Technological Change, Firm Heterogeneity and Wage Inequality*

Guido Matias Cortes† Adrian Lerche‡

Uta Schönberg§ Jeanne Tschopp¶

July 26, 2021

Abstract

We argue that task-biased technological change not only affects wage gaps between workers in different tasks, but also increases wage inequality between workers in different firms, even when looking at individuals in the same task group. Building on a heterogeneous firm framework with labor market frictions, we show that an industry-wide task-biased technological change shock will increase between-firm wage inequality within an industry through four main channels: changes in the abstract wage premium (as in traditional models of technological change); 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. A simultaneous increase in the supply of abstract workers does not offset the technology-induced rise in inequality. Using rich administrative matched employer-employee data from Germany, 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 (Grant 430-2018-00686). 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.

†York University; gmcortes@yorku.ca.

‡Institute for Employment Research (IAB); adrian.lerche@iab.de.

§University College London, Centre for Research and Analysis on Migration (CReAM) and Institute for Employment Research (IAB); u.schoenberg@ucl.ac.uk.

¶University of Bern; jeanne.tschopp@vwi.unibe.ch.

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. The literature has argued that technology has asymmetric impacts across different groups of workers, depending on their skill levels or the tasks that they perform (e.g. Katz & Murphy, 1992; Autor et al., 2003, 2006; Goos et al., 2014; Jaimovich & Siu, 2020). According to this view, technology has impacted wage inequality by changing the demand for different skills and tasks, thus changing the employment structure of the economy and the relative wage returns for different groups.

Intriguingly, however, a recent parallel literature has shown that most of the rise in wage inequality can be traced back 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). This literature has argued that individual-level wages have become increasingly dependent on where people work, rather than the skills that they have or the tasks that they perform.

While the literature on between-firm inequality has documented many novel empirical facts, it has so far been largely silent on the driving forces behind these patterns. In comparison, while the literature on 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 provided limited insights into the growing wage differentials observed *within* groups, across firms.

In this paper, we argue that the development of task-biased automation technologies can account not only for increases in between-task inequality, but also for increases in inequality within-tasks, across workers in different workplaces, as observed in the data. We show this theoretically, using a rich yet tractable heterogeneous firm framework, and empirically, verifying the predictions of the model and quantifying the relative importance of the different channels that it highlights using administrative matched employer-employee data from Germany.

Our analysis draws on data from the so-called Beschäftigtenhistorik (BEH) from the Institute for Employment Research (IAB), covering 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.

In line with the existing literature, we document a substantial shift in employment away from routine occupations (such as office clerks and production workers) and towards abstract occupations (such as managers and engineers). In contrast, the abstract wage premium has remained roughly constant in Germany. The increase in the supply of abstract labor thus appears to have roughly offset the demand effects on the abstract wage premium coming from task-biased technological change. In consequence, essentially all of the increase in wage inequality observed in Germany between 1990 and 2010 is due to growing wage differentials among workers *within* the same broad task group.

In line with existing evidence, we also find that the rise in overall wage inequality is primarily due to growing wage differences between, rather than within establishments. We further show that, in any given year, about 60% of the overall variance of between-establishment log-wages is accounted for by pay differences across establishments within 3-digit industries. This finding indicates that we are far from having a 'representative establishment', in terms of pay, even within detailed industries. Moreover, this heterogeneity between establishments within industries has been growing over time and is an important driver of the overall rise in wage inequality: More than half of the rise in between-establishment wage inequality occurs across establishments operating within the same 3-digit industry. We focus on this growing heterogeneity in (log) wages across establishments within detailed industries in this paper. We further show that a major share of this heterogeneity is due to increased dispersion in the establishment premiums paid to workers conditional on their broad task or detailed occupations.

These patterns motivate us to explore the impact of task-biased technological change in an environment with two worker types and heterogeneous firms within industries, and where the supply of abstract workers is also changing. We do so by building on the major advances made in terms of the analysis of heterogeneous firm frameworks in the international trade literature (e.g. Melitz, 2003; Yeaple, 2005; Egger & Kreickemeier, 2009, 2012; Helpman et al., 2017; Trottner, 2019). In particular, we consider a version of the model in Helpman et al. (2010), which 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). The addition of these frictions in the Melitz (2003) model generates wage differentials between firms within industries for the same worker type in a rich, yet tractable way. The proposed model further enriches standard search and matching models as it features a clear concept of a firm, and allows routine and abstract workers to be imperfect substitutes within firms.[1] While the model has been used to study

[1]In standard search and matching models that allow for ex ante heterogeneity across workers and jobs, there is typically no natural definition of a firm, and worker types are assumed to be perfect substitutes in production.

the impact of trade liberalization, we show that this type of framework can also be very useful in terms of understanding the interplay between firm-level heterogeneity and technological change shocks.

We allow firms in the model to differ in terms of their overall productivity as well as their optimal task input mix. In equilibrium, more productive firms find it optimal to employ more workers of both types, have a higher abstract employment share, and pay higher wages to both types of workers. These cross-sectional patterns are consistent with what we observe in the data.

Our key innovations with regards to the theoretical framework are to introduce an aggregate task-biased technological shock in the spirit of Autor et al. (2003) within this rich heterogeneous firm setting, as well as considering the implications of such a shock if it is accompanied by an increase in the supply of abstract workers. We focus on the relative (rather than absolute) effects across workers and firms (e.g., the wage impacts for workers in more productive relative to less productive firms).[2]

In our setting, in spite of being an aggregate shock that is common across all firms, task-biased technological change (TBTC) leads to an increase in between-firm wage inequality. This occurs through four main channels. The first is an increase in the abstract wage premium (the wage of abstract workers relative to routine workers). This is the sole channel that arises in traditional models of task-biased technological change that feature a competitive labor market and a homogeneous representative firm. In this setting with heterogeneous firms, the change in the abstract wage premium not only increases overall wage inequality, but also wage inequality between firms, as firms with different productivity levels employ different shares of abstract and routine workers.

The remaining three channels are novel to our setting. The first new channel is differential employment growth, whereby the more productive, higher-paying firms are predicted to grow more. Differential employment growth implies that the cross-sectional association between productivity and size becomes stronger as a result of TBTC, and that employment concentration in the more productive firms within the industry increases. This contributes to an increase in worker-weighted measures of between-firm wage inequality. The second new channel is an increase in worker segregation by task, driven by increased sorting of abstract workers to high-productivity (and hence high-wage) firms. Increased worker segregation not only increases between-firm wage inequality, but also implies that the cross-sectional association between productivity and abstract employment shares is also strengthened by TBTC. Thirdly, the model generates endogenous within-firm wage changes, with more productive

[2]Some recent evidence on the absolute impacts of automation technologies are provided, for example, by Acemoglu & Restrepo (2020a,b).

firms disproportionately increasing the wage that they pay to workers of each task type. This implies a strengthening of the cross-sectional association between productivity and wages (overall and conditional on tasks), further contributing to the increase in between-firm wage inequality.

Our framework additionally implies that task-biased technological change will increase the productivity threshold for production, which results in the exit of low-productivity firms. This change in the composition of firms increases the variance of productivity among operating firms. Its impact on the change in the variance of (log) firm wages, however, is ambiguous.

Traditional models of skill- or task-biased technological change posit that there is an ongoing "race" between technological change – which raises the demand for skilled or abstract workers and leads to a rise in the skill premium and wage inequality, all else equal – and increases in educational attainment – which raise the supply of skilled or abstract workers and therefore dampen the effects of technological change (Tinbergen, 1974, 1975; Katz & Murphy, 1992; Goldin & Katz, 2008; Acemoglu & Autor, 2011). In our framework, an increase in the supply of abstract workers will dampen the effect of TBTC on the abstract wage premium, just as in the traditional models. However, the impact of the increased supply on the other channels highlighted by our model is more nuanced. In fact, we show that the increased supply will in general *amplify* (or at least not fully offset) the effects of TBTC on between-firm wage inequality. Hence, even if the increase in the supply of abstract workers fully offsets the impact of TBTC on the abstract wage premium, between-firm wage inequality may still grow due to differential changes in employment and wages across different types of firms.

Guided by the model, we return to the BEH and IABEP data and verify its key predictions. Consistent with the predictions from the model and the presence of task-biased technological change, we find that the within-industry establishment-level associations between productivity, employment, abstract shares and wages have become stronger over our sample period. For example, while a 1% increase in the establishment's labor productivity was associated with a 0.1% increase in establishment size (i.e., number of employees) in the early 1990s, this association quadrupled to 0.4% in 2010. Similarly, while a 1% increase in establishment size was associated with an increase of 0.06% in the wage establishments pay to workers of the same task type in the early 1990s, this association increased to 0.09% in 2010.

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, employing more abstract workers, and increasing the wages that they pay to workers of a given task type. Moreover, in line with recent evidence on increased employment concentration (e.g. Autor et al., 2020), establishments that are more productive, employ more abstract workers or pay higher wages at baseline experience higher employment growth than other establishments in the same industry, a pattern that is consistent with the implications of an aggregate task-biased technological change shock in the model.

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 over time). 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. In consequence, workplace segregation has increased such that workers of the same task type are increasingly clustered in the same workplaces, a pattern that is in line with the literature that has demonstrated that segregation of high-wage workers in high-wage firms is an important proximate reason behind the rise in between-firm wage inequality (Card et al., 2013; Song et al., 2019), and that we can rationalize as being driven by an aggregate technological change shock.

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 according to the model. 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 led to a small increase in between-establishment wage inequality; changes among continuing establishments, however, are quantitatively much more important. Second, we analyze the role of sorting along task dimensions. When considering 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. Workplace segregation, while present, plays only a minor role in the increase in between-establishment inequality, accounting for an additional 3% of the change.[3] The remainder of the change is due to an increase in the variance of the wage premiums that establishments pay conditional on worker task. Around half of the rise in this variance is due to differential employment growth across 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 on wage inequality has so far ignored.

[3] We also perform the decomposition based on detailed occupational categories, and find that, at that level of detail, sorting and segregation account for about 25% and 20% of the increase in within-industry between-establishment wage inequality among continuing establishments.

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 leveraging variation across industries in technology adoption, which we measure in three different ways: 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 (both overall and within task types), abstract share heterogeneity, and the sorting of abstract workers to high-wage establishments. This corroborates the importance of technology adoption (potentially exacerbated by increases in the supply of abstract workers) in driving the establishment-level patterns that we have documented.

Our findings make important contributions and connect to several strands of the literature. First, we provide an important innovation to the literature on skill- or task-biased technological change. Theoretically, this literature has considered representative firm frameworks with perfect competition, and has discussed the "race" between technological change and educational expansion (Katz & Murphy, 1992; Acemoglu & Autor, 2011). Empirically, this literature has leveraged variation in various measures of technology adoption across geographical areas (e.g. Machin & Van Reenen, 1998; Autor et al., 2015; Akerman et al., 2015; Dauth et al., 2021), industries (e.g. Michaels et al., 2014; Graetz & Michaels, 2018) or firms (e.g. Acemoglu et al., 2020; Lindner et al., 2021) to document a strong link between technology adoption and the employment share of high-skilled or abstract workers (as well as in some cases the skill premium) in that area, industry or firm.

By embedding a task-biased technological shock within a framework composed of heterogeneous firms with imperfect competition, we are able to show that task-biased technological change not only differentially 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 thus paint a much richer picture about the individual- and firm-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. Our framework and empirical analysis additionally demonstrates that an industry-wide technology shock has very different impacts on different firms in the industry, with some firms exiting the market and other firms expanding and increasing their wages. Moreover, an important implication of our framework is that there is no "race" between technology and education – an expansion in the supply of abstract workers may in fact exacerbate rather than dampen the effects of technological change on between-firm wage inequality.

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). This literature has been very successful in highlighting the increasing importance of firms for individual wages, and has documented increases in worker sorting and segregation (e.g. Kramarz et al., 1996; Cortes & Salvatori, 2019; Wilmers & Aeppli, 2021). It has, however, been largely silent so far on the underlying driving forces behind these patterns. We provide a tractable theoretical framework that allows us to study the interplay between task-biased technological change and these important workplace-level patterns at the industry level. Guided by the model, we document new empirical findings regarding the strengthening associations between various workplace-level outcomes, as well as quantifying the role of establishment entry and exit, the role of sorting and segregation along task dimensions, and the role of differential employment growth across establishments for wage inequality. Making use of various measures of technology adoption across industries, we show that the increase in between-establishment wage inequality, segregation and sorting was more pronounced in industries that experienced stronger technology adoption, further corroborating the importance of TBTC as a driver of wage inequality not only across workers, but also across workplaces.

An important strand of the labor literature has considered models that generate wage differentials across firms by assuming that workers have idiosyncratic tastes for different workplaces, thus giving employers wage-setting power (e.g. Bhaskar et al., 2002; Card et al., 2018). A recent example of this type of model is by Haanwinckel (2020), who sets up a parsimonious model that incorporates a rich production function with several task types and a minimum wage.[4] Models of this type abstract away from search and matching frictions in the labor market and thus assume that workers instantaneously receive wage offers from all firms in the labor market.[5] The model that we consider provides an alternative rich framework that yields closed-form solutions for key equilibrium outcomes, and is in the tradition of search and matching models, the workhorse models of the labor market.

Our paper also contributes 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

[4] A second recent example of this type of model is by Lindner et al. (2021), who show that firms which implement skill-biased technologies not only employ more skilled workers but also increase wages of skilled workers.

[5] These models further assume that each firm is too small to influence the wage setting behavior of other firms and thus abstract from strategic wage setting behavior of firms. An important exception is Berger et al. (2021), who develop a general equilibrium oligopsony model which considers how non-atomistic firms strategically compete for workers.

employment growth of the most productive workplaces within an industry. Technological change may therefore be at least partly responsible for the rise in employment 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.[6]

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, 2016 version).

We focus on developments after 1990 when wage inequality started to increase sharply in Germany across the entire distribution of wages (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 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 further 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

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

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 worker-level 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.

We follow Acemoglu & Autor (2011) to classify individuals into task groups based on their broad occupation codes.[7] 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).[8] 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.

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 (overall and by task). Our employment counts include part-time workers with a weight of 0.5. Since we do not observe hours worked,

[7] While we could group workers according to their education levels, we prefer, in the German context, to classify individuals according to their occupational affiliation. In Germany, nearly 70% of individuals in a school leaving cohort undergo apprenticeship and vocational training; in consequence, the share of university graduates is smaller than in many other developed countries, and abstract occupations tend to be performed by individuals with apprenticeship training as well as university graduates.

[8] While personal service occupations are typically considered as non-routine manual occupations which are somewhat sheltered from technological change, we group them here with routine occupations given that there is ample overlap in the skill composition across these occupation groups. Below we also present results that distinguish between the 317 detailed occupational categories that are available in the data, rather than distinguishing only between these two broad task groups.

our measures of establishment wages are based on full-time workers only.[9]

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.

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)'; see Appendix B.1 for details.

Since workers may substantially differ within broad task groups, we also 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 we use this in our analysis in order to rule out the possibility that the establishment wage premium computed based on the two broad tasks solely reflects differences in the occupational structure within task groups across establishments. 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

[9]Fitzenberger & Seidlitz (2020) provide evidence that a fraction of part-time workers are misclassified as full-time workers. Even though this affects inequality measures in a given year, the authors show that this misclassification is not driving the rise in inequality over time.

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).[10]

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 (WZ93). 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 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, Abstract Wage Premium and the Importance of Within-Task Inequality. A large literature has documented the decline of employment in routine occupations and the growth of employment in abstract occupations across many developed countries (see Goos et al., 2009; Acemoglu & Autor, 2011). Panel A of Figure 1 verifies this pattern for Germany, and shows that the aggregate employment share of abstract workers steadily rose from about 20% in 1990 to more than 26% in 2010 – a rise of 34% over two decades. This increase was in part driven by differential industry growth: industries which

[10] 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." We use the crosswalk provided by Dauth et al. (2021) to match industry codes in the robot data to industry codes in the BEH social security data.

