Technological Change, Firm Heterogeneity and Wage Inequality*

Guido Matias Cortes[†] Adrian Lerche[‡]

Uta Schönberg[§] Jeanne Tschopp[¶]

December 17, 2025

Abstract

We argue that skill-biased technological change not only affects wage gaps between skill groups but also increases wage inequality within skill groups across workers in different workplaces. Building on a heterogeneous firm framework with labor market frictions, we show that an industry-wide skill-biased technological change shock will increase between-firm wage inequality within the industry through three main channels: increased employment concentration in more productive firms, increased wage dispersion between firms for workers of the same skill type, and increased segregation and sorting of skilled workers in more productive firms. Using rich, matched employer-employee administrative data from Germany, we provide empirical evidence of establishment-level patterns consistent with the model's predictions. We further document that industries with more exposure to technological change exhibit particularly pronounced 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). Schönberg is grateful to the European Research Council for its financial support (Grant Number 818992 (FirmIneq)). We are also grateful to numerous conference and seminar participants for valuable comments and suggestions.

†York University; gmcortes@yorku.ca.

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

§University of Hong Kong, 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@unibe.ch.

1 Introduction

Wage 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 heterogeneous impacts across different groups of workers, depending on their skill levels or the tasks 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, thereby reshaping the employment structure of the economy and the relative wage returns for different groups.

Intriguingly, however, a recent parallel literature has shown that a large share of the rise in wage inequality can be traced 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; OECD, 2021). This literature has argued that individual-level wages have become increasingly dependent on where people work, rather than on the skills they possess or the tasks they perform.

While the literature on between-firm inequality has documented many novel empirical facts, it is not yet entirely clear what the driving forces behind these patterns are. 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 aggregate production functions, so far provided limited insights into the growing wage differentials observed *within* groups, across firms.

In this paper, we argue that the development of skill-biased technologies can account not only for increases in inequality between skill groups but also for increases in inequality within groups, across workers in different workplaces, as observed in the data. We show this theoretically, using a rich yet tractable heterogeneous-firm framework, and empirically, by verifying the model's predictions using matched employer-employee administrative data from Germany.

Our dataset is the *Beschäftigtenhistorik* (BeH) from the Institute for Employment Research (IAB), which covers the universe of workers and establishments covered by the German social security system. We focus on the patterns observed in West Germany between 1990 and 2010, a period of sharply rising wage inequality.

To motivate our analysis, we first show that both the skill premium and the employment of skilled workers relative to unskilled workers rose over this time period. This is consistent with increased relative demand for skilled labor, as posited by the classic skill-biased technological change (SBTC) literature. Furthermore, we show that the increase in the variance of individual log wages is almost entirely driven by growing wage dispersion across establishments, with more than half of this increase occurring within detailed industries. This rise in wage inequality across establishments within industries results from both an increasing dispersion in skill composition across establishments, as well as rising wage differences across establishments conditional on skill.

Guided by this motivating evidence, we set up a theoretical framework that analyzes the impacts of a technological change shock along the lines of the classic SBTC model, but in a setting that allows for heterogeneous firms. Our model incorporates two skill groups. To generate the heterogeneity in skilled employment composition that we document empirically, we allow for skill-biased productivity differences across firms. Moreover, to generate an equilibrium in which firms with different productivity levels pay heterogeneous wages to observationally equivalent workers, our framework incorporates labor market imperfections via search and matching frictions, heterogeneous match-specific ability, and a screening technology, drawing on the closed economy version of the model of Helpman et al. (2010). This supports an equilibrium in which more productive firms find it optimal to employ more workers, have a higher skilled-to-unskilled employment ratio, and pay higher wages, both overall and conditional on skill type.

The SBTC shock that we embed in the model is in the spirit of Katz & Murphy (1992) and Autor et al. (1998), and involves an aggregate increase in the factor-augmenting technology parameter for skilled workers in the production function. As in traditional models of SBTC, this shock leads to an increase in the skilled wage premium and, consequently, increases between-group wage inequality. However, despite being an aggregate shock that is common across all firms, SBTC also induces several endogenous heterogeneous firm-level changes that increase between-firm inequality.

First, the model predicts that SBTC leads to differential employment growth, whereby more productive and higher-paying firms grow faster than less productive firms in the industry. This leads to a rise in employment concentration, which in turn increases worker-weighted measures of between-firm wage inequality. Second, the model predicts that SBTC leads to an endogenous increase in worker segregation by skill, driven by increased sorting of skilled workers to high-productivity (and hence high-wage) firms. Third, the model generates endogenous within-firm wage changes, with more productive firms disproportionately increasing the wages they pay to workers in each skill group, thereby further contributing to the increase in between-firm wage inequality. These changes lead to an increase in between-firm inequality due to rising dispersion in skill composition across firms, as well as rising between-firm inequality in wages conditional on skill—in line with the empirical patterns that we document over time.

In order to test the empirical validity of the firm-level mechanisms highlighted by the model, we derive two model-consistent empirical proxies for the underlying productivity type of each establishment. The first proxy is a normalization of the establishment's log employment ratio between high-skilled and low-skilled workers. Our model demonstrates that this proxy captures the variation across establishments attributable to the productivity parameter while removing the influence of aggregate and industry-specific shocks such as SBTC. As an alternative proxy for establishment productivity, we characterize each establishment by its baseline (1990) log employment level. This proxy relies less on the model structure and, due to its predetermined nature, is also unaffected by shocks such as SBTC.

Consistent with the presence of ongoing SBTC, we find that the establishment-level associations between our productivity proxies and employment, skill intensities, and wages have strengthened over our sample period within industries. For example, while the establishment productivity parameter was not significantly associated with the establishment's employment level in the early 1990s, this association grew to more than 0.1% by 2010. Similarly, a 1% increase in the establishment productivity proxy was associated with an increase of about 0.05% in the average establishment wage in the early 1990s, increasing to almost 0.1% by 2010. This strengthening of the productivity-wage relationship arises from a stronger relationship between productivity and skill intensities, as well as between productivity and establishment wages conditional on worker skill.

These patterns imply that gaps in employment, skill intensities, and wages between low- and high-productivity establishments have widened over time, and are consistent with the channels highlighted by the model: differential increases in employment, skill intensity, and wages in more productive firms relative to less productive ones, driven by SBTC. We also find similar, if not larger, increases in the strength of the association over time between our establishment productivity proxies and average *residual* log wages for both skilled and unskilled workers, suggesting that the widening gaps in average skill-specific wages between low- and high-productivity establishments are not solely due to increased sorting based on observable characteristics, but also involve genuine changes in establishment pay premia conditional on observable factors.

In the final step of the empirical analysis, we leverage variation in exposure to SBTC across industries to provide more direct evidence on the link between technological change and the various drivers of between-establishment wage inequality. We measure exposure to SBTC using industry-level changes over time in the stock of capital related to information and communication technologies (ICT) per worker, as captured in the EUKLEMS dataset. We first show that industries with more adoption of ICT experience larger increases in wage inequality between establishments, compared to industries with less adoption. This is due to both larger increases in the variance of relative skill intensities across establishments in high-adoption industries, as well as larger increases in the variance of establishment-level wages conditional on worker skill. We then show that the strengthening of the establishment-level relationships between productivity, employment, skill intensities, and wages is more pronounced in high-adoption industries, thus providing support for the workplace-level mechanisms highlighted by the model. An alternative measure of exposure to SBTC based on changes in industry skill premia yields similar patterns. Moreover, our findings persist when we control for differential trade and offshorability exposure across industries, corroborating the importance of SBTC as a driver of between-establishment inequality.

