatives of (C.16) we obtain:

$$\begin{aligned}
\frac{\partial w_r(\theta)}{\partial \mu_s} &= \left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k\nu}{\delta\Lambda}\mu_s^{-1}w_r(\theta)\frac{\varphi(\theta)}{1+\varphi(\theta)} \quad > \quad 0 \\
\frac{\partial^2 w_r(\theta)}{\partial \mu_s \partial \theta} &= \left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k\nu}{\delta\Lambda}\mu_s^{-1}\frac{1}{1+\varphi(\theta)}\left[\varphi(\theta)\frac{\partial w_r(\theta)}{\partial\theta}+\frac{w_r(\theta)}{1+\varphi(\theta)}\cdot\frac{\partial\varphi(\theta)}{\partial\theta}\right] \quad > \quad 0
\end{aligned}$$

Hence, TBTC increases firm wages of routine workers, and more so for more productive firms. Similarly, taking the derivatives of (C.18):

$$\begin{aligned}\frac{\partial w_s(\theta)}{\partial \mu_s} &= \frac{k\nu}{\delta\Lambda}\mu_s^{-1} w_s(\theta)\left[1+\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{\varphi(\theta)}{1+\varphi(\theta)}\right] \quad > \quad 0 \\ \frac{\partial^2 w_s(\theta)}{\partial \mu_s \partial\theta} &= \frac{k\nu}{\delta\Lambda}\mu_s^{-1}\left\{\frac{\partial w_s(\theta)}{\partial\theta}\left[1+\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{\varphi(\theta)}{1+\varphi(\theta)}\right]+w_s(\theta)\left(\frac{\beta-\nu}{\nu\Gamma}\right)[1+\varphi(\theta)]^{-2}\frac{\partial\varphi(\theta)}{\partial\theta}\right\} \quad > \quad 0\end{aligned}$$

Thus, the increase in firm wages of abstract workers is disproportionately larger for more productive firms. Given the rise in firm abstract share and wages for both types of workers, TBTC also leads to an increase in firm average wages for all firms. In addition, if TBTC increases sorting, then $[\partial^2 w_\ell(\theta)/\partial\mu_s\partial\theta] > 0$, where $\ell = \{s, r\}$, unambiguously implies that the rise of firm average wages is monotonically increasing in firm productivity.

Figure A.1: Abstract Wage Premium

Note: The figure shows the evolution of the abstract wage premium overall, and holding the industry structure constant based on 1990 industry employment shares. The abstract premium is calculated as the difference between the average log wage of full-time abstract and routine workers separately for each industry and year, and then averaged across industries using the actual or 1990 industry structure. The mapping of detailed occupation codes to broad task categories is detailed in Appendix Table A.1.