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.2 for details), the employment share of abstract workers rose substantially by about 18%.

Whereas the abstract employment share strongly increased between 1990 and 2010, Panel B of Figure 1 shows that the abstract wage premium (i.e. the wage gap between abstract and routine workers; see Appendix B.3 for details) remained roughly constant over our sample period. This suggests that the rise in demand for abstract tasks driven by task-biased technological change was, in Germany at least, largely offset by an expansion in the supply of abstract workers (due e.g. to increased educational attainment or increased training). It also implies that there will be limited scope for changes in wage inequality *between* task groups to account for much of the rise in overall inequality.

We confirm this in Panel C of Figure 1. The black circles show that overall wage inequality in West Germany, measured as the variance of individual-level 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). The light grey diamonds, meanwhile, indicate that this is almost entirely a within-task phenomenon (see Appendix B.4 for details of the computation of within-task wage inequality). In other words, essentially all of the increase in wage 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 through the between-task component (as this is the only relevant dimension in a representative firm framework with a perfectly competitive labor market), this component – which is equal to the gap between the black circles and the light grey diamonds in the Figure – remained stable in Germany over this time period.

The grey triangles in Panel C of Figure 1 reflect the evolution of the within-*occupation* variance of log-wages (using 317 occupational 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 within-occupation variance is much lower than the within-task variance, wage differences within detailed occupational categories account for the majority of the overall log-wage variance in the cross-section, and account for more than half of the increase in the variance over time. Thus, there is substantial heterogeneity in wages across workers within the same detailed occupations, and wage inequality rose sharply also within these detailed occupational groups.

Wage Inequality 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 task group does not mean that TBTC is not an important driver of wage inequality in Germany. As we show below, in a setting that departs from the traditional representative firm framework with perfect competition in the labor market, TBTC can lead to an increase in inequality across workers in different workplaces (even if the supply of abstract workers keeps up with the rise in demand). Panel A of Figure 2 shows that the increase in overall wage inequality is nearly entirely driven by increasing wage differences between establishments – a pattern that 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). The figure decomposes the variance of individual log wages, denoted Var_t, into a within-establishment and a between-establishment component as follows:

$$
\begin{aligned}
Var_t &= \frac{1}{n_t}\sum_i (\ln w_{it} - \overline{\ln w}_t)^2 \\
&= \underbrace{\frac{1}{n_t}\sum_f \sum_{i \in i_{ft}} (\ln w_{it} - \overline{\ln w}_{ft})^2}_{\text{within establishments}} + \underbrace{\frac{1}{n_t}\sum_f n_{ft}(\overline{\ln w}_{ft} - \overline{\ln w}_t)^2}_{\text{between establishments } (Var_t^{BE})}, \qquad (1)
\end{aligned}
$$

where i denotes an individual and f indexes establishments. $\ln w_{it}$ is the log wage of individual i at time t, $\overline{\ln w}_t$ is the average log wage in period t, and $\overline{\ln w}_{ft}$ is the average log wage in establishment f in period t. n_t is the total number of workers and n_{ft} is the total number of workers at establishment f in year t (i_{ft} denotes this set of individuals). The figure highlights that the increase in the log-wage variance observed between 1990 and 2010 occurred almost entirely between establishments; increases in within-establishment wage differentials account for only 1.6% of the rise.

Panel B of Figure 2 explores the extent to which between-establishment wage differentials, as captured by Var_t^{BE} in Equation (1), are due to differences between establishments in the same 3-digit industry or due to differences between establishments in different industries:

$$
Var_t^{BE} = \underbrace{\frac{1}{n_t}\sum_k \sum_{f \in f_{kt}} n_{ft}(\overline{\ln w}_{ft} - \overline{\ln w}_{kt})^2}_{\text{between establishments, within industries}} + \underbrace{\frac{1}{n_t}\sum_k n_{kt}(\overline{\ln w}_{kt} - \overline{\ln w}_t)^2}_{\text{between establishments, between industries}}, \qquad (2)
$$

where k indexes industries, f_{kt} is the set of establishments in industry k in year t, n_{kt} is the total number of workers in industry k and year t, and $\overline{\ln w}_{kt}$ is the average log wage in industry k at time t.

While both of the components are important, within-industry differences account for more than half of the between establishment variance in the cross-section, and more than half of its change over time (compare the black dots and the light-grey diamonds in Panel B). These results indicate that there is heterogeneity in pay across establishments that goes well beyond the differentials predicted by their industry affiliation, implying that we are far from having a 'representative establishment', in terms of pay, even within detailed industries. Moreover, this heterogeneity between establishments within industries has been growing over time and is an important driver of the overall rise in wage inequality. It is these within-industry differences across establishments that are the focus of our paper.[11]

The within-industry increase in between-establishment wage inequality could in principle be driven by industries with higher between-establishment wage inequality growing at a faster rate than the average industry. To rule out this possibility, the mid-grey triangles in Panel B of Figure 2 display the counterfactual within-industry increase in the variance of (log) establishment wages holding the industry structure constant at its 1990 level; see Appendix B.5 for details. While the counterfactual increase in between-establishment within-industry wage inequality is slightly less pronounced than the actual increase – indicating that industries with above average within-industry between-establishment wage variances have grown in relative terms – there is a clear increase also in counterfactual inequality between 1990 and 2010 of about 40%.

Panel C of Figure 2 links the findings on the importance of wage differentials between establishments to the findings on the importance of wage differentials within task and occupation groups (as shown in Panel C of Figure 1). Specifically, the figure displays the within-industry between-establishment wage variance (based on the 1990 industry structure), as well as the variance of *establishment wage premiums*, which capture establishment-level differences in wages for workers within the same task (or occupation); see Appendix B.6 for details.[12] The figure clearly highlights that heterogeneity in establishment wage premiums are a major component of within-industry between-establishment inequality, and these premiums have become increasingly dispersed over time, regardless of whether they are computed conditional on workers' broad tasks (the light-grey diamonds in the figure) or on their detailed occupations (the mid-grey triangles).

[11] Recent work by Haltiwanger et al. (2021), in contrast, focuses on the importance of the between-industry component in accounting for the rise in between-firm wage inequality in the U.S. This study finds that in the U.S., 25% of the increase in between-firm wage inequality over the past three decades occurred within very detailed 4-digit industries.

[12] As mentioned above, establishment wage premiums are computed in each year as the average residual for each establishment from an individual-level log wage regression on a task indicator (or a full set of 3-digit occupation dummies) interacted with a full set of 3-digit industry fixed effects, and hence capture within-industry pay differentials across establishments conditional on worker composition.

Importance of Establishment Premiums Relative to Task or Occupational Premiums. To further gauge the relative importance of establishments versus tasks or occupations in determining individual workers' wages, Panel A of Figure 3 displays the difference between the 90th and 10th and the 80th and 20th percentile in within-industry establishment wage premiums (tasks) over time alongside the abstract wage premium, averaged across industries using 1990 industry employment shares. Whereas the 80th-20th and 90th-10th gaps in establishment 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 40 log points over the same period. Hence, the wage differentials between high- and low-premium establishments for workers in a given task are large and are becoming even larger over time.

Panel B of Figure 3 reveals a similar picture when we focus on establishment wage premiums that account for each establishment's occupation structure at a detailed level (and not only their broad task structure), and contrasts these to the observed wage differences between detailed occupational groups. Whereas the wage gap between the 10% lowest and 10% highest paying occupations increased by about 5 log points between 1990 and 2010, the wage gap between the 10% lowest and 10% highest paying establishments for workers in the same occupation increased by more than 10 log points over the same period. The increase is similarly larger if we focus on the 80-20 gap across establishments (conditional on occupational composition) and compare it to the 80-20 gap across occupations.

Overall, Figure 3 clearly illustrates that wages have become increasingly dependent on where workers work and (in relative terms) less dependent on what workers do.

Establishment Productivity, Task Usage and Wages. As a final piece of motivating evidence before setting up our theoretical framework, we explore the cross-sectional link between productivity, size, task usage and wages at the establishment level. Using sales data from the IABEP, Panel A of Table 1 explores the relationship between establishments' log productivity (total sales per full-time equivalent worker) and various establishment-level outcomes. This is analyzed by running a set of regressions which include fully interacted 3-digit industry and year fixed effects, so that identification is limited to cross-sectional 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.

Columns (1) and (2) of Panel A show that more productive establishments employ more workers of both routine and abstract task type, and hence are larger in terms of total employment.[13] The coefficient for abstract employment in Column (1) is larger than the one for routine employment in Column (2), suggesting that more productive establishments have a higher abstract employment share – a pattern that we verify directly in Column (3).

Column (4) shows that more productive establishments pay, on average, higher wages. In Column (5), we regress the establishment's wage premium (two tasks) 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 (4), which is in line with more productive establishments employing more abstract workers. Finally, Column (6) confirms that more productive establishments pay higher wages also to workers within the same detailed occupation group; hence, the reason why more productive establishments pay higher wages is partly, but not entirely due to the fact that they employ a higher share of workers in higher paying occupations.

Panel B of Table 1 uses log establishment size as the key regressor of interest. Column (1) confirms the positive and statistically significant relationship between establishment size and establishment productivity – in line with the evidence in Columns (1) and (2) of Panel A. The remaining columns of Panel B draw on the full BEH records. In Column (2) we show that larger establishments employ a higher share of abstract workers. Larger establishments also pay higher wages on average not only overall (Column (3)), but also to workers of the same task type (Column (4)) and to workers within the same detailed occupation (Column (5)).

Hence, there is an empirical link between establishment productivity, size and task usage that will motivate our model setup. Note that these relationships occur within 3-digit industries and are thus not accounted for by differences across industries in establishment sizes, wages or productivities.

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 within industries, with a production structure that distinguishes between two tasks, and where firm productivity and task usage are linked. We use the model in order to investigate the ways in which task-biased technological change may impact wage inequality

[13] For the analysis in these two columns, establishments with no workers of a given type are imputed to have one part-time worker (i.e., 0.5 full-time equivalent workers) of that type in order to be able to compute log employment.

between firms within industries.

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 differentials across firms for the same worker type within industries. Helpman et al. (2010) extend the Melitz (2003) model by deviating from the benchmark of perfect competition in the labor market, instead 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).

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).[14] Our key innovations relative to the analysis in Helpman et al. (2010) are in Sections 4.3 and 4.4. In Section 4.3 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. Given the evidence that the supply of abstract workers has also increased over time (see our discussion above), in Section 4.4 we consider whether the implications of task-biased technological change are mitigated or amplified by a simultaneous expansion in the supply of workers in these tasks.

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

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

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

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

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

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

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 $z > 2$.[15] Once firms observe θ, they decide whether to exit or produce. Production involves 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

$$y = \left[(\theta \mu_s \overline{a}_s h_s^\gamma)^\nu + (\mu_r \overline{a}_r h_r^\gamma)^\nu \right]^{1/\nu} , \tag{3}$$

where $0 < \nu < \beta$, and μ_s and μ_r are aggregate task-augmenting technology parameters.[16] 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

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

[16]The assumption that $\nu < \beta$ ensures that employment and wages of both types of workers are increasing in θ, in line with the empirical evidence presented above.

to be matched with n_ℓ workers, $\ell = \{s, r\}$.[17] Consistent with the empirical evidence, we assume that abstract workers are relatively scarce 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 $k > 1$.[18] 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 above or below an (endogenously chosen) cutoff $\widetilde{a}_\ell$, where $\ell = \{s, r\}$, $c > 0$, and $\delta > k$.[19]

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 sample (n_r and n_s), and what the match-specific ability cutoffs that they will screen to should be ($\widetilde{a}_r$ 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 choices. 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.

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

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

[19] The assumption that $\delta > k$ is also needed in order to ensure that employment and wages of both types of workers are increasing in θ.

4.2 Key Equilibrium Properties

Closed-form solutions can be obtained for the equilibrium values of firm-level employment, wages, revenue, 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.

Firm-Level Employment

Firm-level employment for routine workers is given by

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

where

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

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

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

The firm's abstract share is therefore

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

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

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

$$\frac{\partial h_r(\theta)}{\partial\theta} > 0, \qquad \frac{\partial h_s(\theta)}{\partial\theta} > 0, \qquad \frac{\partial h_s(\theta)/h(\theta)}{\partial\theta} > 0. \tag{4}$$

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. These predictions are in line with the motivating evidence presented in Table 1.

Firm-Level Wages

Firm-level wages for routine workers are

$$w_r(\theta) = w_{dr} \left[1 + \varphi(\theta)\right]^{\left(\frac{\beta\Lambda}{\nu\Gamma}-1\right)\frac{k}{\delta}}. \tag{5}$$

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

$$w_s(\theta) = \frac{b_s}{b_r}\varphi(\theta)^{k/\delta} w_r(\theta). \tag{6}$$

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

$$\frac{\partial w_r(\theta)}{\partial \theta} > 0, \qquad \frac{\partial w_s(\theta)}{\partial \theta} > 0. \tag{7}$$

The model therefore generates wage differences between firms, with more productive firms paying higher wages for workers in a given task. This is again in line with the motivating evidence presented in Table 1. Intuitively, the wage differentials arise in the model 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 (employment), but there would be no wage inequality between firms conditional on task.[20] 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.

Productivity Threshold

Finally, 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).

To summarize, the cross-sectional predictions of the model are that more productive firms are larger, have a higher abstract share, and 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. All of these relationships are in line with the motivating empirical evidence presented in Table 1.

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 (3). 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. Further note 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

[20] Felbermayr et al. (2011) is an example of such a framework.

and employment levels.

The key implications of an increase in μ_s are the following:[21]

Prediction 1: *Increased Abstract Wage Premium* – Task-biased technological change increases the abstract wage premium within all firms, and in the aggregate.

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

$$\frac{\partial \left[w_s(\theta)/w_r(\theta)\right]}{\partial \mu_s} > 0$$

Implications: As in traditional models with perfect competition and homogeneous firms, holding the supply of abstract workers constant, the rise in demand for abstract tasks induced by TBTC leads to a rise in overall inequality due to an increasing between-task component, i.e. a higher wage differential between abstract and routine workers. In our setting, given that more productive (higher wage) firms have a higher abstract share, increased abstract wage premiums also lead to higher between-firm inequality in average wages, all else equal.

Prediction 2: *Selection* – TBTC 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.[22] 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. This

[21] 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).

[22] 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. The Pareto assumption, however, is standard in the literature, and is supported by empirical evidence (see e.g. Axtell, 2001; Corcos et al., 2012).

increase in the variance of productivity may (but does not necessarily) lead to an increase in the variance of log wages among firms operating in the market, hence potentially contributing to the rise in between-firm inequality.

Prediction 3: *Differential Employment Growth* – TBTC strengthens the cross-sectional association between employment and productivity.

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

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

Implications: This prediction implies that more productive firms become disproportionately larger in terms of employment relative to less productive firms. TBTC therefore leads to increased employment concentration in more 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).

Prediction 4: *Increased Sorting and Segregation by Task* – TBTC strengthens the cross-sectional association between productivity and abstract employment shares, provided that firms employ relatively more routine than abstract workers at baseline (the empirically relevant case).[23]

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

$$\text{If} \quad \frac{h_s(\theta)}{h_r(\theta)} < 1, \;\; \text{then} \;\; \frac{\partial\left(\frac{\partial h_s(\theta)/h(\theta)}{\partial\theta}\right)}{\partial\mu_s} > 0.$$

Implications: This prediction implies that 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, resulting in more segregation of workers by task. Moreover, abstract (high-wage) workers will increasingly sort into more productive (and hence high-wage) firms. Both increased segregation and increased sorting will contribute to the overall increase in between-firm wage inequality.

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

Prediction 5: *Differential Wage Growth* – TBTC strengthens the cross-sectional association between productivity and wages conditional on worker type.