Our findings contribute to several strands of the economic literature. First, we introduce a significant innovation to the study of how technological change affects wage inequality. Previous research has primarily relied on frameworks with aggregate production functions and perfectly competitive labor markets, focusing on how technological change influences inequality through shifts in the skill or task structure of the economy (see, e.g. Katz & Murphy, 1992; Machin & Van Reenen, 1998; Acemoglu & Autor, 2011; Michaels et al., 2014; Autor et al., 2015; Akerman et al., 2015; Graetz & Michaels, 2018; Dauth et al., 2021). In contrast, our framework and empirical analysis demonstrate that an industry-wide technological shock can produce markedly different effects across firms within the same industry and consequently among workers within the same skill group. These results broaden our understanding of the impact of SBTC on wage inequality by emphasizing that the effects of technological change on workers depend not only on their skill level, but also on the specific firm to which they are matched.

Our analysis also provides an important contribution to the literature on between-firm wage inequality. Empirically, this literature has been highly successful in documenting the increasing importance of firms for individual wages (e.g. Card et al., 2013; Song et al., 2019; Barth et al., 2016; Helpman et al., 2017; Criscuolo et al., 2020), and has provided evidence of the rise of worker sorting and segregation (e.g. Kramarz et al., 1996; Cortes & Salvatori, 2019; Wilmers & Aeppli, 2021). From a theoretical perspective, various approaches have been proposed in order to model wage differences between firms, such as fair wage considerations (e.g. Egger & Kreickemeier, 2009), idiosyncratic heterogeneities in worker preferences for firm amenities (e.g. Card et al., 2018), and frictions related to search and matching and screening (Helpman et al., 2010). Papers in this literature, however, have not focused on the interplay between skill-biased technological change and workplace heterogeneity.[1] We study

[1]Song et al. (2019) quantify the extent to which the aggregate increase in the skill premium can account for the rise of segregation and sorting in the U.S., but do not consider how skill-biased technological change could have broader impacts on workplace heterogeneity. Other papers in the literature have provided a rich analysis of how technology affects the sorting of workers to jobs (e.g. Lindenlaub, 2017). However, these types

this interplay within a tractable theoretical framework, and highlight the empirical relevance of our theoretically-identified channels. We further show that sorting and wage dispersion patterns across establishments are more pronounced within industries that have been more exposed to technological change, consistent with the idea that SBTC is an important factor in the rise in between-workplace inequality.

Our paper also relates 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., 2020, 2022; Bajgar et al., 2019). We show that skill-biased technological change increases employment concentration in highly productive firms. Even in the absence of any wage changes within firms, this rise in concentration will imply an increase in worker-weighted measures of between-establishment wage inequality.[2]

Finally, we relate to recent studies that investigate the impact of *firm-level* adoption of industrial robots (e.g. Acemoglu et al., 2020; Bonfiglioli et al., 2020; Koch et al., 2021), automation expenditures (Bessen et al., 2020; Aghion et al., 2020) or innovation (Lindner et al., 2021) on firm outcomes. These papers generally find that technology adoption is associated with increases in employment, sales, and skill intensities at the firm level. While we are also interested in the firm-specific impacts of technological advances, our study highlights that an *industry-wide* shock can have differential effects on firms.

The paper most related to ours is Haanwinckel (2020). While both studies aim to understand the role of technological change in driving wage inequality, our frameworks and approaches differ substantially. Haanwinckel (2020) employs a model that generates wage differences between firms by assuming that workers have idiosyncratic tastes for different workplaces (as in e.g. Bhaskar et al., 2002; Card et al., 2018). Our model, in contrast, generates wage heterogeneity due to search and matching frictions and match-specific worker ability. Haanwinckel (2020) disentangles the effects of various shocks, including skill-biased technological change, changes in the supply of skilled workers, and minimum wages, by calibrating their model to match the increase in wage inequality in Brazil. We, on the other hand, use our model to derive comparative statics that clarify the mechanisms through which technological change influences between-firm inequality, and assess the empirical relevance of these mechanisms. Moreover, we leverage variation in technology exposure across industries to provide additional evidence on the role of technological change in exacerbating wage inequality between establishments.

---

of models often have no natural definition of a firm and assume that worker types are perfect substitutes in production.

[2]See Webber (2015); Mueller et al. (2017); Rinz (2022) and Cortes & Tschopp (2023) for more detailed analyses of the link between rising concentration and rising wage inequality.

2 Data

2.1 Social Security Records (Beschäftigtenhistorik)

Our main data are drawn from social security records for Germany provided by the Institute for Employment Research (IAB)—the so-called *Beschäftigtenhistorik* (BeH, 2016 version). This data set includes 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.

We focus on developments after 1990, when wage inequality began to increase sharply across the entire wage distribution in Germany (see, for example, Card et al., 2013; Dustmann et al., 2014). We end the analysis in 2010 due to structural breaks in key variables, such as the full-time employment status of workers, which affect comparability with subsequent years.

We begin by selecting all full- and part-time employment spells that refer to June 30 of each year. We then restrict the sample to workers who are not currently in an apprenticeship, are aged 16-65, and are employed in West Germany. We exclude primary-sector industries and some small industries, such as private households and international organizations. We also drop workers with missing data for their occupation or their education, as well as spells with implausibly low wages below the limit for which social security contributions have to be paid, establishments with missing industry affiliation, and establishments employing only part-time workers.

As is common in administrative data sources, wages are censored at the highest social security limit, affecting, on average, about 9% 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.[3]

We classify individuals as either "skilled" or "unskilled", using information on their education and apprenticeship occupation. Specifically, technical college and university graduates are classified as skilled, while individuals with no completed vocational or post-secondary education are classified as unskilled. Individuals with apprenticeship or vocational training—which comprise the largest share of the workforce in Germany (around 75%)—are categorized as skilled or unskilled depending on whether their apprenticeship occupation is a primarily "skilled" (e.g. technicians) or "unskilled" (e.g. manual labor) occupation, following the map-

[3]Note that individual wages in the data set always refer to a single establishment and are never averaged across establishments.

ping of Blossfeld (1987); see Appendix Table A.1.[4] This classification of skill groups mimics the distinction between college and non-college graduates typically used in the U.S. context, as those apprenticeship graduates whom we classify as skilled would typically enroll in college in the U.S. In contrast, those whom we classify as unskilled would typically enter the labor market without a post-secondary education.

We make use of the unique establishment identifiers available in the data to aggregate the worker-level information to the establishment level in each year. Our establishment-level employment counts include part-time workers with a weight of 0.5. Since we do not observe hours worked, our measures of establishment wages are based on full-time workers only.[5]

We focus on establishment patterns within industries, where industries are defined as 3-digit NAICS codes that distinguish between 196 sectors. To address the change in the industry classification in the social security data that occurred in 1999, we harmonize industry codes as described in Appendix A.2.