Figure A.2: Within-Industry Between-Establishment Wage Inequality by Task Group

Panel A: Within-Industry Between-Establishment Wage Inequality: Routine Workers

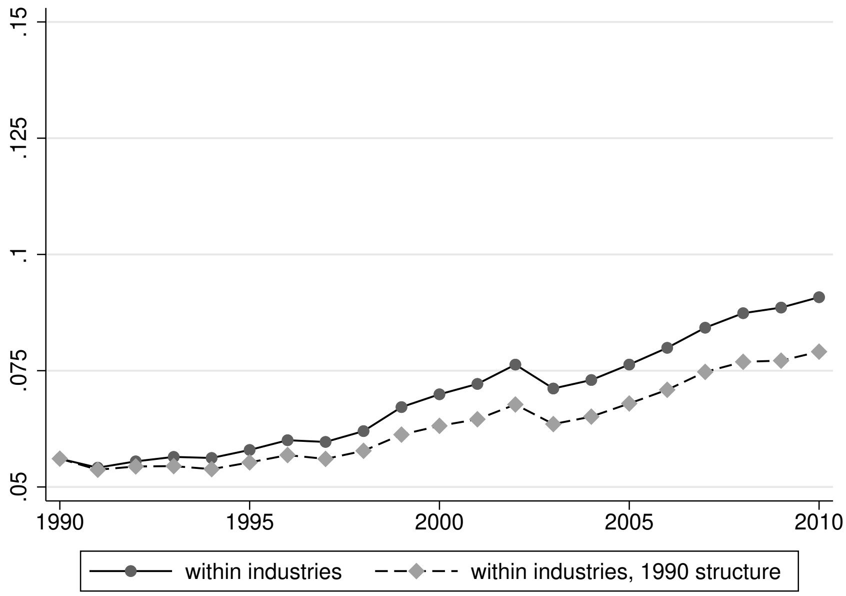

Panel B: Within-Industry Between-Establishment Wage Inequality: Abstract Workers

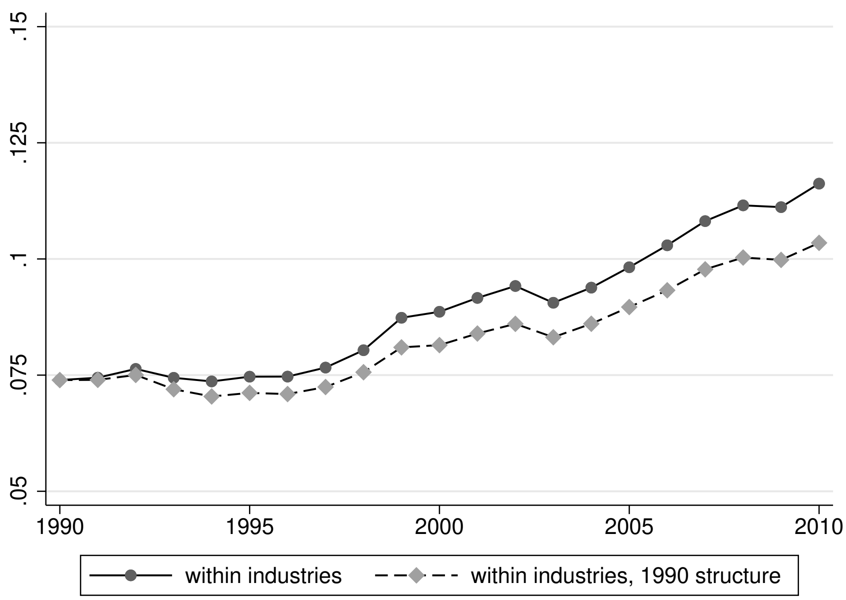

Note: The figure shows the evolution of the variance of average log wages of routine (Panel A) and abstract (Panel B) workers across establishments within industries. The solid line uses the contemporaneous industry structure in each year, while the dashed line fixes the industry composition using 1990 employment shares to average across industries. The mapping of detailed occupation codes to broad task categories is detailed in Appendix Table A.1.

Figure A.3: Decomposition of Changes in the Within-Industry Variance of Log Wages Across Establishments: 3-Digit Occupations

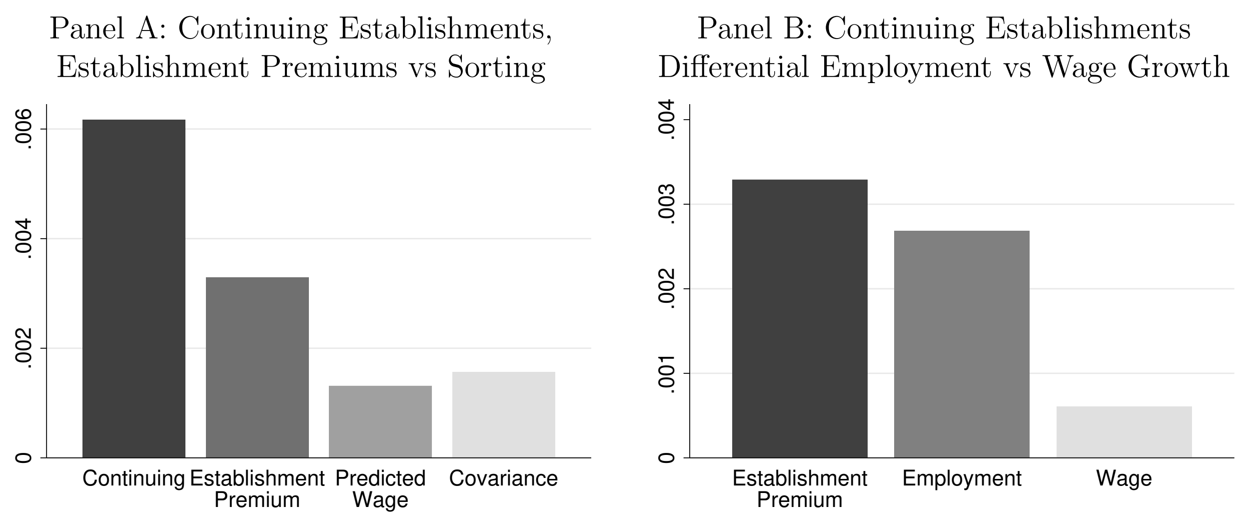

Note: Panel A decomposes changes in the within-industry variance of log wages among continuing establishments into components related to establishment premiums, dispersion in task usage (predicted wage) and sorting (covariance); see Equation (20). Panel B decomposes changes in the within-industry variance of establishment wage premiums among continuing establishments into a differential employment and wage growth component; see Equation (21). The decompositions are performed over non-overlapping 5-year windows within industries and then averaged across industries and time using 1990 employment shares as weights and giving equal weight to each time period.

Table A.1: Mapping of Occupation Codes to Task Groups

Task Group	Occupation Codes (KldB88)	Education Shares (at most high school / university)	Most Common Occupations
Abstract	303, 304, 600-635, 684, 751-763, 811-893	1.91% / 36.97%	nurses (8.68%), managers (7.11%)
Routine	71-302, 305-549, 681-683, 685-744, 771-805, 901-937	11.81% / 6.76%	office clerks (16.85%), salespersons (5.87%)
Excluded (Agric/ Forestry/ Unpaid)	11-62, 555, 666, 971-999		

Note: Education shares and shares of most common occupations (shown in brackets) weighted by employment status, with part-time employees accounting for half of full-time employees.