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

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

Implications: As a result of TBTC, wages for both types of workers disproportionately increase within more productive firms relative to less productive firms. Thus, firm wage premiums – wages firms pay to workers of the same task type – become more dispersed, leading 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. The first is an extension to the heterogeneous firm setting of the channel highlighted by traditional models of TBTC with competitive markets and no firm heterogeneity; that is, a rise in the wage of abstract workers relative to routine workers. The other three channels are novel to our setting; they are: differential employment growth, segregation and sorting, and differential within-firm wage growth. All of the channels compound each other in driving the increase in wage inequality between firms. Selective entry and exit of firms may additionally contribute to the rise in between-firm inequality.

4.4 Impacts of an Increase in the Supply of Abstract Workers

The evidence presented in Figure 1 – and in particular the fact that the abstract wage premium has remained stable over time in Germany – suggests that the rise in demand for abstract tasks due to TBTC has been accompanied by a rise in the supply of workers in these tasks. In this section we consider whether a rise in the supply of abstract workers fully mitigates the impacts of TBTC on inequality – as would be the case in a representative firm framework with competitive labor markets.

We model the rise in the supply of abstract workers as an exogenous fall in the abstract worker search cost. Intuitively, when abstract workers become more abundant, it becomes easier for firms to find workers of this type, therefore reducing this task-specific search cost.[24]

[24]Technically, the search cost is an endogenous variable which depends on labor market tightness. However, as Helpman et al. (2010) discuss, the equilibrium search cost is pinned down solely by expected income in

Proofs of all of the results discussed in this section are provided in Appendix C.4.

In the model, an increase in the supply of abstract workers mitigates the rise in the abstract wage premium induced by TBTC, as in traditional models. The implications regarding the other channels through which wage inequality may increase are, however, more subtle. First, the supply of abstract workers further increases the productivity threshold for production θ_d, thus amplifying the increase in the variance of productivity among operating firms induced by the technology shock, which in turn may contribute to the rise in between firm wage inequality.

Furthermore, the inequality-enhancing effects of technology that operate through differential employment growth and through increased sorting and segregation by task are also amplified if the technological change shock is accompanied by an increase in the supply of abstract workers. Intuitively, the reduced cost of hiring abstract workers disproportionately benefits more productive firms (which employ relatively more workers of this type). These more productive firms therefore expand (in relative terms) and also further increase their specialization in abstract tasks.

Finally, regarding the differential wage growth channel, an increase in the supply of abstract workers amplifies the prediction that between-firm wage inequality increases for routine workers, but (under reasonable parameter assumptions) dampens the prediction that between-firm wage inequality increases for abstract workers.

To summarize, a simultaneous expansion in the supply of abstract workers counteracts the rise in the relative wage of abstract workers as in traditional models of TBTC, but generally exacerbates the other channels. Hence, between-firm (and overall) wage inequality may rise even if TBTC occurs alongside an increase in the supply of abstract workers. In consequence, the notion of a 'race' between technology and the supply of skills, which is present in traditional models of technological change based on representative firms and perfectly competitive labor markets, is no longer present in this richer type of model.

5 Empirical Evidence

In this section we return to the BEH and IABEP data in order to test the various predictions of the model, and to decompose the relative empirical importance of the different channels that it highlights.

the outside sector, which is exogenous. We can think of an increase in the economy-wide supply of abstract workers as reducing the wage of abstract workers in the outside sector (all else equal) and hence reducing the cost of searching for an abstract worker for firms in the differentiated sector.

5.1 Associations over Time and Longitudinal Changes within Establishments

We begin by exploring the evolution over time of the cross-sectional associations between various establishment characteristics. The model predicts that, due to ongoing task-biased technological change and the differential changes that it induces for workplaces with different productivity levels, we should observe a strengthening of the cross-sectional relationship between establishment productivity and size (Prediction 3), abstract share (Prediction 4), and wage (Prediction 5). A simultaneous increase in the supply of abstract workers amplifies these predictions. To test whether this is indeed the case, we estimate the associations from Table 1 separately for each year, controlling for 3-digit industry fixed effects, thereby focusing once again on within-industry associations.

Figure 4 plots the coefficients from these yearly regressions, using data from the IABEP and (log) productivity as the 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, Panel B shows a strengthening of the relationship between productivity and abstract employment shares, which is indicative of increased sorting over time of abstract workers towards high-productivity establishments. Panel C shows that 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 5 confirms these findings drawing on the larger BEH 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 2 complements this evidence by showing estimates based on a set of regressions that consider changes within establishments over (non-overlapping) 5-year windows. Panel A regresses changes in various establishment outcomes (conditional on survival) on baseline establishment size, plus a set of fully interacted 3-digit industry and year fixed effects. The results show that surviving establishments that are larger (within their industry) 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 their industry (in terms of their productivity, abstract share and the wages they pay). If viewed through the lens of the model, TBTC therefore amplifies, rather than reduces, differences in productivity, task usage and pay across establishments within industries. A simultaneous expansion in the supply of abstract workers may have exacerbated these differences across establishments.

Panel B provides further evidence of differential employment growth across establishments, considering also establishments that exit the market. This panel shows the results of a set of regressions that use within-establishment percentage changes in employment over 5-year windows as the dependent variable, and link these changes to various baseline establishment characteristics (controlling for a set of fully interacted 3-digit industry and year fixed effects). Column (1) uses establishment productivity as the regressor of interest and confirms that establishments that are more productive at baseline (within their 3-digit industry) grow significantly more than less productive establishments in the industry over subsequent years. The remaining columns show that establishments with initially higher abstract shares and establishments that pay higher wages at baseline – overall and to workers of the same task and occupation type – also exhibit significantly larger employment growth. This evidence is consistent with the idea that TBTC (as well as a simultaneous increase in the supply of abstract workers) shifts employment toward more productive, higher wage establishments, as predicted by the model.

5.2 Segregation and Sorting

As shown above, larger establishments (which pay higher wages) have disproportionately increased their abstract employment shares. This suggests an increase in the sorting of abstract workers towards higher paying establishments, and an increase in worker clustering by task. We verify this in Figure 6.

Panel A plots the within-industry variance of establishments' abstract employment shares over time, averaged across industries using either the contemporaneous or the 1990 industry structure (see Appendix B.7 for details). 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. Put differently, segregation by task has increased across establishments within industries.

Panel B of Figure 6 shows the evolution of 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 Appendix B.8 for details). This co-variance also shows a clear positive trend over time: abstract (high-wage) workers increasingly sort into establishments that pay higher wage premiums.

This evidence, which our model rationalizes as being driven by task-biased technological change and potentially exacerbated by a simultaneous increase in the supply of abstract workers, 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.

5.3 Decompositions of Changes in Wage Inequality

In this section we perform a series of decompositions in order to assess the relative quantitative importance of the different channels highlighted by the model: changes in the abstract wage premium (Prediction 1; which according to Panel B of Figure 1 has remained largely stable and according to Panel C of Figure 1 appears to play a negligible role in the increase of wage inequality in Germany), selective firm exit (Prediction 2), differential employment growth (Prediction 3), sorting and segregation (Prediction 4), and differential wage growth conditional on worker type (Prediction 5).

In what follows we decompose the change in $Var_{kt}(\overline{\ln w_{ft}})$, the within-industry between-establishment variance in log wages (see Equation B.8), focusing on non-overlapping 5-year windows in our data.

Selection. 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, thereby leading to an increase in the variance of productivity among operating firms (as firm productivity is drawn from a Pareto distribution), and, potentially, to an increase in the variance of log wages (Prediction 2).

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. 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 of establishment wages among continuing establishments:

$$\Delta Var_{kt}(\overline{\ln w}_{ft}) = \underbrace{Var_{kt} - Var^{con}_{kt} + Var^{con}_{kt-5} - Var_{kt-5}}_{\text{selection}} + \underbrace{\Delta Var^{con}_{kt}}_{\text{continuing establishments}}, \tag{8}$$

where for clarity $Var_{kt}(\overline{\ln w}_{ft}) = Var_{kt}$ and Var^{con}_{kt} is the variance of establishment wages among continuing establishments in industry k. 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 productivity among continuing establishments exceeds the variance among all establishments, and hence the variance of establishment wages among continuing establishments (Var^{con}_{kt-5}) may be higher than 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 window in our data and average across industries using 1990 industry employment shares as weights. Panel A of Figure 7 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 driven by continuing establishments.

Sorting and Segregation. Next, we further decompose the change in the within-industry between-establishment variance among continuing establishments, ΔVar^{con}_{kt}, to determine the role of changes in worker sorting and segregation (Prediction 4).

To begin with, note that the establishment wage can be written as the sum of the establishment wage premium and the establishment predicted wage, i.e.: $\overline{\ln w}_{ft} = \widetilde{\ln w}_{ft} + \widehat{\ln w}_{ft}$, where $\widetilde{\ln w}_{ft}$ is the establishment wage premium and $\widehat{\ln w}_{ft}$ is the establishment predicted wage, which in turn is a function of the establishment's task input mix and the abstract wage premium in the industry. Given this relationship, we can decompose the change in the wage variance among continuing firms in industry k as follows:

$$\Delta Var^{con}_{kt}(\overline{\ln w}_{ft}) = \underbrace{\Delta Var^{con}_{kt}(\widetilde{\ln w}_{ft})}_{\text{within-task}} + \underbrace{\Delta Var^{con}_{kt}(\widehat{\ln w}_{ft})}_{\text{segregation}} + \underbrace{2\Delta Cov^{con}_{kt}(\widetilde{\ln w}_{ft}, \widehat{\ln w}_{ft})}_{\text{sorting}}. \tag{9}$$

The first component, $\Delta Var^{con}_{kt}(\widetilde{\ln w}_{ft})$, captures changes in the variance of the wage premiums paid by continuing establishments conditional on worker task. We displayed the evolution of this variance, averaged across industries using the 1990 industry structure, in Panel C of Figure 2 (the light-grey diamonds), with the modification that in the figure we considered all and not only continuing establishments.

The second component, the change in the variance of predicted establishment wages $\Delta Var_{kt}^{con}(\widehat{\ln w}_{ft})$, captures both changes in the abstract wage premium and changes in the segregation of worker types across establishments. 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 ($\overline{\ln w}_{kt}^{r}$), plus the share of abstract workers in the establishment (S_{ft}) multiplied by the industry-year-specific abstract wage premium ($AbPrem_{kt}$); i.e. $\widehat{\ln w}_{ft} = \overline{\ln w}_{kt}^{r} + S_{ft} \cdot AbPrem_{kt}$. Thus, the within-industry variance in predicted establishment wages among continuing establishments equals:

$$\begin{aligned} Var_{kt}^{con}(\widehat{\ln w}_{ft}) &= Var_{kt}^{con}(\overline{\ln w}_{kt}^{r} + S_{ft} \cdot AbPrem_{kt}) \\ &= AbPrem_{kt}^{2} \cdot Var_{kt}^{con}(S_{ft}). \end{aligned}$$

Since the abstract wage premium remained roughly constant over time (Panel B of Figure 1), whereas the variance of establishments' abstract employment shares increased over time (Panel A of Figure 6), we expect any changes in the within-industry variance of predicted establishment wages to primarily reflect changes in the variance of establishments' abstract employment shares. This observation motivates us to refer to this term as changes in worker segregation by task.

In turn, the third component in Equation (9), $\Delta Cov_{kt}^{con}(\widetilde{\ln w}_{ft}, \widehat{\ln w}_{ft})$, captures the increased sorting of abstract workers into establishments paying higher wage premiums. To see this, note that the co-variance between establishments' predicted wages and their wage premiums equals:

$$\begin{aligned} Cov_{kt}^{con}(\widehat{\ln w}_{ft}, \widetilde{\ln w}_{ft}) &= Cov_{kt}^{con}(\overline{\ln w}_{kt}^{r} + S_{ft} \cdot AbPrem_{kt}, \widetilde{\ln w}_{ft}) \\ &= AbPrem_{kt} \cdot Cov_{kt}^{con}(S_{ft}, \widetilde{\ln w}_{ft}). \end{aligned}$$

Based again on the finding of a stable abstract wage premium, changes in this term will be primarily driven by changes in the co-variance between establishment abstract shares and establishment premiums, i.e. by the increased sorting of abstract workers to establishments that pay higher wage premiums.

The results of this decomposition are presented in Panel B of Figure 7, 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 9) can account for about 12% of the overall increase in within-industry between-establishment wage inequality among continuing establishments. While dispersion in task usage or segregation has increased over time across establishments within the same industry (see Panel A of Figure 6), its contribution to the overall increase in between-establishment wage inequality is minor (the second component in Equation 9). Not surprisingly, given that we only distinguish between two task groups, the within-task component (i.e., changes in the variance of establishment wage premiums) accounts for the majority (86.1%) of the change in the within-industry variance among continuing establishments.

In Appendix Figure A.1, we repeat the exercise distinguishing between 317 occupations, rather than two tasks. As expected, increased segregation (or dispersion in the occupational structure) across establishments (i.e., the second component in Equation 9) and increased sorting of workers in high-paying occupations into establishments paying high establishment premiums (i.e., the third component in Equation 9) become quantitatively more important, accounting for about 20% and 25% of the overall increase in the within-industry wage variance among continuing establishments, respectively. The change in the variance of establishment wage premiums, however, remains the dominant component also when considering this detailed occupational level.

Differential Employment Growth vs Differential Wage Growth. In a final step, our goal is to gauge the importance of differential employment growth for the increase in between-establishment wage inequality. In line with the prediction of the model, we have documented that the association between establishment productivity and size has become stronger over time (Panel A of Figures 4 and 5) and that establishments that were initially more productive (and those that paid higher wages) grow at faster rates (Panel B of Table 2). These differential employment growth rates will result in an increase in wage inequality across establishments, even if establishment wage premiums had remained unchanged.

To assess the importance of this channel, we focus on the change in the variance of the establishment wage premiums, $\Delta Var_{kt}^{con}(\widetilde{\ln w_{ft}})$, and decompose this variance into the role that is due to differential employment growth (Prediction 3), and a residual component that is attributable to differential changes in wage premiums within establishments (Prediction 5). This final decomposition is given by:

$$\Delta Var_{kt}^{con}(\widetilde{\ln w_{ft}}) =$$
$$\underbrace{\sum_{f \in f_{kt}^{con}} \left[\Delta\left(\frac{n_{ft}}{n_{kt}^{con}}\right)\right]\left(\widetilde{\ln w_{ft-5}} - \overline{\widetilde{\ln w_{kt-5}}}^{con}\right)^2}_{\text{differential employment growth}} + \underbrace{\sum_{f \in f_{kt}^{con}} \frac{n_{ft}}{n_{kt}^{con}}\Delta\left(\widetilde{\ln w_{ft}} - \overline{\widetilde{\ln w_{kt}}}^{con}\right)^2}_{\text{residual (differential wage growth)}}, \quad (10)$$

where f_{kt}^{con} is the set of continuing firms in industry k at time t, n_{kt}^{con} denotes total employment among continuing establishments in industry k and time t and $\overline{\widetilde{\ln w_{kt}}}^{con}$ denotes the average wage premium in industry k in period t among continuing establishments. 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 7 shows the results from this decomposition. The figure shows that differential employment growth is an important driver behind the increase in the variance of establishment wage premiums among continuing establishments, accounting for around half of the increase. Thus, the fact that establishments grow at different rates accounts for an important share of the increase in wage inequality, even abstracting from any changes in wages within establishments. The residual component is also important, indicating that the increased dispersion in establishment wage premiums is also quantitatively relevant. Panel B of Appendix Figure A.1 presents the analogous results based on the establishment wage premiums that control for detailed occupations, rather than two tasks, and also shows a very important role for differential employment growth.