2.2 Industry-Level Technology Exposure Measures

We supplement our primary data source with information from the 2009 release of the EU-KLEMS dataset. We measure exposure to SBTC using the industry-level change over time in the real fixed capital stock of information and communication technologies (ICT) per worker. ICT capital per worker is obtained by dividing the total stock of computing and communication equipment and software, at 1995 prices, by total industry-level employment, as reported in the EUKLEMS data. The 2009 release of EUKLEMS reports industry codes from ISIC Revision 3, allowing us to match the technology exposure measures to the BeH social security data at the two-digit industry level.

Motivated by the standard prediction that SBTC increases the relative demand for skilled labor and, hence, the skill premium, we use the industry-level increase in the skill premium as an alternative measure of industry-level technology exposure.

3 Motivating Evidence

In this section, we confirm the empirical insights highlighted by the classic SBTC literature and, consistent with recent studies, show that between-establishment wage differences are a

[4] If an individual with apprenticeship education is not observed at the time of their training, we use the occupation when they are first observed in the data to classify them as skilled or unskilled.

[5] 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 does not drive the rise in inequality over time.

major component of overall wage inequality and of its increase over time. Additionally, we show that within 3-digit industries, average wages across establishments vary due to both differences in skill composition and differences in establishment wages conditional on worker skill. These findings inform and motivate our modeling choices in the theoretical framework presented in Section 4.

Evidence Consistent With SBTC The traditional SBTC literature (e.g. Katz & Murphy, 1992) documents a concurrent rise in both relative employment and relative wages of skilled workers, consistent with the idea that technological change increases the relative demand for skilled workers and that this increase outpaces the growth in relative supply arising from increasing levels of educational attainment.

The top two panels of Figure 1 corroborate these trends in our data. The black line in Panel A shows that the relative employment of skilled workers rose steadily in Germany, from a log ratio of -0.45 in 1990 to a log ratio of nearly -0.11 in 2010—a rise of about 40% over two decades. This increase was in part driven by differential industry growth: industries employing relatively more skilled workers expanded more rapidly than those predominantly employing unskilled workers. Yet, even when holding the industry structure fixed at its 1990 level, the relative employment of skilled workers rose substantially, by 20%, as indicated by the gray line (see Appendix B.1 for a formal expression).

Panel B of Figure 1 indicates that the surge in the relative employment of skilled workers coincided with a rise in the overall log skill premium, measured as the difference between the average log wages of skilled and unskilled workers. The rise in the skill premium is observed regardless of whether we allow the industry composition by skill to vary over time or hold it constant at 1990 levels (see Appendix B.2 for a formal expression). Overall, the evidence is consistent with growing relative demand for skilled workers in Germany, as predicted by the traditional SBTC model.

Importance of Between-Establishment Wage Inequality Next, we document that the substantial increase in wage inequality observed in Germany from the mid-1990s onward primarily occurred between rather than within establishments. Panel C of Figure 1 highlights that the variance of individual log wages increased by around 33% over our two-decade study period—a pattern also documented by Dustmann et al. (2009, 2014). This increase is entirely driven by a widening of wage dispersion between establishments, while within-establishment inequality remained fairly stable over the same time period (see Appendix B.3 for details on the decomposition). These results align with existing evidence for Germany (e.g. Card et al., 2013).[6]

In Panel D, we further decompose the between-establishment variance of log wages into a between-industry and a within-industry component (see Appendix B.3 for details on the decomposition). More than half of the cumulative change over time in between-establishment wage inequality can be accounted for by rising wage differences across establishments within detailed industries. A similar pattern holds when we keep the industry structure constant at its 1990 level, indicating that the within-industry rise over time is not primarily due to the growing employment shares of industries with higher levels of between-establishment heterogeneity. These results indicate that establishments differ substantially in average wages, even within the same detailed industry, and that this heterogeneity is growing over time. In the remainder of the paper, we focus on these differences in pay across establishments within industries.[7]

Panels E and F provide evidence on the reasons why establishments within the same industry pay different average wages. Panel E shows that establishments vary in their skill composition relative to others in their industry, and increasingly so over time, as indicated by the rise in the variance of establishment-level log skilled-to-unskilled employment ratios between 1990 and 2010.[8] A corollary of this pattern is that skilled and unskilled workers increasingly sort into distinct establishments within industries, consistent with previous findings on worker-firm sorting (e.g. Barth et al., 2018; Card et al., 2013; Ferrer & Lluis, 2008; OECD, 2021). Finally, Panel F plots the variance of log wages across establishments within industries separately for skilled and unskilled workers (aggregated using the 1990 industry structure). The figure shows that establishments in the same industry differ in the wages they pay conditional on worker skill, and that skill-specific wages have become increasingly dispersed across establishments over time.

In sum, Panels E and F imply that the cross-sectional dispersion of average wages across establishments within industries, and the rise in this dispersion over time, are driven both by heterogeneity in the skill input mix across establishments and by heterogeneity in

[6] For related evidence for the U.S. and other OECD countries, see Barth et al. (2016), Song et al. (2019), and OECD (2021).

[7] 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. Their study finds that in the U.S., 25% of the increase in between-firm wage inequality over the past three decades occurred within 4-digit industries.

[8] The figure plots the within-industry variance of the log skilled-to-unskilled employment ratio among establishments with at least one worker of each skill type. We obtain qualitatively similar results if we use *all* establishments by either (i) imputing an employment level of one part-time worker (i.e. 0.5 full-time equivalent workers) for establishments with no workers of a given type, or (ii) using the variance of the establishment-level skilled employment share, which can range from zero to one and does not require any imputation.

establishment-level wages conditional on worker skill.[9]

4 Theoretical Framework

Guided by the motivating evidence presented in Section 3, we consider a theoretical framework in which firms within industries differ in terms of their skill composition and the wages that they pay conditional on skill type. We then study the impact of an aggregate SBTC shock (along the lines of Katz & Murphy, 1992) within this heterogeneous firm environment.[10]

Our model builds on the framework of Helpman et al. (2010), which adds search and matching frictions, match-specific abilities, and screening to the heterogeneous-firm environment of Melitz (2003), generating an equilibrium in which firms differ in their skill input mix as well as their pay. We outline the key features of the model below, and provide details in Appendix C.

4.1 Model Setup

Within each sector of the economy, 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 $0 < \beta < 1$.

As in Melitz (2003), firms that enter the market produce a single variety. Entry requires paying a fixed cost $f_e > 0$. After paying this fixed cost, the firm, denoted by f, draws an idiosyncratic productivity parameter θ_f from a common Pareto distribution $G_\theta(\theta)$ with scale parameter $\theta_{\min}$ and shape parameter z. Producing a variety requires a fixed cost $f_d > 0$. After observing their productivity draw, firms can decide whether to exit or produce.

There are two types of workers: skilled (s) and unskilled (u). The labor market features standard Diamond-Mortensen-Pissarides search and matching frictions, with firms required to pay a search cost of $b_\ell n_{\ell,f}$ to be matched with $n_{\ell,f}$ workers, $\ell = \{s, u\}$.[11] Skilled workers

[9] Existing research has further documented substantial dispersion in pay across firms conditional on much more granular individual controls (e.g. Card et al., 2013). We discuss the issue of within-skill worker heterogeneity in our empirical analysis in Section 5.

[10] While all of our empirical analysis is at the establishment level, in the theoretical framework we use the term 'firm', to parallel the theoretical literature in the field.

[11] Search and matching frictions in the labor market were first proposed by Mortensen (1970), Pissarides (1974), Diamond (1982b,a), and Mortensen & Pissarides (1994), and have subsequently been widely used in

are relatively scarce and hence command a higher search cost, that is, $b_s > b_u$.

Workers of a given skill type are ex-ante identical, but upon matching with a firm, they draw match-specific abilities from a Pareto distribution with shape parameter k and scale parameter a_{min}. Individual ability draws are not observable by the firm or the worker. However, a screening technology is available. By paying a screening cost $c\widetilde{a}_{\ell,f}^{\delta}/\delta$, the firm can determine whether a worker's ability exceeds a chosen threshold $\widetilde{a}_{\ell,f}$, where $\ell = \{s, u\}$, $c > 0$, and $\delta > 0$.[12]

The production function of firm f is given by:

$$y_f = \left[\left(\theta_f^{1+\alpha\mu_s}\mu_s\overline{a}_{s,f}h_{s,f}^{\gamma}\right)^{\nu} + \left(\overline{a}_{u,f}h_{u,f}^{\gamma}\right)^{\nu}\right]^{1/\nu}, \tag{1}$$

where $h_{s,f}$ and $h_{u,f}$ are the firm's input of skilled and unskilled labor (which differ from $n_{s,f}$ and $n_{u,f}$ because firms might not choose to hire all workers that they match with), $\overline{a}_{s,f}$ and $\overline{a}_{u,f}$ represent the average match-specific ability of skilled and unskilled workers hired by the firm, μ_s is an aggregate SBTC term, and ν, γ and α are parameters that are common across all firms, with $\nu > 0$, $0 < \gamma < 1$ and $\alpha \geq 0$.

This production function features a constant elasticity of substitution between the (effective) skilled and unskilled labor inputs, along the lines of the classic SBTC literature (e.g. Katz & Murphy, 1992; Autor et al., 1998), but allows for firm heterogeneity driven by the firm-specific productivity term, θ_f. This parameter generates skill-biased, rather than Hicks-neutral, productivity differences between firms. With Hicks-neutral productivity differences, all firms would optimally choose the same mix of skilled and unskilled workers, a pattern that is inconsistent with the empirical regularities highlighted in Section 3. For simplicity, we refer to firms with a higher value of θ_f as "more productive".

In addition to entering the production function as a skill-biased term, the parameter θ_f in Equation (1) features an exponent that depends on the aggregate state of technology μ_s and a parameter α. Setting $\alpha = 0$ yields a simpler production function that is identical to that in Section 5.2 of Helpman et al. (2010). While the equilibrium of the model is qualitatively identical in the case where $\alpha = 0$, a positive α will allow for a richer interaction between aggregate technological change (i.e. changes in μ_s) and firm-specific productivity differences, as we discuss in further detail in Section 4.3.

Note also that the production function features complementarity between workers' match-

---

theoretical and empirical work (e.g. Acemoglu, 1999; Cahuc et al., 2006; Manning & Petrongolo, 2017).

[12] These modeling choices are in line with a large empirical literature showing that worker ability cannot be costlessly observed (Gibbons & Katz, 1991; Altonji & Pierret, 2001; Schönberg, 2007; Carranza et al., 2022), and firms spend substantial resources to screen candidates in the hiring process (Barron et al., 1985; Autor & Scarborough, 2008; Hoffman et al., 2017). The functional form of the screening cost function captures the idea that screening at a higher ability threshold requires more complex, and hence costlier tests.

specific abilities ($\overline{a}_{s,f}$ and $\overline{a}_{u,f}$) and the firm type (θ_f). This gives an incentive for more productive firms to hire higher-ability workers by screening more intensively (i.e. by choosing a higher screening threshold) in equilibrium.[13]

The final key element of the model is the process by which wages are determined. As in Helpman et al. (2010), we assume that the firm and the workers engage in Stole & Zwiebel (1996a,b) bargaining. Since the screening process reveals only whether a worker's ability exceeds the threshold chosen by the firm—not the exact value of this ability—all hired workers of type ℓ in a given firm are treated as having the same expected ability, equal to $\overline{a}_{\ell,f}$, and hence they all receive the same wage.

To summarize, a firm that decides to pay the fixed cost of entry will observe its productivity and then decide whether to produce. If it decides to produce, it will choose the number of workers of each type to match with and the match-specific ability threshold to screen to. It makes these decisions in order to maximize profits, anticipating the outcome of the bargaining game.

4.2 Equilibrium

Firm-Level Outcomes The model yields closed-form expressions for firm-level employment and wages as a function of the firm-specific productivity parameter θ_f. Here we discuss the key qualitative predictions of the model, relegating the formal derivations to Appendix C. Since the motivating evidence as well as our subsequent empirical analysis considers firm outcomes in logarithmic terms (e.g. log employment or log skill-specific wages), and to maintain a tight link between the theoretical framework and the empirical application, we present a graphical illustration of the log-log relationships between θ_f and the various firm-level outcomes, derived from a model simulation. The simulation uses parameter values that are either standard in the literature or selected to satisfy key theoretical restrictions. Full details are provided in Appendix D.[14]

The solid lines in Figure 2 represent the equilibrium relationships for a baseline level of technology (i.e. a baseline value of the parameter μ_s). We defer the discussion of the dashed lines—which we use to illustrate the impacts of SBTC—to Section 4.3. Panel A of Figure 2 shows that firms that draw a higher value of θ_f will be larger in equilibrium, a

[13] As discussed by Helpman et al. (2010), this modeling feature reflects human capital complementarities, whereby a worker's productivity increases with the abilities of coworkers—a feature emphasized in Lucas (1978) and Rosen (1982), and more recently by Garicano (2000) and Moretti (2004). Note that the assumption that $0 < \gamma k < 1$ is required for firms to have an incentive to screen.

[14] Note that the simulation is not intended to target specific moments in the data; rather, it is used to provide a visual illustration of the qualitative predictions of the model. The simulation assumes that $\alpha > 0$; we discuss the role of α in detail at the end of Section 4.3.

standard feature of Melitz-type heterogeneous firm models. Panel B further illustrates that higher-productivity firms employ disproportionately more skilled workers, as measured by their log skilled-to-unskilled employment ratio (log skill ratio). For this outcome, the model yields the simple expression:

$$\ln\left(\frac{h_{s,f}}{h_{u,f}}\right) = -(1+\gamma\xi)\ln\left(\frac{b_s}{b_u}\right) + \xi\left(1+\alpha\mu_s\right)\ln\theta_f + \xi\ln\mu_s, \tag{2}$$

where ξ is a function of parameters of the model, such that $\xi > 0$.[15] Equation (2) shows that the firm-specific log skill ratio is decreasing in the relative search cost of skilled workers and increasing in the aggregate and firm-specific skill-biased technology parameters (μ_s and θ_f).