Overall, we can conclude that all of the channels highlighted by the model have contributed to the rise in wage inequality between establishments within industries. While changes in the abstract wage premium (the channel that emerges from traditional models of task-biased technological change) and changes in segregation play only a minor role, changes in the composition of operating establishments, sorting of abstract workers to high-wage establishments, differential employment growth, and differential within-establishment wage growth are all quantitatively important, with the latter two being of primary importance.

5.4 Technology Adoption: Industry-Level Analysis

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

We first consider industry-level variation in the change of the abstract worker share between 1990 and 2010. If we think of changes in the supply of abstract workers as an aggregate common shock impacting all industries, we can interpret differential changes in industry-level abstract employment shares as being driven by differential exposure to task-biased technological change. 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.[25]

Panel A of Figure 8 shows the evolution over time of the overall variance of establishment (log) wages for these two groups of industries. In line with the theoretical model, we find that industries that experience larger increases in their abstract employment shares 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, establishments have become increasingly heterogeneous in their task mix 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 larger increases in sorting of abstract workers to high wage premium establishments in industries characterized by larger increases in their abstract employment shares.

Figure 9 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 (log) wages and establishment wage premiums that adjust for the task and occupation structure in the establishment. The remaining panels show that establishments are becoming increasingly heterogeneous in terms of their task mix particularly in industries with above-median robot adoption (Panel E), and that the sorting of abstract workers into high-wage establishments is also particularly pronounced in these industries (Panel F).

Finally, Figure 10 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

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

1991 and 2007 from the EUKLEMS data. Industries with more technology adoption tend to experience larger increases in their abstract shares (Panel A), larger increases in the dispersion of average establishment (log) wages (Panel B) and establishment wage premiums (Panels C and D), larger increases in the dispersion of establishments' task usage (Panel E), and more sorting of abstract workers towards high wage establishments (Panel F).

To summarize, in line with the predictions from the theoretical model, patterns of overall between-establishment wage inequality, segregation, sorting, and dispersion in establishment wage premiums appear to be stronger in industries more affected by TBTC, providing additional supporting evidence for the hypothesis that TBTC is an important driver of increased wage inequality between 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 representative firm frameworks with perfectly competitive labor markets and has hence had implications solely in terms of wage differentials *between* workers in different task groups. Empirically, however, a major component of the increase in wage inequality is observed *within* task groups, across establishments within industries.

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, thereby generating a rise in between-firm wage inequality. Using detailed administrative social security data from Germany, we document a number of novel empirical patterns at the establishment level, and show that these patterns are in line with the predictions of the model. The model highlights that the rise in inequality occurs due to endogenous changes in worker sorting, establishment size and establishment wages paid to the same worker type, as well as possibly endogenous changes in the composition of operating firms. We find that all of these channels are empirically relevant. 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.

The literature on technological change has long thought about increases in educational attainment as being a useful tool 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 (if large enough) 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 is no longer valid. We show that even though an increase in educational attainment mitigates the rise in the abstract wage premium (as is indeed observed in the German case over our sample period), it actually compounds the effects of technological change on wage inequality that operate through some of the other channels highlighted by the model. Hence, expanding educational attainment may no longer be sufficient to dampen the rise in wage inequality induced by task-biased technological change.

Overall, our results point to 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 (e.g. Cortes, 2016; Blien et al., 2021), our findings indicate that the type of firm that individuals are matched to is at least as important: 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 firms. Understanding what type of policies can mitigate the negative impacts of technological change on some groups of workers, within the context of a more realistic environment with heterogeneous firms and various market frictions such as the ones considered in this paper, remains a crucial direction 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., Lelarge, C., & Restrepo, P. (2020). Competing with robots: Firm-level evidence from France. *AEA Papers and Proceedings*, 110, 383–88.

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. *AEA Papers and Proceedings*, 110, 356–61.

Akerman, A., Gaarder, I., & Mogstad, M. (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics*, 130(4), 1781–1824.

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., Dorn, D., & Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584), 621–646.

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.

Axtell, R. L. (2001). Zipf distribution of U.S. firm sizes. *Science*, 293(5536), 1818–1820.

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.

Berger, D. W., Herkenhoff, K. F., & Mongey, S. (2021). Labor market power. *NBER Working Paper No. 25719*.

Bhaskar, V., Manning, A., & To, T. (2002). Oligopsony and monopsonistic competition in labor markets. *The Journal of Economic Perspectives*, 16(2), 155–174.

Blien, U., Dauth, W., & Roth, D. H. (2021). Occupational routine intensity and the costs of job loss: Evidence from mass layoffs. *Labour Economics*, 68, 101953.

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1), S13–S70.

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.

Corcos, G., Gatto, M. D., Mion, G., & Ottaviano, G. I. (2012). Productivity and firm selection: Quantifying the 'new' gains from trade. *The Economic Journal*, 122, 754–798.

Cortes, G. M. (2016). Where have the middle-wage workers gone? A study of polarization using panel data. *Journal of Labor Economics*, 34(1), 63–105.

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*.

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 10.1093/jeea/jvab012.

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.

Fitzenberger, B. & Seidlitz, A. (2020). The 2011 break in the part-time indicator and the evolution of wage inequality in Germany. *Journal for Labour Market Research*, 54(1), 807–836.

Goldin, C. & Katz, L. F. (2008). Long-run changes in the wage structure: Narrowing, widening, polarizing. *Brookings Papers on Economic Activity*, 2007(2), 135–165.

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.

Haanwinckel, D. (2020). Supply, demand, institutions, and firms: A theory of labor market sorting and the wage distribution. *Working Paper*.

Haltiwanger, J., Hyatt, H., & Spletzer, J. R. (2021). Industries, mega firms, and increasing inequality. *NBER Wage Dynamics in the 21st Century Conference Working Paper*.

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.

Kramarz, F., Lollivier, S., & Pele, L.-P. (1996). Wage inequalities and firm-specific compensation policies in France. *Annales d'Economie et de Statistique*, (pp. 369–386).

Lindner, A., Muraközy, B., Reizer, B., & Schreiner, R. (2021). Firm-level technological change and skill demand. *Working Paper*.

Machin, S. & Van Reenen, J. (1998). Technology and changes in skill structure: Evidence from seven OECD countries. *The Quarterly Journal of Economics*, 113(4), 1215–1244.

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*, 86(1), 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: Evolution of Task-Specific Labor Market Outcomes

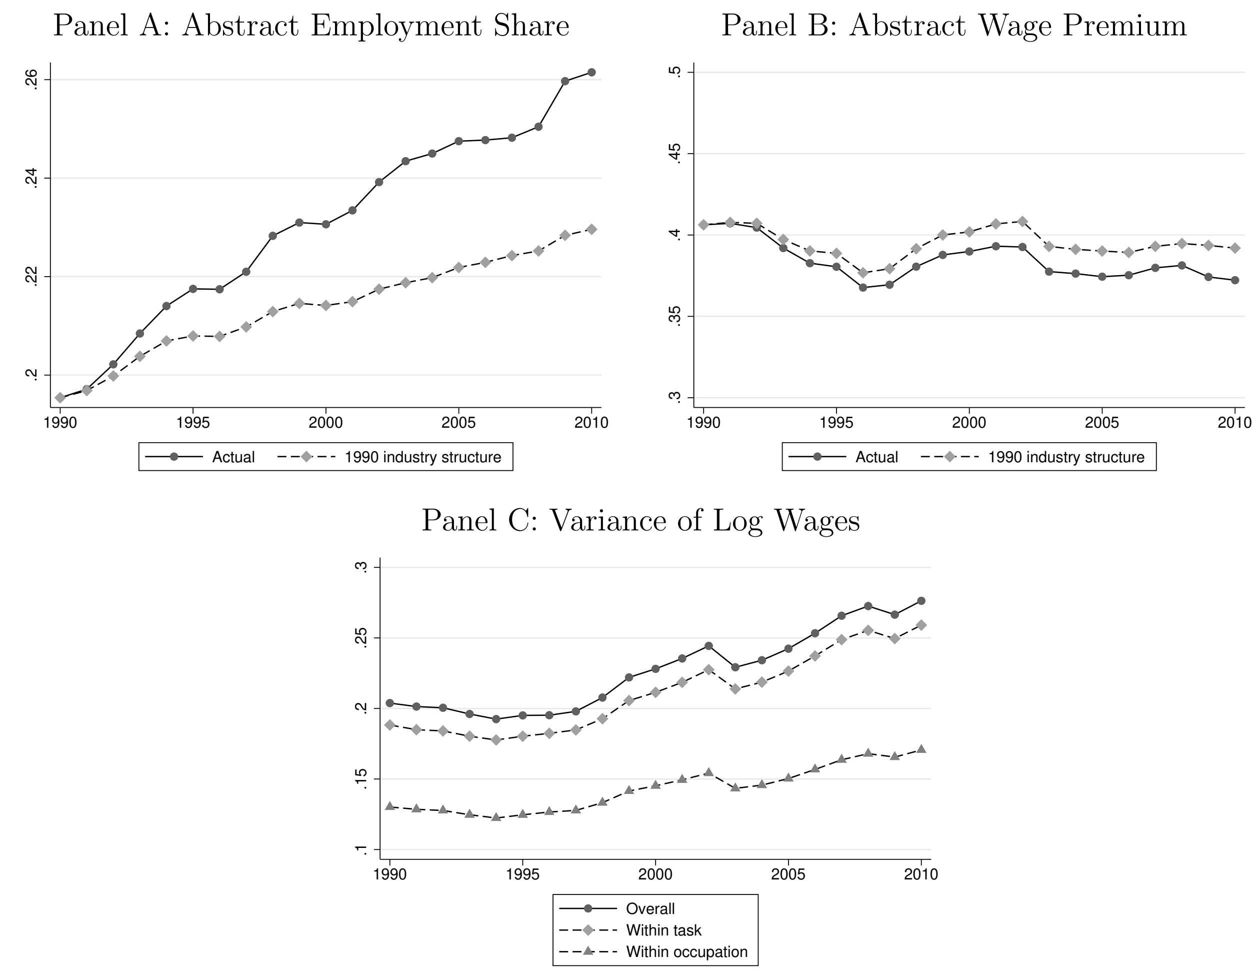

Note: Panel A shows the evolution of the share of aggregate employment in abstract occupations in Germany between 1990 and 2010 based on data from the Beschäftigtenhistorik (BEH). The solid line uses the contemporaneous industry structure in each year, while the dashed line fixes the industry composition using 1990 employment shares and therefore captures only changes in the task composition of employment within industries (see Appendix B.2). Panel B shows the evolution of the abstract wage premium, 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 (see Appendix B.3). Panel C displays the evolution of the overall, within-task (based on two broad task groups) and within-occupation (based on 317 occupations) variance of individual log wages (see Appendix B.4). The mapping of detailed occupation codes to broad task categories is presented in Appendix Table A.1.

Figure 2: Evolution of the Variance of Log Wages

Note: Panel A displays the evolution of the overall variance of individual log wages and its within- and between-establishment components, as described in Equation (1). Panel B displays the overall and the within-industry between-establishment wage variance, computed using the actual and 1990 industry structure, as shown in Equation (2) and Appendix B.5. Panel C displays the within-industry variance in establishment wage premiums based on either two broad tasks or 317 detailed occupation groups, averaged across industries using the 1990 industry structure (see Appendix B.6). Establishment wage premiums adjust for the task and occupation 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.).

Figure 3: Establishment vs Task and Occupation Premiums

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

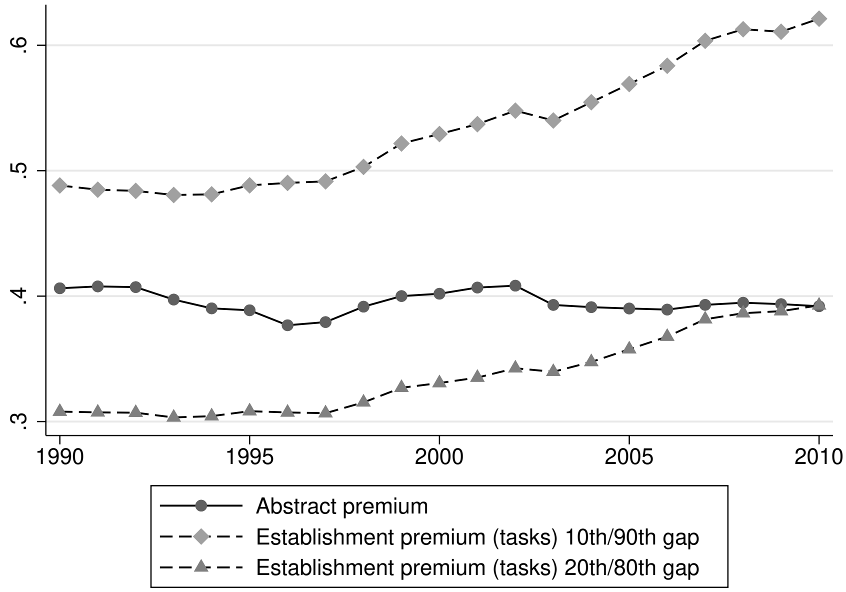

Panel B: Establishment Premium Gaps (317 Occupations) and Occupation Wage Premiums

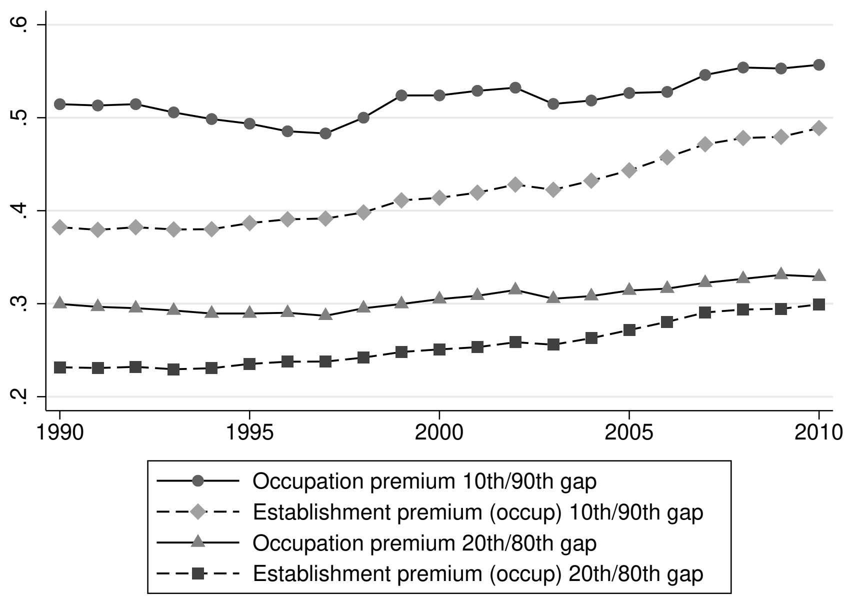

Note: The figure compares the within-industry gaps in wages for workers in different tasks (occupations), relative to the gaps in wages for workers in the same task (occupation) but in different establishments. Panel A shows the evolution of the within-industry abstract wage premium and of the 90-10 and 80-20 percentile gap in establishment wage premiums (broad tasks). Panel B displays the within-industry evolution of the 90-10 and 80-20 percentile gaps in establishment wage premiums (317 occupations) and in occupational wages. 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. Within-industry gaps are averaged across industries using 1990 industry shares.

Figure 4: Year-by-Year Associations between Establishment Productivity and Other Establishment Characteristics (Within Industries)

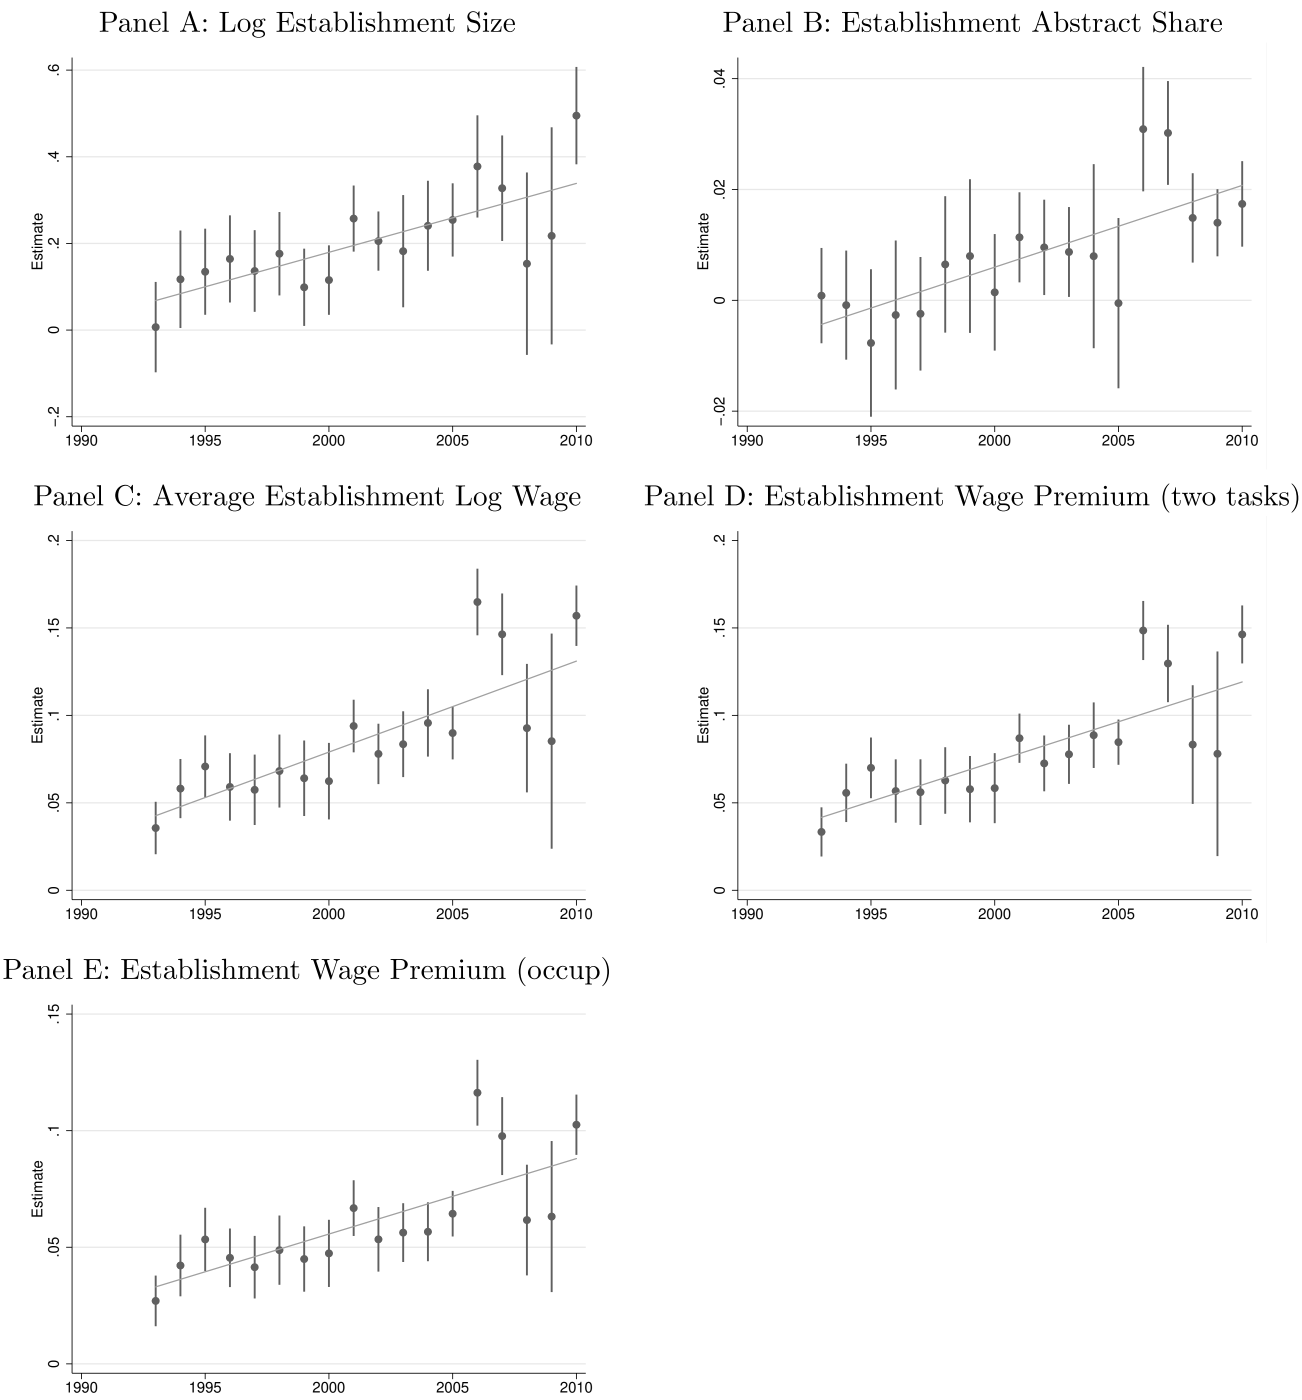

Note: The figure shows estimated coefficients and 95% confidence intervals from regressions of the outcome that appears in the title of each panel on log establishment productivity and a full set of 3-digit industry fixed effects, with the regressions estimated separately for each year. Results are based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

Figure 5: Year-by-Year Associations between Establishment Size and Other Establishment Characteristics (Within Industries)

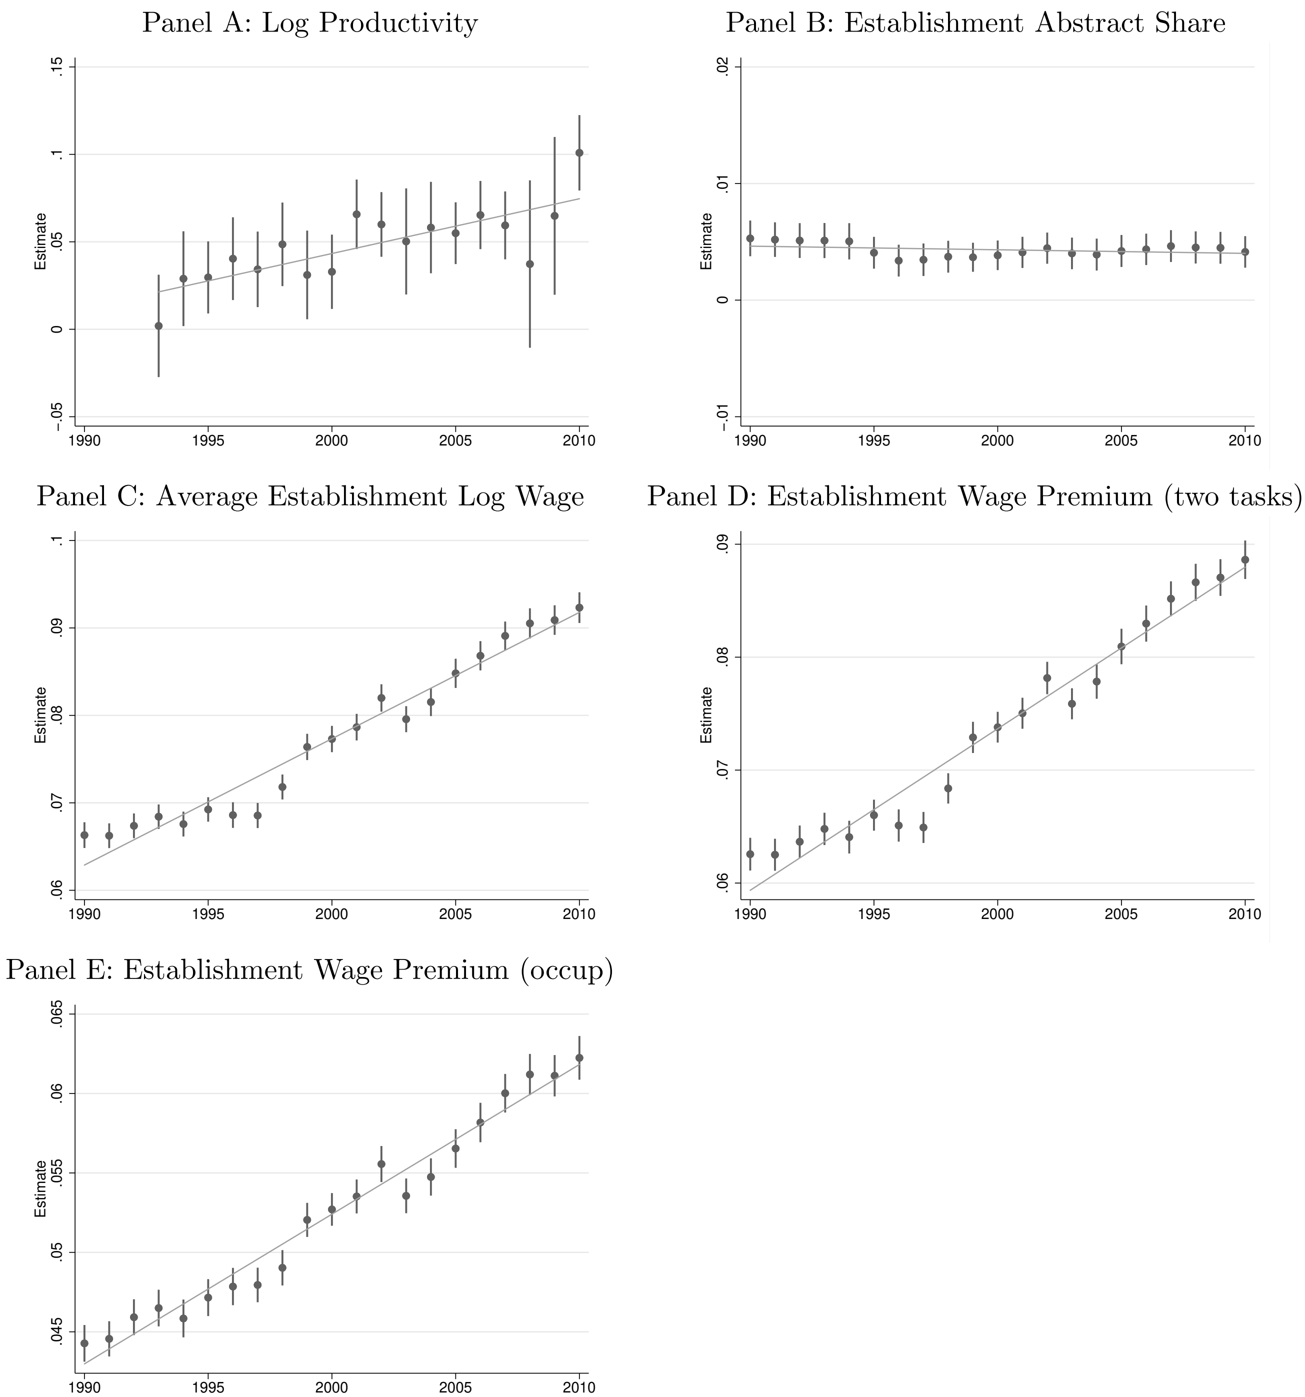

Note: The figure shows estimated coefficients and 95% confidence intervals from regressions of the outcome that appears in the title of each panel on log establishment size and a full set of 3-digit industry fixed effects, with the regressions estimated separately for each year. Results are based on 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 6: Abstract Share Heterogeneity and Sorting (Within Industries)

Panel A: Variance of Abstract Employment Shares

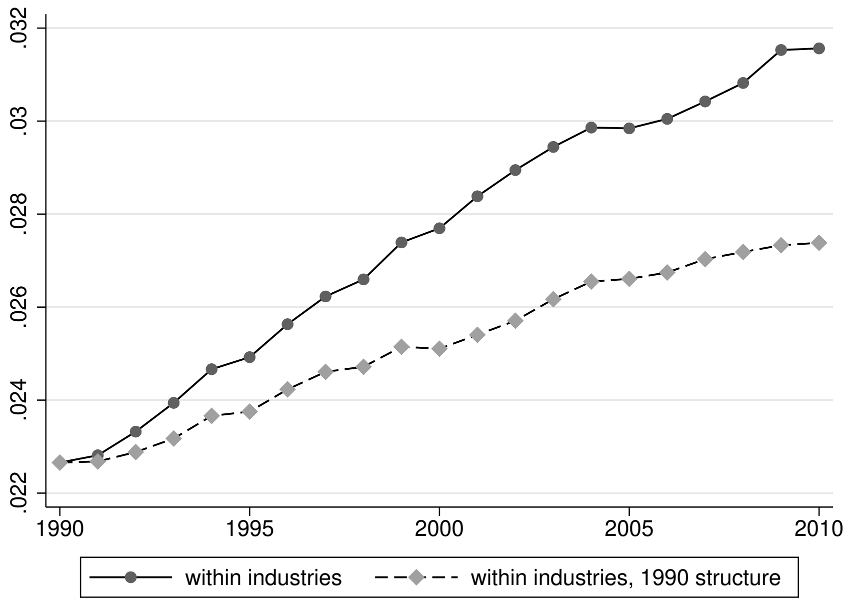

Panel B: Co-variance between Abstract Share and Establishment Wage Premium

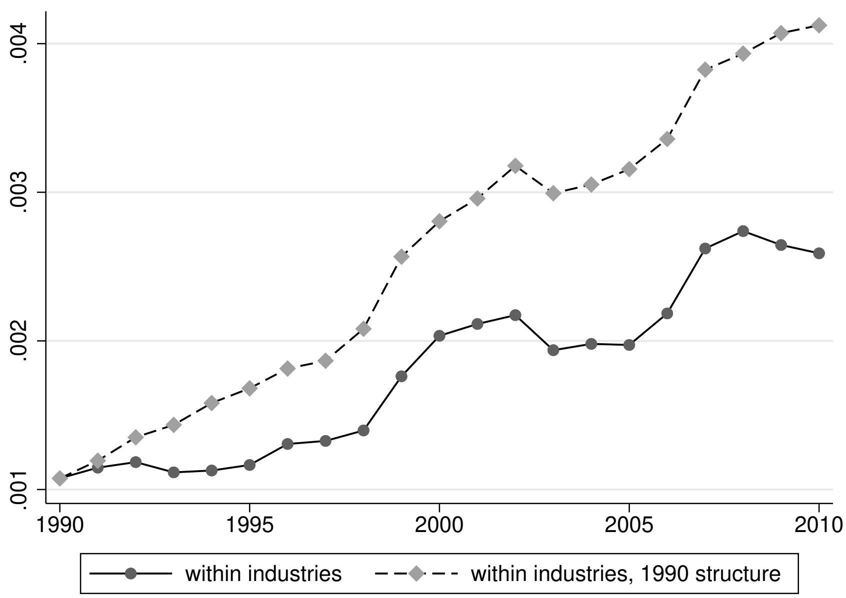

Note: Panel A shows the evolution of the variance of the abstract employment share across establishments within industries, averaging across industries using either observed industry employment shares in each year (solid line) or constant 1990 industry employment shares in all periods (dashed line); see Appendix B.7. Panel B shows the co-variance between establishments' abstract employment shares and their wage premium (two tasks); see Appendix B.8.