Next, we turn to wages. The equilibrium firm-specific log wage for unskilled workers can be expressed as:

$$\ln w_{u,f} = \ln w_{dr} + \left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}\ln\left[1+\varphi(\theta_f;\mu_s)\right], \tag{3}$$

where w_{dr} is a term that is common across firms, Γ is a function of parameters of the model such that $\Gamma > 0$, and $\varphi(\theta_f;\mu_s)$ is a function that is increasing in θ_f and μ_s and decreasing in the relative search cost for skilled workers.[16] As in Helpman et al. (2010), the nature of the relationship between firm productivity θ_f and the log wage of unskilled workers is governed by the sign of $\beta-\nu$. As argued by Helpman et al. (2010), and as we verify in Section 5, the empirically-relevant pattern corresponds to the case where $\beta > \nu$, such that unskilled wages are increasing in firm productivity. Panel C of Figure 2 provides a visual illustration of the relationship between $\ln w_{u,f}$ and $\ln\theta_f$.

The equilibrium log wage of skilled workers can be expressed as:

$$\ln w_{s,f} = \ln\left(\frac{b_s}{b_u}\right) + \frac{k}{\delta}\ln\varphi(\theta_f;\mu_s) + \ln w_{u,f}. \tag{4}$$

This expression is also increasing in θ_f, both directly through $\varphi(\theta_f;\mu_s)$ and indirectly through $\ln w_{u,f}$. Hence, in equilibrium, more productive firms will also pay higher wages to their skilled workers, as illustrated in Panel D of Figure 2.

Since more productive firms are more skill-intensive and pay higher wages to both types of workers, the average firm wage is increasing in productivity, as illustrated in Panel E.

[15] Specifically, $\xi \equiv \frac{\nu}{\Lambda}\left(1-\frac{k}{\delta}\right)$, where $\Lambda \equiv 1-\nu\gamma-\frac{\nu}{\delta}(1-\gamma k) > 0$. As discussed by Helpman et al. (2010), $\delta > k$ must be assumed for the model to match the well-documented positive empirical correlation between employer size and wage. Hence, $\left(1-\frac{k}{\delta}\right) > 0$ and $\xi > 0$.

[16] w_{dr} depends on model parameters as well as on an endogenously determined aggregate sectoral shifter (see Appendix C), $\Gamma \equiv 1-\beta\gamma-\frac{\beta}{\delta}(1-\gamma k) > 0$, and $\varphi(\theta_f;\mu_s) = \left(\theta_f^{1+\alpha\mu_s}\mu_s\right)^{\nu/\Lambda}\left(\frac{b_s}{b_u}\right)^{-\gamma\nu/\Lambda}$.

Finally, it follows from Equation (4) that the firm-specific log skill premium, $\ln\left(\frac{w_{s,f}}{w_{u,f}}\right)$, is increasing in productivity. We illustrate this in Panel F of Figure 2.

Sectoral Outcomes As is standard in this class of models, there is an endogenously determined productivity threshold θ_d, below which firms optimally choose not to operate. This threshold is pinned down by two conditions: a Zero-Cutoff Profit (ZCP) condition, requiring that the firm at the cutoff θ_d makes zero profits, and a Free Entry (FE) condition, requiring that the expected profits for a potential entrant equal the fixed entry cost (see Appendix C for details).

The firm-level outcomes discussed above depend on search costs b_s and b_u, which are proportional to the expected income of workers outside the sector. Helpman et al. (2010) discuss conditions under which these search costs are invariant to aggregate shocks (such as trade opening, in their setting). In what follows, we assume that such conditions hold, so that search costs remain unaffected by shocks such as SBTC.[17]

Summary and Intuition The model delivers an equilibrium in which operating firms that draw a higher value of θ_f employ relatively more skilled workers, pay higher wages conditional on skill type, and are larger.[18] Consistent with empirical evidence, this generates variation between firms in average log wages due to (i) differences in skill intensity and (ii) differences in wages for workers with a given level of skill.

Intuitively, differences in skill intensity between firms arise from the skill-biased nature of productivity in the production function, while between-firm wage differentials conditional on skill arise from the complementarity in the production function between workers' match-specific abilities and firm productivity θ_f. Screening is costly, but higher θ_f firms have an incentive to screen more intensively in order to restrict hiring to workers with higher match-specific abilities. In equilibrium, wages are bargained down to the replacement cost of the worker. Given that higher θ_f firms incur higher screening costs when searching for a new worker, their workers are more costly to replace and are therefore paid a higher wage.

This intuition makes clear why both heterogeneity in match-specific ability and the screening technology are crucial elements of the model—in addition to productivity differences between firms—in order to generate wage differences between firms for workers of a given skill type. If there were no match-specific abilities, firms would have no incentive to screen,

[17]See Helpman et al. (2010), Section 2.3, for full details on the determination of the sectoral variables and Section 6 of their paper for how they close the model in general equilibrium.

[18]As mentioned above, all of these equilibrium conditions would still hold in a simpler model where $\alpha = 0$. This parameter will only play a role in terms of the impact of an SBTC shock, as discussed in the next section.

and wages would be bargained down to the replacement cost of a worker, which would be the search cost b_ℓ which is common across firms for workers of a given type. Firms would be heterogeneous along the employment margin, but there would be no wage inequality between firms conditional on skill (as in Felbermayr et al., 2011). If screening were not feasible, the expected average worker ability would be common across all firms and equal to the average match-specific ability in the population. In this case, the bargaining process would lead to a uniform wage across firms for all workers of a given type, and once again, there would be no wage inequality conditional on skill. Although both the screening technology and heterogeneity in match-specific ability are necessary to generate wage differences between firms for workers of a given type, the ultimate source of between-firm wage inequality is differences in firm productivity. If firms exhibited uniform productivity—that is, if θ_f were common among all firms—all firms would pay the same wage.

4.3 Impacts of Skill-Biased Technological Change

We now consider how the model equilibrium is affected by skill-biased technological change (SBTC), modeled— as in the literature—as an exogenous increase in the parameter μ_s. Our goal is to trace how such a shock affects between-firm wage inequality. We thus focus on the *relative* effects of SBTC across worker and firm types, rather than the *absolute* effects on overall wage and employment levels.

We first consider how the equilibrium firm-level relationships described in the previous section and illustrated in Figure 2 are impacted by SBTC. Since it is not always straightforward to analytically obtain second-order derivatives of log firm outcomes with respect to log θ_f and μ_s, we rely on our simulation to characterize the model's predictions. The dashed lines in Figure 2 illustrate how the equilibrium log-log relationships are changed by SBTC (i.e. by an increase in μ_s).

Panel A shows that an increase in μ_s induces relative employment growth among high-productivity firms (which were already larger at baseline). Therefore, SBTC increases sectoral employment concentration. Intuitively, SBTC exacerbates the effective productivity differences between firms, due to the way in which θ_f and μ_s interact in the production function in Equation (1), leading to increased heterogeneity in firm size. This differential employment growth of high-productivity (and high-wage) firms will lead to an increase in worker-weighted between-firm wage inequality, all else equal.