Figure 7: Decomposition of Changes in the Within-Industry Between-Establishment Variance of Log Wages

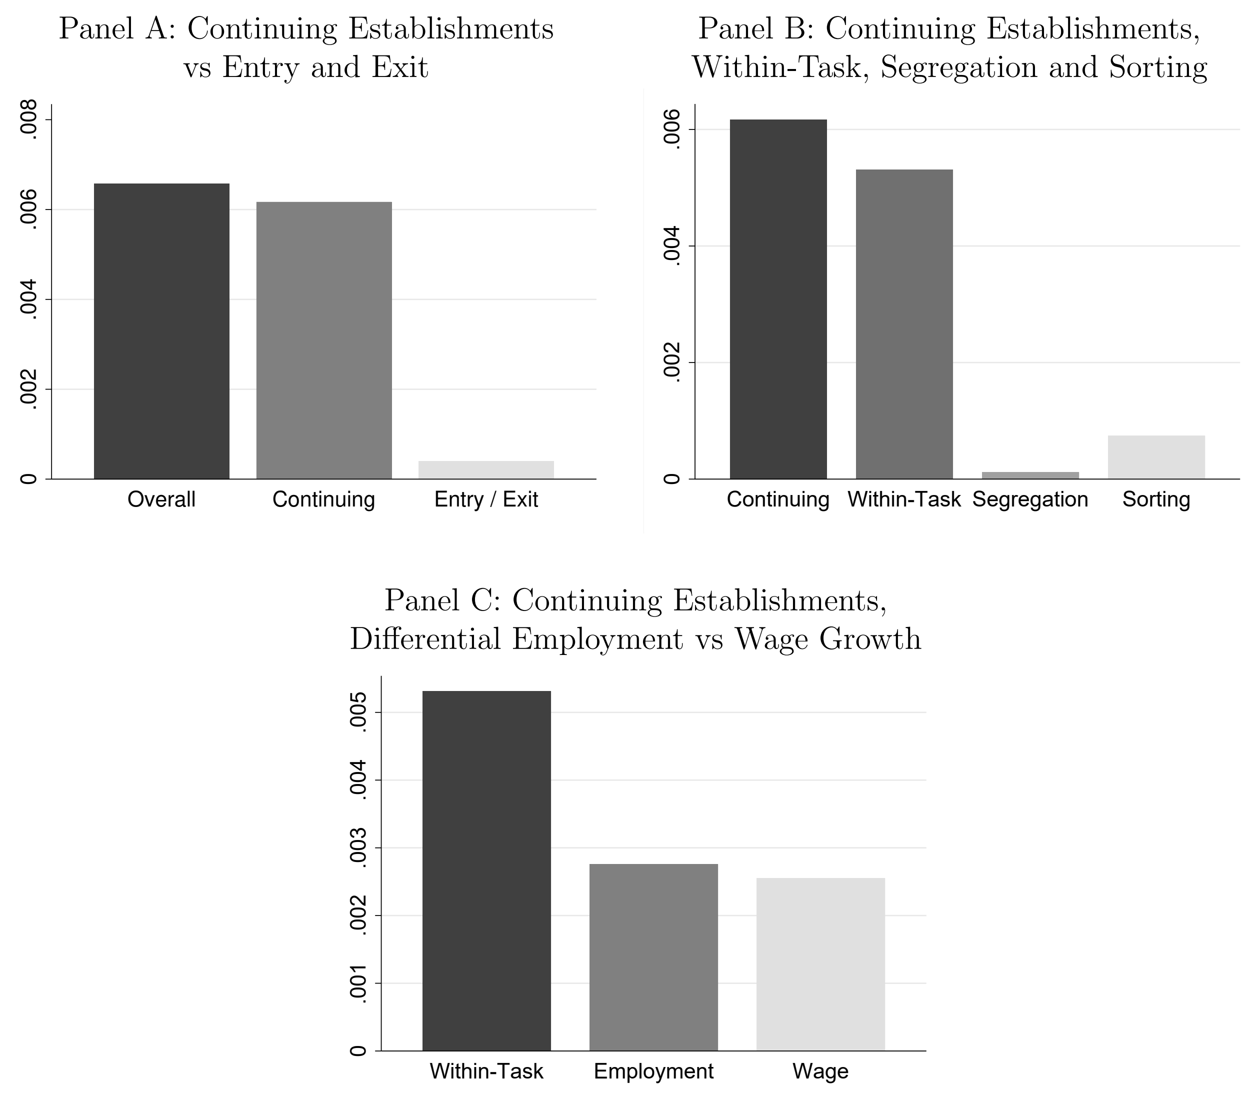

Note: Panel A decomposes changes in the within-industry, between-establishment variance of log wages into two components: selection due to establishment entry and exit, and changes among continuing establishments; see Equation (8). Panel B decomposes changes in the within-industry variance of log wages among continuing establishments into within-task, segregation and sorting components; see Equation (9). Panel C decomposes changes in the within-task component among continuing establishments into differential employment and wage growth components; see Equation (10). 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 8: Industries with Below vs Above Median Increases in Abstract Employment Shares

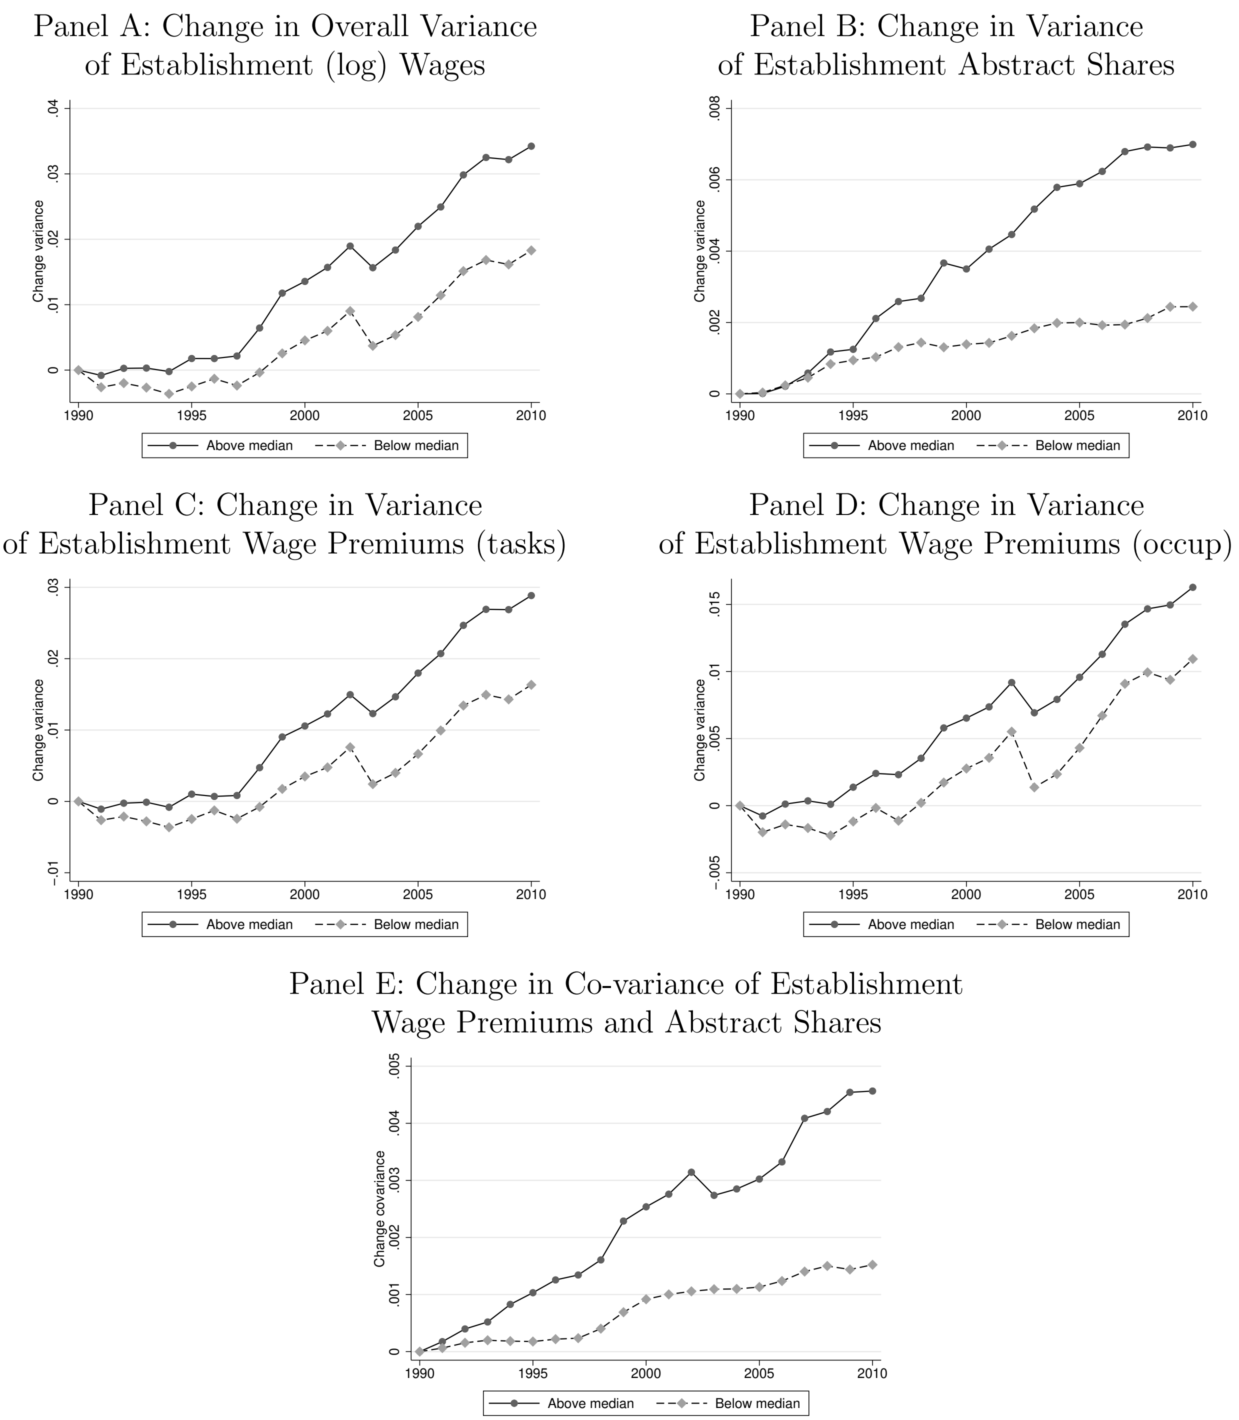

Note: The figures contrast the evolution of the overall variance of average establishment (log) wages (Panel A), the variance of establishments' abstract employment shares (Panel B), the variance of establishments' wage premiums adjusting for their broad task and detailed occupation structure (Panels 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 9: Industries with Below vs Above Median Robot Adoption

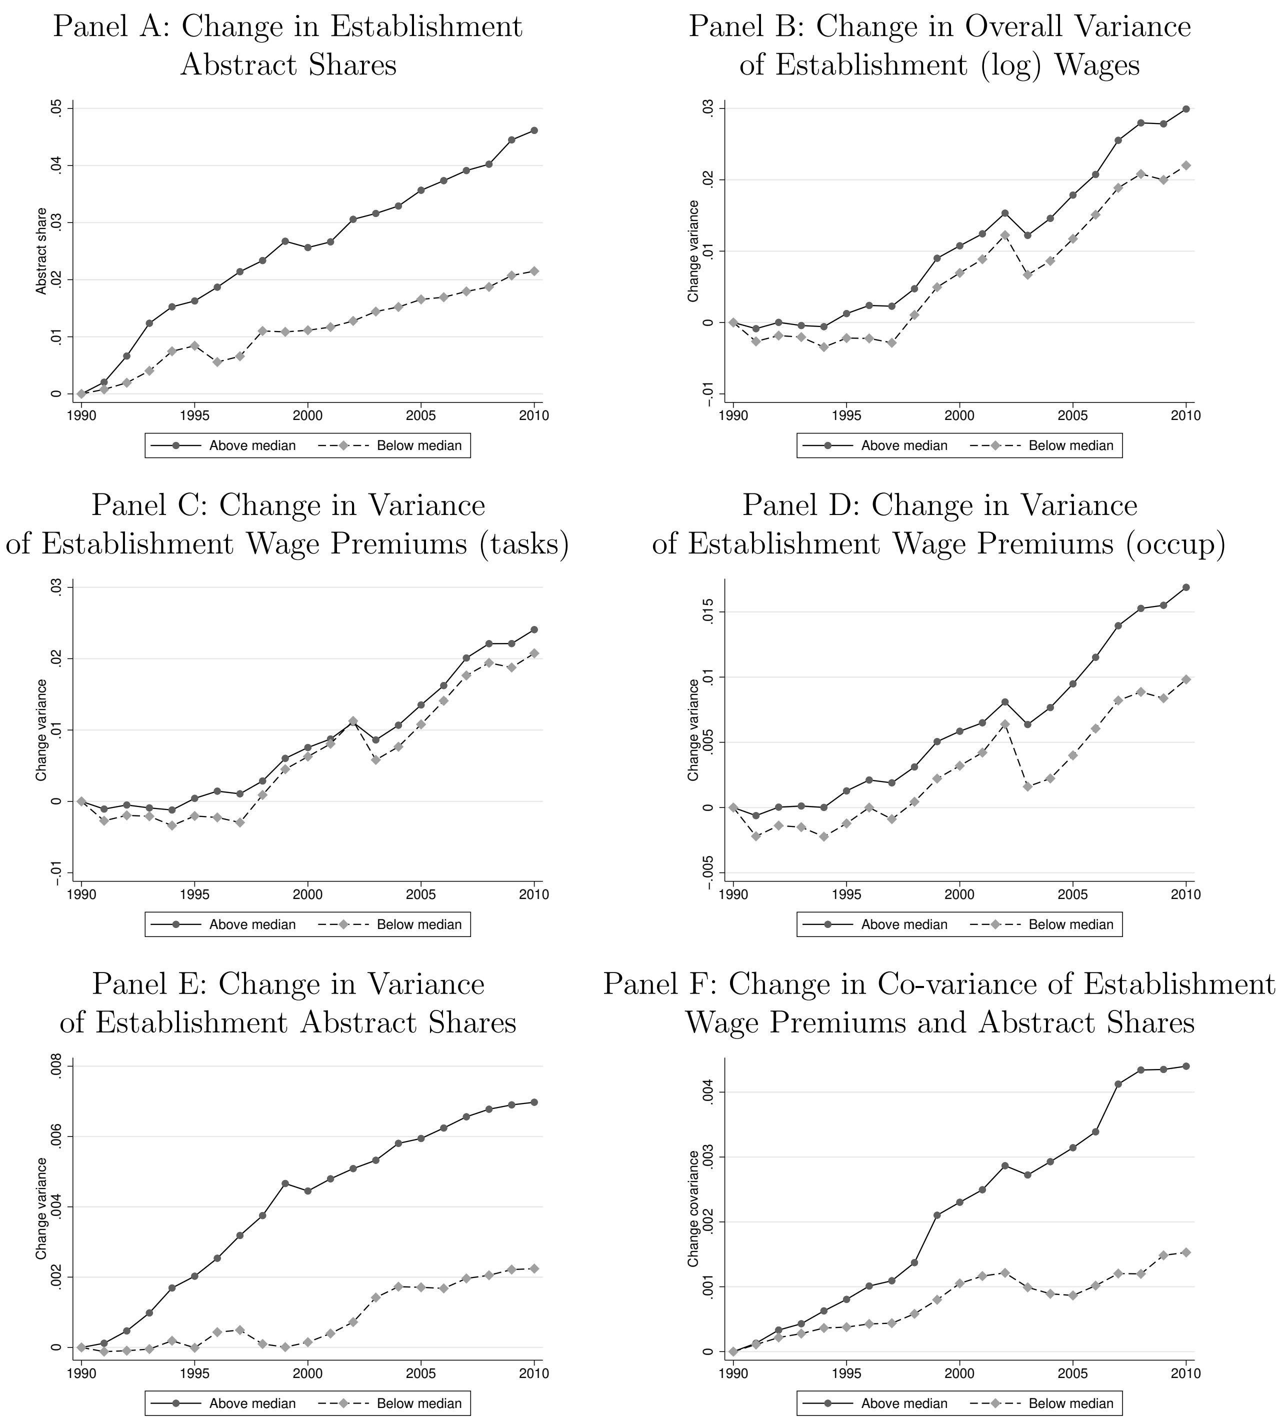

Note: The figures contrast the evolution of the increase in the abstract employment share (Panel A), the variance of average establishment (log) wages (Panel B), the variance of establishments' wage premiums adjusting for their broad task and detailed occupation structure (Panels C and D), the variance 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 based on data from the International Federation of Robotics. We average across industries using the 1990 industry employment structure as weights.