Panel B shows that SBTC strengthens the link between log θ_f and the firms' log skill ratios. It can be easily inferred from Equation (2) that this will occur as long as $\alpha > 0$ (a point we return to below). This differential growth of log skill ratios implies a stronger sorting of skilled workers into more productive (high-wage) firms, which will exacerbate the average wage gap between low- and high-productivity firms.

Panels C and D of Figure 2 highlight that SBTC leads to wage increases for high-productivity firms relative to low-productivity firms, for both unskilled and skilled workers. This implies increased dispersion in log wages within skill groups, across firms. Intuitively, since SBTC exacerbates the effective productivity differences between firms, it induces higher-θ_f firms to increase both the quantity and the quality of the workers that they hire (given the complementarity between firm productivity and workers' match-specific abilities). The increase in worker quality is achieved by increasing the screening threshold, which leads to higher wages for workers. Together with the increase in sorting, the skill-specific wage changes within firms lead to an increase in the average log wage in more productive firms relative to less productive firms, as shown in Figure 2, Panel E. Finally, as illustrated in Panel F, the model further predicts that SBTC increases the log skill premium in all firms, with disproportionately large increases among high-productivity firms.

From the patterns in Figure 2, one can infer that if the set of operating firms remained constant after SBTC, wage inequality would increase due to the widening wage and employment differences between firms with low and high productivity levels. However, the set of operating firms will generally change in response to SBTC, as the endogenous productivity threshold below which firms choose not to operate (θ_d) shifts. As we show in Panel A of Appendix Figure A.1, the productivity threshold is increasing in μ_s. Intuitively, since SBTC disproportionately benefits high-productivity firms, those with a low value of θ_f cannot compete and therefore choose not to operate.

The remaining panels of Appendix Figure A.1 illustrate how various sector-level outcomes evolve as μ_s increases, after taking into account both the changes in outcomes conditional on θ_f (Figure 2) and the change in the composition of operating firms. As shown in Panel B, SBTC (i.e. an increase in μ_s) leads to an increase in the aggregate log skill premium—consistent with the key prediction of the standard SBTC model, and with the empirical pattern observed over time in Panel B of Figure 1.

Panel C of Figure A.1 shows that the (employment-weighted) variance of the log skill ratio between firms increases as μ_s increases, even after considering changes in the composition of operating firms. This is in line with the empirical patterns observed over time in Panel E of Figure 1. Finally, as shown in Panels D to F of Figure A.1, wage inequality between firms (overall and conditional on skill) increases as SBTC intensifies, in line with the empirical patterns observed over time in Panels D and F of Figure 1.

To summarize, the model highlights three reinforcing channels through which SBTC raises employment-weighted between-firm wage inequality: (i) the disproportionate employment growth of high-productivity (high-wage) firms, (ii) the stronger sorting of skilled workers to high-productivity firms, and (iii) the differential wage growth of skilled and unskilled workers in these firms.

The Role of α Figure 3 complements the findings above by tracking how the strength of the relationships between log productivity and the various log firm outcomes in Figure 2 evolves over a wide range of values of μ_s. This ensures that the observed steepening of the relationships observed in Figure 2 is systematic, rather than an artifact of comparing two particular values of μ_s. The figure also helps clarify the role of the α parameter in the model by reporting results separately for our baseline simulation, which uses $\alpha = 0.05$ (dashed lines), and for an alternative simulation in which we set $\alpha = 0$, holding all other model parameters constant (solid lines).

Each point on the lines represents the slope that would be obtained from a regression of the log firm outcome indicated in the panel title on log θ_f, for the given value of μ_s on the horizontal axis. When $\alpha = 0.05$, slopes rise monotonically with μ_s for all outcomes, confirming that the patterns in Figure 2 hold qualitatively for any pair of values of μ_s. In the alternative simulation with $\alpha = 0$, slopes are uniformly smaller, but still increase with μ_s for log employment and for the average firm log wage (overall and by type).[19] Thus, in the simpler case with $\alpha = 0$, SBTC continues to increase between-firm wage inequality due to differential growth in employment and differential growth in wages conditional on skill.

However, if $\alpha = 0$, SBTC does not differentially impact the log skill ratio or the log skill premium across firms (see the flat lines in Panels B and F of Figure 3). Thus, in this case, SBTC does not induce more dispersion in log skill ratios across firms—in contrast to what we observe empirically in Panel E of Figure 1.[20]

5 Empirical Evidence

In this section, we empirically validate the predictions of our theoretical framework regarding the workplace-level mechanisms through which SBTC influences wage inequality between establishments.

[19] This is because log employment and log wage levels are a function of $\ln[1+\varphi(\theta_f;\mu_s)]$, where an interaction term between θ_f and μ_s remains even if $\alpha = 0$.

[20] The series in Panel E of Figure 1 is constructed using employment weights for each firm. We have verified that the non-employment-weighted series is also rising strongly over time; hence, the increase over time is not (solely) driven by differential employment growth across workplaces.

5.1 Measuring Productivity

Our theoretical framework discusses the implications of SBTC on workplaces that differ in terms of their underlying productivity parameter θ_f—which is not directly observable in the data. We use two complementary approaches to obtain theoretically-consistent empirical proxies for the underlying productivity type of each establishment. First, using the structure of the model, we show that one can approximate log θ_f by normalizing the establishment-specific log skill ratio. As an alternative approach, we characterize each workplace according to its (log) employment level in 1990—a measure that is proportional to θ_f and is not affected by subsequent SBTC given its fixed baseline nature.

I. Proxy Derived from Establishment Skill Ratio: To obtain our first proxy for $\log \theta_f$, we start from Equation (2), which defines the log skill ratio for each firm in equilibrium. To simplify notation and to allow for differences in the SBTC term across industries and over time, we rewrite the equation as:

$$\ln s_{f(k)t} = -(1+\gamma\xi)\ln b + \xi(1+\alpha\mu_{kt})\ln\theta_{f(k)} + \xi\ln\mu_{kt}, \tag{5}$$

where k represents the industry that firm f operates in, $s_{f(k)t}$ denotes the firm-specific skill ratio at time t, b denotes the search cost ratio (b_s/b_u), and μ_{kt} reflects the level of SBTC, which is allowed to vary between industries (k) and over time (t).

Let Ω_k denote the set of continuing firms in industry k that are always in operation. We can write the mean (E) and the standard deviation (SD) of the log skill ratio at time t among the set of firms in Ω_k as:

$$E_{kt}\left[\ln s_{f(k)t} \mid \Omega_k\right] = -(1+\gamma\xi)\ln b + \xi(1+\alpha\mu_{kt})E_k\left[\ln\theta_{f(k)} \mid \Omega_k\right] + \xi\ln\mu_{kt},$$

$$SD_{kt}\left[\ln s_{f(k)t} \mid \Omega_k\right] = \xi(1+\alpha\mu_{kt})SD_k\left[\ln\theta_{f(k)} \mid \Omega_k\right],$$

where the conditional expectation and the conditional standard deviation on the right-hand side of these equations do not carry a time subscript as they capture the mean and standard deviation of log productivity among a fixed set of firms.