Figure 10: Industries with Below and Above Median ICT Capital Adoption

Panel A: Change in Establishment Abstract Share

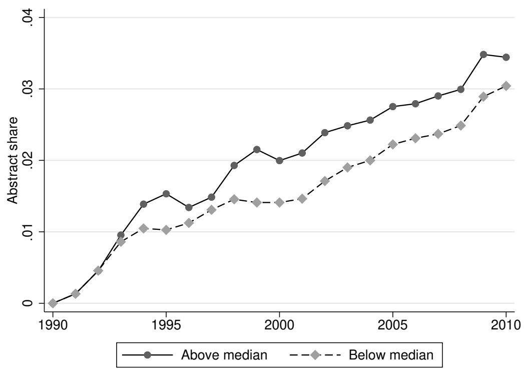

Panel B: Change in Variance of Establishment Wages

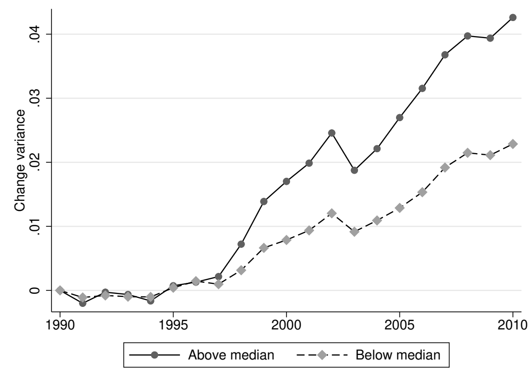

Panel C: Change in Variance of Establishment Wage Premium (tasks)

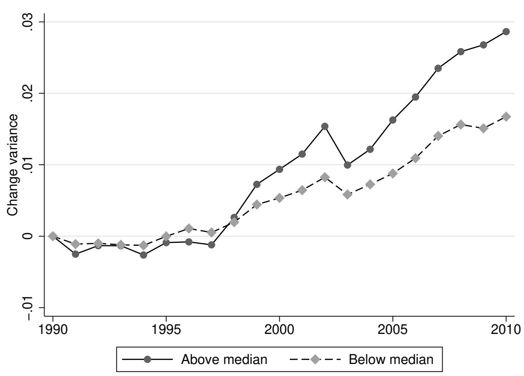

Panel D: Change in Variance of Establishment Wage Premium (occup)

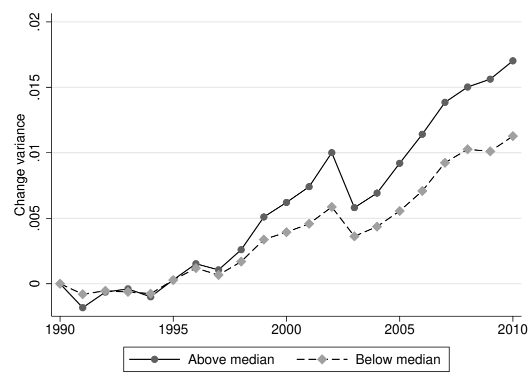

Panel E: Change in Variance of Establishment Abstract Share

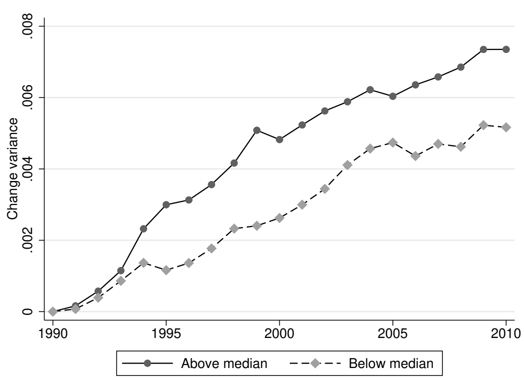

Panel F: Change in Co-variance of Establishment Wage Premium and Abstract Share

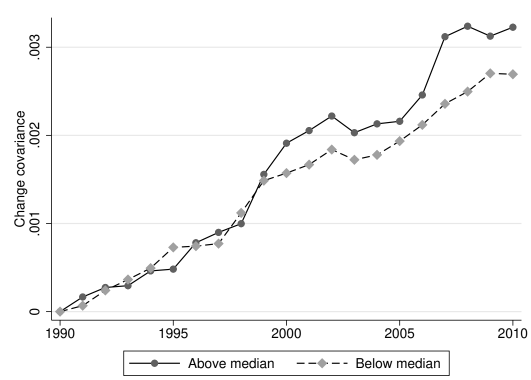

Note: The figures contrast the evolution of the increase in the abstract employment share (Panel A), the variance of average establishment (log) wages (Panel B), the variance of establishments' wage premiums adjusting for their broad task and detailed occupation structure (Panels C and D), the variance 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 1991 and 2007 based on EUKLEMS data. We average across industries using the 1990 industry employment structure as weights.

Table 1: Cross-Sectional Relationships between Wages, Productivity and Employment (Within Industry and Year)

Panel A: Relationship with (Log) Productivity

	Log Abstract Workers	Log Routine Workers	Abstract Share	Avg. Log wage	Establishment Premium (tasks)	Establishment Premium (occup)
	(1)	(2)	(3)	(4)	(5)	(6)
Log Productivity	0.27***	0.19***	0.0065**	0.080***	0.075***	0.056***
(Rev. p. Worker)	(0.031)	(0.025)	(0.003)	(0.0050)	(0.0047)	(0.0035)
N	86,883	86,883	86,883	86,883	86,883	86,883

Panel B: Relationship with Establishment Size

	Productivity	Abstract Share	Avg. Log wage	Establishment Premium (tasks)	Establishment Premium (occup)
	(1)	(2)	(3)	(4)	(5)
Log Estab Size	0.047***	0.0043***	0.077***	0.073***	0.052***
(# of employees)	(0.006)	(0.001)	(0.00064)	(0.00061)	(0.00049)
N	86,883	26,814,744	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. For Columns (1) and (2) of Panel A establishments with no workers of a given type are imputed to have one part-time workers (i.e. 0.5 full-time equivalent workers) of that type in order to be able to compute log employment. All columns in Panel A and Column (1) in Panel B are based on establishments observed in the IABEP; observations are weighted by establishment size and survey weights. All other columns of Panel B are based on establishments observed in the full BEH data, and observations are weighted by establishment size. Standard errors are clustered at the establishment level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Baseline Establishment Characteristics and Within-Establishment Changes (Within Industries)

Panel A: Baseline Establishment Size and Longitudinal Changes in Other Outcomes

	Dependent Variable:				
	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

Panel B: Baseline Establishment Characteristics and Longitudinal Changes in Size

	Dependent Variable: Δ Employment (Incl. Exits)				
	Independent Variable: Baseline Level of ...				
	Estab Productivity	Abstract Share	Avg. Log Wage	Estab Premium (Tasks)	Estab Premium (Occup)
	(1)	(2)	(3)	(4)	(5)
$\hat{\beta}$	0.049***	0.037***	0.11***	0.12***	0.14***
	(0.0084)	(0.0079)	(0.0042)	(0.0041)	(0.0039)
N	15,782	5,107,149	5,107,149	5,107,149	5,107,149

Note: Panel A shows estimated coefficients from regressions of within-establishment changes in the outcome variable shown in each column of the table on baseline establishment size (conditioning on surviving establishments). Panel B shows estimated coefficients from regressions of within-establishment employment growth (including exiting establishment) on the baseline establishment characteristic shown in each column. Within-establishment changes are taken over non-overlapping 5-year windows. Regressions include a set of fully interacted 3-digit industry and year fixed effects. With the exception of Column (1), results are based on establishments in the full BEH and observations are weighted by establishment size. 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 observations 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; see text for details. Standard errors are clustered at the establishment level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 Additional Details on the Empirical Analysis and the Decompositions

Appendix B.1 Establishment Wage Premiums

To compute the establishment wage premiums (two tasks), which we denote as $\widetilde{\ln w_{ft}}$, we first run a regression, separately for each year, of individual log wages on an indicator variable that is equal to 1 if the individual is employed in an abstract task occupation in year t, interacted with K indicator variables that are equal to 1 if the individual is employed in industry k in that year:

$$\ln w_{i(k)t} = \sum_k \beta_{kt} A_{i(k)t} \times D_{i(k)t} + \epsilon_{i(k)t}, \tag{B.1}$$

where $\ln w_{i(k)t}$ is the log wage of individual i employed in industry k at time t, $A_{i(k)t}$ denotes the abstract indicator, $D_{i(k)t}$ is the industry indicator and $\epsilon_{i(k)t}$ is the error term. The establishment wage premium (tasks) of establishment f in year t is then computed as the residuals from the estimation of equation (B.1), averaged across individuals in the establishment:

$$\widetilde{\ln w_{ft}} = \frac{1}{n_{ft}} \sum_{i \in i_{ft}} \epsilon_{i(k)t}, \tag{B.2}$$

where i_{ft} is the set of individuals working in establishment f in year t.

We proceed similarly to compute the establishment wage premiums (occupations), replacing the indicator variable $A_{i(k)t}$ in Equation (B.1) with a full set of occupation fixed effects, thus allowing for occupation wage premiums to differ across industries in each year.

Appendix B.2 Abstract Share (Figure 1, Panel A)

The actual abstract share in year t (the black line in Figure 1, Panel A), S_t, is a weighted average of S_{kt}, the abstract share in each industry k:

$$S_t = \sum_k \frac{n_{kt}}{n_t} S_{kt}, \tag{B.3}$$

where n_{kt} and n_t denote the number of workers employed in industry k at time t and the total number of employed workers at time t, respectively. Hence, (n_{kt}/n_t) denotes industry k's share of employment in year t.

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

$$S_t^{1990} = \sum_k \frac{n_{k1990}}{n_{1990}} S_{kt}, \tag{B.4}$$

where (n_{k1990}/n_{1990}) captures industry k's share of employment in 1990.

Appendix B.3 Abstract Wage Premium (Figure 1, Panel B)

The abstract wage premium, $AbPrem_t$, at time t (the black line in Figure 1, Panel B) is computed as follows:

$$AbPrem_t = \sum_k \frac{n_{kt}}{n_t} \underbrace{\left(\overline{\ln w}_{kt}^{s} - \overline{\ln w}_{kt}^{r}\right)}_{AbPrem_{kt}}, \tag{B.5}$$

where $\overline{\ln w}_{kt}^{s}$ and $\overline{\ln w}_{kt}^{r}$ denote the average log wage of abstract and routine workers in industry k at time t, respectively. Thus, the abstract wage premium is a weighted average of the difference between the average log wage of abstract and routine workers by industry and year, averaged across industries using current industrial employment shares as weights. The term $AbPrem_{kt}$ denotes the industry-year-specific abstract wage premium.

The counterfactual abstract wage premium in year t holding the industry structure constant at its 1990 employment level (the grey line in Figure 1, Panel B) is computed as:

$$AbPrem_t^{1990} = \sum_k \frac{n_{k1990}}{n_{1990}} \left(\overline{\ln w}_{kt}^{s} - \overline{\ln w}_{kt}^{r}\right). \tag{B.6}$$

Appendix B.4 Within-Task Wage Inequality (Figure 1, Panel C)

We compute within-task wage inequality in year t, denoted Var_t^{WT}, as follows:

$$\begin{aligned} Var_t^{WT} &= \sum_\ell \frac{n_{\ell t}}{n_t} \left[\frac{1}{n_{\ell t}} \sum_{i \in i_{\ell t}} (\ln w_{it} - \overline{\ln w_{\ell t}})^2 \right] \\ &= \frac{1}{n_t} \sum_\ell \sum_{i \in i_{\ell t}} (\ln w_{it} - \overline{\ln w_{\ell t}})^2, \end{aligned} \tag{B.7}$$

where ℓ represents a task, i denotes an individual, $i_{\ell t}$ is the set of individuals in task ℓ at time t, $n_{\ell t}$ is the total number of workers in task ℓ at time t, $\ln w_{it}$ is the log wage of individual i at time t and $\overline{\ln w_{\ell t}}$ is the average log wage in task ℓ at time t. Thus, within-task wage inequality is a weighted average of the variance of individual wages by task, averaged over tasks using the actual share of employment in each task. Within-occupation wage inequality is computed similarly, with ℓ denoting 3-digit occupations instead of the two broad task categories.

Appendix B.5 Within-Industry Between-Establishment Wage Inequality (Figure 2, Panel B)

We compute the counterfactual within-industry between-establishment wage inequality using the 1990 industry structure, denoted by $Var_t^{WIBE,1990}$, as follows:

$$Var_t^{WIBE,1990} = \sum_k \frac{n_{k1990}}{n_{1990}} \underbrace{\sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (\overline{\ln w_{ft}} - \overline{\ln w_{kt}})^2}_{Var_{kt}(\overline{\ln w_{ft}})} . \tag{B.8}$$

where f indexes establishments, f_{kt} is the set of establishments in industry k in year t, n_{ft} is the total number of workers at establishment f in year t, $\overline{\ln w_{ft}}$ is the average log wage in establishment f at time t and $\overline{\ln w_{kt}}$ is the average log wage in industry k at time t. The term $Var_{kt}(\overline{\ln w_{ft}})$ refers to the variance of establishment log wages in industry k at time t.

Appendix B.6 Within-Industry Between-Establishment Wage Premium Inequality (Figure 2, Panel C)

We compute the counterfactual variance of within-industry between-establishment wage premiums using the 1990 industry structure, denoted $Var_t^{W\widetilde{IB}E,1990}$, as follows:

$$Var_t^{W\widetilde{IB}E,1990} = \sum_k \frac{n_{k1990}}{n_{1990}} \sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (\widetilde{\ln w}_{ft} - \overline{\widetilde{\ln w}_{kt}})^2,$$

where $\widetilde{\ln w}_{ft}$ denotes the wage premium of establishment f at time t, and $\overline{\widetilde{\ln w}_{kt}}$ is the average establishment wage premium in industry k at time t. The wage premiums are computed as detailed in Appendix B.1.

Appendix B.7 Within-Industry Abstract Share Heterogeneity (Figure 6, Panel A)

Within-industry heterogeneity in establishments' abstract share, denoted Var_t^A, is given by the within-industry variance in establishments' abstract share, averaged over industries using industrial employment shares as weights:

$$Var_t^A = \sum_k \frac{n_{kt}}{n_t} \sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (S_{ft} - S_{kt})^2,$$

where S_{ft} is establishment f's abstract employment share at time t.

The counterfactual within-industry variance in establishments' abstract employment shares, holding the industry structure constant at its 1990 level, $Var_t^{A,1990}$, equals:

$$Var_t^{A,1990} = \sum_k \frac{n_{k1990}}{n_{1990}} \sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (S_{ft} - S_{kt})^2.$$

Appendix B.8 Within-Industry Sorting (Figure 6, Panel B)

We capture the extent of sorting of abstract workers into high-wage establishments using the within-industry co-variance between establishments' abstract employment shares and their wage premiums, averaged across industries using industrial employment shares. This co-variance, denoted Cov_t, is computed as follows:

$$Cov_t = \sum_k \frac{n_{kt}}{n_t} \sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (S_{ft} - S_{kt})(\widetilde{\ln w_{ft}} - \overline{\widetilde{\ln w_{kt}}}),$$

and the corresponding counterfactual co-variance, holding the industry structure constant at its 1990 level, denoted Cov_t^{1990} , equals:

$$Cov_t^{1990} = \sum_k \frac{n_{k1990}}{n_{1990}} \sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (S_{ft} - S_{kt})(\widetilde{\ln w_{ft}} - \overline{\widetilde{\ln w_{kt}}}).$$

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 also refer the reader 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 $\phi_\ell(\theta)$, where $\ell \in \{s, r\}$, are defined as follows:

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

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 revenue, employment and wages by tasks:

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

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

Revenue As Helpman et al. (2010) mention in Appendix 5.4 footnote 1, revenue 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 $\kappa_y \equiv \frac{k a_{min}^{\gamma k}}{k-1}$. 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):

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

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

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

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:

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

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:

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

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 pinned 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:[26]

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

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

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

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:

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

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 (4) and (7) 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 $1 - \frac{k}{\delta} > 0$.