Consider a standardization of the log skill ratio of each firm relative to the mean and standard deviation among the set of firms in Ω_k. Given the equations above, this yields:

$$\frac{\ln s_{f(k)t} - E_{kt}\left[\ln s_{f(k)t} \mid \Omega_k\right]}{SD_{kt}\left[\ln s_{f(k)t} \mid \Omega_k\right]} = \frac{\ln\theta_{f(k)} - E_k\left[\ln\theta_{f(k)} \mid \Omega_k\right]}{SD_k\left[\ln\theta_{f(k)} \mid \Omega_k\right]} \equiv \hat{\theta}_{f(k)}. \tag{6}$$

By standardizing the firm-specific log skill ratio in this way, we recover—according to the model—a standardized measure of log θ_f that is not affected by SBTC (or any other aggregate or industry-level shocks). Intuitively, this transformation leverages the relative ranking of firms within industries in terms of their log skill ratio and, by standardizing, eliminates the influence of other shocks and parameters, thereby isolating the component of variation attributable solely to $\log \theta_f$.[21]

Although according to the model, the productivity proxy should be time-invariant for a given firm, in practice, idiosyncratic fluctuations at the firm level—for instance, due to short-term shocks or measurement error in observed skill ratios—as well as genuine changes in the underlying productivity of a firm over time, may introduce time variation in this measure. Empirically, we compute the average standardized log skill ratio from the left-hand side of Equation (6) for each establishment over rolling 5-year windows around each year t.[22] In this way, we smooth out short-term fluctuations in the establishment's log skill ratio while allowing for longer-run changes in $\log \theta_f$ not captured by the model. To reflect the empirical variation over time in this measure, we now denote it as $\hat{\theta}_{f(k)t}$.

II. Baseline Establishment Employment: As an alternative proxy for establishment productivity, we use log employment in 1990. This measure is positively related to $\log \theta_f$ and is not affected by subsequent SBTC (unlike any contemporaneous establishment-level measures that one might consider as proxies). However, when using this proxy, we must limit the sample to establishments observed in 1990, thereby excluding any establishments that enter into operation after this year. Moreover, to avoid the influence of changes in the composition of operating workplaces over time due to exiting establishments, we focus only on those that are active throughout the entire sample period when using this proxy.

Appendix Figure A.2 illustrates the relationships between log firm outcomes and baseline log employment obtained from our simulation, both at baseline and after SBTC (using the same parameter values as in Figure 2). All relationships are qualitatively the same as

[21]Note that we could allow b, α and ξ to be industry- and/or time-specific, and the standardization would still yield the same result. Note also that it is essential to standardize relative to a constant set of firms. Otherwise, the mean and standard deviation on the right-hand side of the equation would vary over time, due to changes in the composition of operating firms. Finally, note that $\hat{\theta}_{f(k)}$ can be recovered for *all* firms (including entering and exiting firms), even though the standardization is done relative to the set of continuing firms in Ω_k.

[22]To leverage the whole sample period, we use shorter windows at the boundaries: for 1990, we use the 1990–1992 average; for 1991, the 1990–1993 average; for 2009, the 2007–2010 average; and for 2010, the 2008–2010 average. Moreover, to be able to compute the log skill ratio (and the productivity proxy) for all operating establishments in all periods, establishments with zero workers of a given type in a given period are imputed to have one part-time worker (i.e. 0.5 full-time equivalent) of that type. The set of establishments in Ω_k is kept constant across all years. When computing the mean and the standard deviation of the log skill ratio among establishments in set Ω_k, we weight each establishment by its 1990 employment level.

in Figure 2, highlighting the validity of baseline employment as an alternative proxy for establishment productivity.

5.2 Cross-Sectional Relationships between Establishment Productivity Proxies and Other Workplace Outcomes

We begin by empirically validating the equilibrium cross-sectional relationships between the establishment's productivity and its employment and wage levels, as predicted by the model. To do so, we estimate a series of regression specifications of the following form:

$$x_{f(k)t} = \beta \widehat{\theta}_{f(k)t} + d_{kt} + \epsilon_{f(k)t}, \tag{7}$$

where x denotes a particular establishment-level outcome, d_{kt} are 3-digit industry-year fixed effects, and $\epsilon_{f(k)t}$ is an error term. β captures the within-industry association between $\widehat{\theta}_{f(k)t}$ and $x_{f(k)t}$ which, according to the model, should be positive for all of our outcomes of interest (log employment, log wages, the log skill ratio, and the log skill premium).

Column 1 of Panel A in Table 1 confirms that our model-consistent proxy, $\hat{\theta}_{f(k)t}$, is positively correlated with the establishment's log employment. As expected, given how the proxy is constructed, it is also highly correlated with the log skill ratio (Column 2). The productivity proxy is also positively associated with the average log wage in the establishment (Column 3). In Columns 4 to 6, we restrict the sample to establishments that employ at least one full-time worker of each skill type, and show that establishments with higher estimated productivity pay higher wages to both types of workers, but particularly so to their skilled workers. In consequence, the log skill premium in the establishment is increasing in productivity—in line with the model's predictions.

In Panel B of Table 1, we replicate the analysis using the establishment's baseline (1990) log employment as an alternative proxy for the underlying unobserved parameter log θ_f. As noted earlier, in this analysis, we restrict the sample to establishments that operate throughout the entire sample period. The results are consistent with those in Panel A and show that an establishment's skill intensity and its wages—both overall and by skill type—are significantly and positively correlated with its baseline size, as is its log skill premium.

Robustness: Controlling for Worker Heterogeneity within Skill Groups So far, our empirical approach has abstracted from worker heterogeneity within skill groups. This is consistent with our theoretical framework in which workers differ in match-specific ability, but these differences are not observable ex-ante and are not individually rewarded within firms. As such, in the model, workers of a given skill within a given firm are treated as homogeneous and perfectly substitutable.

In reality, workers differ along various dimensions—such as age, gender, immigration status, or occupation—for which returns may vary. If more productive firms tend to hire a higher share of workers with highly rewarded characteristics, then cross-firm differences in average wages by skill may in part reflect sorting along those dimensions rather than the mechanism emphasized in our model.

To assess the extent to which this type of sorting may influence our baseline results, we estimate residual wages by running year-specific regressions of individual log wages on controls for gender, nationality, a cubic function of age, and a full set of 3-digit occupation fixed effects. We then calculate the average residual wages by skill type in each establishment as an alternative measure of workplace pay. This approach implicitly treats workers as embodying different bundles of effective labor within each skill group, and assumes that these bundles are perfect substitutes in production. Establishment wage premia are thus allowed to differ across workers' broad skill groups (unskilled vs skilled), but not across other worker characteristics.

Columns 7 and 8 of Table 1 confirm that workplaces with higher values of our productivity proxies ($\widehat{\theta}_{f(k)t}$ and 1990 log establishment employment) pay higher wages to both skilled and unskilled workers, also when focusing on their residualized wages. These firms also display a larger log skill premium (computed based on residualized wages; Column 9). These findings support the theoretical hypothesis that more productive workplaces pay higher wages and pay a higher skill premium even when controlling for worker composition.