Equation (4) Result: More productive firms are larger, employing more workers of both types, and have a higher abstract employment share.

Proof: Taking the derivative of equations (C.10) and (C.12), we obtain:

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

Taking the derivative of equation (C.13), we have that:

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

Equation (7) Result: More productive firms pay higher wages to both types of workers.

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

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

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 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:

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

and

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

Prediction 1: *Increased Abstract Wage Premium* – Task-biased technological change increases the abstract wage premium within all firms, and in the aggregate.

Proof: Taking the first-order derivative of (C.16), we have:

$$\frac{\partial}{\partial \mu_s}\left[\frac{w_s(\theta)}{w_r(\theta)}\right] \quad = \quad \frac{k\nu}{\delta\Lambda}\mu_s^{-1}\frac{w_s(\theta)}{w_r(\theta)} \quad > \quad 0 \tag{C.31}$$

Prediction 2: *Selection* – TBTC increases the productivity threshold for production θ_d.

Proof: We prove Prediction 2 by contradiction. Consider equation (C.24), which pins down the equilibrium threshold as a function of parameters of the model:

$$f_d \int_{\theta_d}^{\infty} \left(\left[\frac{1+\varphi(\theta)}{1+\varphi(\theta_d)} \right]^{\frac{\beta\Lambda}{\nu\Gamma}} - 1 \right) dG(\theta) = f_e \tag{C.32}$$

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.32) 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.32) 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.32) is an increase in θ_d when μ_s increases. Therefore:

$$\frac{\partial \theta_d}{\partial \mu_s} \quad > \quad 0 \tag{C.33}$$

Prediction 3: *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, and 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, and more so for more productive firms.

Prediction 4: *Increased Sorting and Segregation by Task* – TBTC strengthens the cross-sectional association between productivity and abstract employment shares, provided that firms employ relatively more routine than abstract workers at baseline (the empirically relevant case).

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,

$$\frac{\partial^2}{\partial\mu_s\partial\theta}\left[\frac{h_s(\theta)}{h(\theta)}\right] \quad > \quad 0 \quad \text{if} \quad \frac{h_s(\theta)}{h_r(\theta)} < 1 \tag{C.34}$$

Prediction 5: *Differential Wage Growth* – TBTC strengthens the cross-sectional association between productivity and wages conditional on worker type.

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.

Appendix C.4 Impact of an Increase in the Supply of Abstract Workers

In this section we study the implications of an increase in the supply of abstract workers, modeled as a fall in the search costs of abstract workers. Intuitively, when the supply of abstract workers increases, it becomes easier for firms to fill their vacancies, therefore reducing the cost of searching for workers. For simplicity, we normalize the search cost of routine workers to one, i.e. $b_r = 1$, and set $b_s = b$, and study the impacts of a fall in b. It is useful to note that this normalization implies that κ_r, h_{dr} and w_{dr} are independent of b, the search costs for abstract workers. In addition, note that $\varphi(\theta) = \mu_s^{\nu/\Lambda}\left(\frac{\lambda_s}{\lambda_r}\right)^{1/\Lambda} b^{-\gamma\nu/\Lambda}\theta^{\nu/\Lambda}$, which implies $\frac{\partial\varphi(\theta)}{\partial b} < 0$.

In what follows, we proceed in the same order as we did for TBTC, first evaluating how the shock affects firm-level outcomes and then examining the heterogeneity of the effects across the productivity distribution. In doing so we make some parameter restrictions to ensure that, consistent with the existing literature, wages of abstract workers fall when the supply of abstract workers rises. In particular, we assume that $1 - \frac{k\gamma\nu}{\delta\Lambda}\left(1 + \frac{\beta-\nu}{\nu\Gamma}\right) > 0$.

Prediction 1: *Decreased Abstract Wage Premium* – An increase in the supply of abstract workers decreases the abstract wage premium within all firms, and in the aggregate.

Proof:

$$\begin{aligned}
\frac{\partial w_r(\theta)}{\partial b} &= -\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k\gamma\nu}{\delta\Lambda}b^{-1}\frac{\varphi(\theta)}{1+\varphi(\theta)}w_r(\theta) \quad < \quad 0 \\
\frac{\partial w_s(\theta)}{\partial b} &= w_r(\theta)\varphi(\theta)^{k/\delta}\left\{1 - \frac{k\gamma\nu}{\delta\Lambda}\left[1 + \left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{\varphi(\theta)}{1+\varphi(\theta)}\right]\right\} \quad > \quad 0 \quad \text{if} \quad 1 - \frac{k\gamma\nu}{\delta\Lambda}\left(1 + \frac{\beta-\nu}{\nu\Gamma}\right) > 0
\end{aligned}$$

When the search cost of abstract workers falls, wages of routine workers will increase and (under our parameter restrictions) wages of abstract workers will fall. Intuitively, wages are bargained down to the replacement cost of a worker. This cost is a function of both, (i) the search cost and (ii) a firm's (endogenously chosen) ability threshold, which in turn is positively related to revenue. Thus, a change in the search cost impacts wages directly and, via its effect on the ability threshold, indirectly.

A fall in b makes it cheaper for firms to find abstract workers and hence increases firms' revenues. As a result, firms screen more intensively and choose a higher ability threshold not only for abstract but also for routine workers. Thus, even if the search cost of routine workers remains unchanged, a higher ability threshold translates into higher routine wages. For abstract workers, the wage effects are ambiguous because the direct and indirect wage effects of a fall in b work at cross-purposes. On one hand, the increase in revenue leads firms to choose a higher ability threshold. This effect pushes wages upwards for both abstract and routine workers (although to varying extents). On the other hand, abstract workers become cheaper to replace. Given that wages are adjusted down to the replacement cost of workers, this direct effect (which is absent for routine types) gives rise to a fall in abstract wages.

Prediction 2: *Selection* – A positive shock to the supply of abstract workers increases θ_d.

Proof: Prediction 2 can be proven by contradiction, as we do for TBTC. Noting that b only appears in $\varphi(\theta)$, it is straightforward to show that a fall in b will increase θ_d, the productivity threshold for production, hence leading to the exit of the least productive firms.

Prediction 3: *Differential Employment Growth* – A positive shock to the supply of abstract workers strengthens the cross-sectional association between employment and productivity.

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

$$\frac{\partial h_r(\theta)}{\partial b} = -\frac{\gamma\nu}{\Lambda} b^{-1} \left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right) h_r(\theta) \left[1+\varphi(\theta)\right]^{-1} \varphi(\theta) \quad < \quad 0.$$

and recalling that $\frac{\partial\varphi(\theta)}{\partial\theta} > 0$ and noting that $1+\left[\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)-1\right]\frac{\varphi(\theta)}{1+\varphi(\theta)} > 0$:

$$\frac{\partial^2 h_r(\theta)}{\partial b \partial\theta} = \frac{\partial h_r(\theta)}{\partial b} \varphi(\theta)^{-1} \frac{\partial\varphi(\theta)}{\partial\theta} \left\{1+\left[\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)-1\right]\frac{\varphi(\theta)}{1+\varphi(\theta)}\right\} \quad < \quad 0$$

Hence, an abstract-biased supply shock increases routine employment for all firms, but more so for more productive firms. Similarly, taking the derivatives of (C.12) and using the previous results:

$$\frac{\partial h_s(\theta)}{\partial b} \quad = \quad \frac{1}{b}\varphi(\theta)^{1-k/\delta}\left\{\frac{\partial h_r(\theta)}{\partial b} - \frac{1}{b}\left[1+\frac{\gamma\nu}{\Lambda}\left(1-\frac{k}{\delta}\right)\right]\right\} \quad < \quad 0$$

and recalling that $\frac{\partial h_r(\theta)}{\partial\theta} > 0$:

$$\frac{\partial^2 h_s(\theta)}{\partial b \partial\theta} \quad = \quad \left(1-\frac{k}{\delta}\right)\varphi(\theta)^{-1}\frac{\partial\varphi(\theta)}{\partial\theta}\frac{\partial h_s(\theta)}{\partial b} + \varphi(\theta)^{1-k/\delta}\left\{-\frac{1}{b^2}\left[1+\frac{\gamma\nu}{\Lambda}\left(1-\frac{k}{\delta}\right)\right]\frac{\partial h_r(\theta)}{\partial\theta} + \frac{1}{b}\frac{\partial^2 h_r(\theta)}{\partial b\partial\theta}\right\} \quad < \quad 0$$

Hence, a positive shock to the supply of abstract workers also increases abstract employment for all firms, but more so for more productive firms.

Prediction 4: *Increased Sorting and Segregation by Task* – An increase in the supply of abstract workers strengthens the cross-sectional association between productivity and abstract employment shares.

Proof: Taking the first-order derivative of (C.13), we get:

$$\frac{\partial}{\partial b}\left[\frac{h_s(\theta)}{h(\theta)}\right] = \frac{1}{\left[b+\varphi(\theta)^{1-k/\delta}\right]}\frac{h_s(\theta)}{h(\theta)}\left[\left(1-\frac{k}{\delta}\right)\varphi(\theta)^{-1}\frac{\partial\varphi(\theta)}{\partial b}b-1\right] \quad < \quad 0$$

Hence, an abstract-biased supply shock increases the share of abstract workers for all firms. Finally, given our assumption that firm abstract employment is lower than routine employment at baseline (i.e. $b-\varphi(\theta)^{1-k/\delta} > 0$) and taking the second-order derivative of firm abstract employment share, we obtain:

$$\frac{\partial^2}{\partial b\partial\theta}\left[\frac{h_s(\theta)}{h(\theta)}\right] = \frac{\partial}{\partial b}\left[\frac{h_s(\theta)}{h(\theta)}\right]\cdot\frac{1}{\left[b+\varphi(\theta)^{1-k/\delta}\right]}\left(1-\frac{k}{\delta}\right)\varphi(\theta)^{-1}\frac{\partial\varphi(\theta)}{\partial\theta}\left[b-\varphi(\theta)^{1-k/\delta}\right] \quad < \quad 0$$

This result implies that the increase in the share of abstract workers will be larger for more productive firms assuming that the number of routine workers outweighs the number of abstract workers at baseline, an assumption supported by our data. Therefore, a positive shock to the supply of abstract workers increases sorting of abstract workers into firms with relatively higher productivity levels.

Prediction 5: *Differential Wage Growth* – An increase in the supply of abstract workers unambiguously leads to an increase in the wages of routine workers and this increase is disproportionately larger for firms with relatively higher productivity levels. Under some reasonable parameter restrictions, wages of abstract workers will fall and more so for more productive firms.

Proof: Taking the second-order derivative of (C.16) and recalling that $\frac{\partial w_r(\theta)}{\partial \theta} > 0$ and $\frac{\partial \varphi(\theta)}{\partial \theta} > 0$, we obtain:

$$\frac{\partial^2 w_r(\theta)}{\partial b \partial \theta} = -\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k\gamma\nu}{\delta\Lambda}b^{-1}\frac{1}{1+\varphi(\theta)}\left[\varphi(\theta)\frac{\partial w_r(\theta)}{\partial \theta}+\frac{w_r(\theta)}{1+\varphi(\theta)}\cdot\frac{\partial \varphi(\theta)}{\partial \theta}\right] < 0$$

Hence, an increase in the supply of abstract workers unambiguously raises firm wages of routine workers (see Prediction 1), and more so for more productive firms.

By contrast, whether wage inequality between firms increases for abstract workers depends on parameter restrictions. Taking the second derivative of abstract wages with respect to the search costs yields:

$$\begin{aligned}\frac{\partial^2 w_s(\theta)}{\partial b \partial \theta} &= \left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}\left[1+\varphi(\theta)\right]^{-1}\frac{\partial \varphi(\theta)}{\partial \theta}w_r(\theta)\varphi(\theta)^{k/\delta}\left\{1-\frac{k\gamma\nu}{\delta\Lambda}\left[1+\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{\varphi(\theta)}{1+\varphi(\theta)}\right]-\frac{\gamma\nu}{\Lambda}\left[\frac{k}{\delta}+\frac{1}{1+\varphi(\theta)}\right]\right\}\\ &+ \frac{k}{\delta}\frac{\partial \varphi(\theta)}{\partial \theta}w_r(\theta)\varphi(\theta)^{\frac{k}{\delta}-1}\end{aligned}$$

Under the assumption that $1-\frac{k\gamma\nu}{\delta\Lambda}\left(1+\frac{\beta-\nu}{\nu\Gamma}\right)>0$, wages of abstract workers fall (see Prediction 1). In addition, assuming that $1-\frac{k\gamma\nu}{\delta\Lambda}\left(1+\frac{\beta-\nu}{\nu\Gamma}\right)-\frac{\gamma\nu}{\Lambda}\left(1+\frac{k}{\delta}\right)>0$, wage changes of abstract workers would be disproportionately larger for top firms.

Figure A.1: Decomposition of Changes in the Within-Industry Between-Establishment Variance of Log Wages: 3-Digit Occupations

Panel A: Continuing Establishments, Within-Occupation, Segregation and Sorting

Panel B: Continuing Establishments, Differential Employment vs Wage Growth

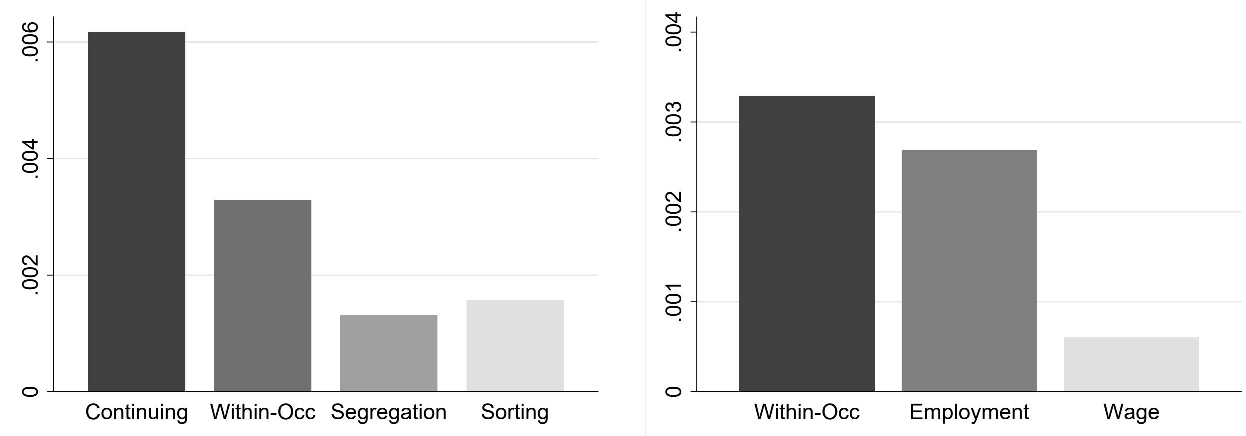

Note: Panel A decomposes changes in the within-industry variance of log wages among continuing establishments into within-occupation, segregation and sorting components; see Equation (9). Panel B decomposes changes in the within-occupation component among continuing establishments into differential employment and wage growth components; see Equation (10). 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 counting as 0.5 and full-time employees counting as 1.