5.3 Establishment-Level Patterns Over Time

In the next step, we present empirical evidence consistent with the model's predictions regarding the workplace-level adjustments induced by SBTC. Specifically, as illustrated in Figures 2 and 3, the model predicts that SBTC will strengthen the cross-sectional relationships between firms' underlying productivity parameters (log θ_f) and their log employment, their skill intensities, and their log wages—both overall and conditional on worker skill.

Assuming that SBTC has been an ongoing phenomenon between 1990 and 2010, this implies that the cross-sectional relationships in Equation (7) should become stronger over time. To test this prediction, we modify our previous estimation specification by now allowing for a year-specific coefficient on $\widehat{\theta}_{f(k)t}$; that is:

$$x_{f(k)t} = \beta_t \widehat{\theta}_{f(k)t} + d_{kt} + \epsilon_{f(k)t}. \tag{8}$$

Figure 4 plots the coefficient estimates for β_t using our model-consistent proxy for establishment productivity, $\hat{\theta}_{f(k)t}$. This figure represents the data counterpart of Figure 3. Each panel considers a different outcome variable. Panel A shows that, while the establishment productivity proxy was not significantly associated with employment in the early 1990s, this association increased to more than 0.1% in 2010. This finding suggests that more productive firms have grown faster than less productive firms within the same industry, leading to increased employment concentration. Panel B shows a similar strengthening of the relationship between productivity and the log skill ratio, pointing toward increased sorting of skilled workers into high-productivity establishments. Meanwhile, Panels C and D indicate that the association between $\hat{\theta}_{f(k)t}$ and log establishment wages by worker type grew stronger over the time period. Focusing on average log wages in the establishments in Panel E, the regression coefficient doubled from slightly below 0.05 in the early 1990s to nearly 0.1 by 2010. These findings highlight that wages—both overall and conditional on worker skill—increased relatively more in high-productivity establishments, compared to low-productivity establishments within the same industry, leading to more wage inequality between high- and low-productivity establishments within industries. It should be noted that the strengthening of the relationships between establishment productivity and wages was similar in magnitude for skilled and unskilled workers. As a result, establishments' log skill premia have not become increasingly heterogeneous across establishments with different productivity levels over time (Panel F). One explanation for this finding is that unions and work councils restrict establishments in Germany from differentially adjusting the wages of their skilled and unskilled workers.[23]

Appendix Figure A.3 shows the corresponding results when using the establishment's log employment level in 1990 rather than $\hat{\theta}_{f(k)t}$ as a proxy for productivity in the estimation of Equation (8). Here, we restrict the sample to establishments that remain in operation throughout the entire sample period. The findings corroborate those of Figure 4: The relationships between baseline log employment and the log skill ratio, as well as log wages (overall and conditional on skill), have become stronger over time within industries.

Figure 5 highlights that we obtain qualitatively similar results when using residualized log wages (rather than raw log wages) by skill for each establishment. The three panels on the left focus on our model-consistent productivity proxy derived from establishments' log skill ratios, $\hat{\theta}_{f(k)t}$, while the three panels on the right use baseline (1990) log establishment employment as a proxy for productivity and restrict the sample to a balanced panel of

[23] For example, union agreements typically stipulate the same relative wage increases for workers with different skill levels. It is worth noting that a constant relationship between the establishment's productivity and its log skill premium over time is also consistent with the model if α in Equation (1) is equal to zero (see Panel F of Figure 3). However, in this case, we would expect a constant relationship also between the establishment's productivity and its log skill ratio—contrary to what we find in Panel B.

workplaces operating throughout the entire sample period. Regardless of which productivity proxy we consider, we find that the associations between productivity and establishment-level wages by skill are increasing over time. We also find evidence of a slight strengthening of the relationship between productivity and the log skill premium when focusing on the log skill premium in residual wages.

To summarize, the evidence in Figures 4 and 5 supports the key predictions of the model: With ongoing SBTC, high-productivity workplaces grow disproportionately; they increasingly employ skilled workers; and they increasingly pay higher wages to both skilled and unskilled workers.

5.4 Industry-Level Variation in Technology Exposure

To provide more direct evidence on the role of skill-biased technological change in driving the patterns that we have documented, we leverage variation in exposure to SBTC across industries. We first focus on the adoption of capital in information and communication technologies (ICT)—technologies commonly interpreted as skill-biased—as our measure of industry-level SBTC exposure. We classify industries according to whether they experienced above- or below-median changes in their ICT capital stock per worker between 1991 and 2007, as measured in the EUKLEMS dataset.[24] We compute industry-level outcomes within each 3-digit industry, and aggregate to these two broad groups using 1990 industry employment shares as weights.

Panels A and B of Figure 6 confirm that industries with above-median ICT adoption exhibit larger increases in their overall log skill ratios and log skill premia over time, consistent with the standard predictions of the SBTC model. Next, we consider the evolution of between-establishment wage inequality within each of the two industry groups. According to the model, we expect larger increases in between-establishment wage inequality in industries that adopt more ICT—both because of increased clustering of skilled workers in high-productivity establishments, and because of increased wage dispersion across establishments conditional on worker skill.

Panels C, E and F of Figure 6 indicate that high-adoption industries experienced a larger rise in between-establishment wage inequality from 1990 to 2010 compared to low-adoption industries, both overall (Panel C) and when separately considering skilled and unskilled wages (Panels E and F). Furthermore, Panel D shows that establishments became more dispersed

[24]We use 1991 and 2007 as these are the earliest and latest years, respectively, that are available in the EUKLEMS data. The median change is computed using 1990 industry employment weights (so that roughly 50% of 1990 employment is in industries with above-median changes and 50% in industries with below-median changes).

in their log skill ratios over time, particularly in high-adoption industries, suggesting an increased clustering of skilled workers within the same establishments in these industries. This evidence is consistent with our model predictions.

To further corroborate the validity of the workplace-level mechanisms highlighted by the model, we investigate whether the strengthening of the relationships between establishment productivity and other establishment outcomes over time is more pronounced in industries more exposed to SBTC. We do this by modifying Equation (8) to allow the coefficients β_t to vary also by industry group K (where K indicates whether the establishment's industry k is a high- or a low-exposure industry); that is:

$$x_{f(k)t} = \beta_{Kt}\widehat{\theta}_{f(k)t} + d_{kt} + \epsilon_{f(k)t}. \tag{9}$$

Figure 7 plots the change in the estimated β_{Kt} coefficients relative to 1990 separately for low-exposure industries (the light gray markers) and high-exposure industries (the black markers); that is, it plots estimates of $\beta_{Kt} - \beta_{K1990}$ for each industry group K and year t. Panel A shows that the association between our model-consistent proxy for establishment productivity ($\widehat{\theta}_{f(k)t}$) and employment has strengthened mainly in high-exposure industries, with little change observed in low-exposure industries. The relationship between $\widehat{\theta}_{f(k)t}$ and the log skill ratio also increased more in high-exposure industries (Panel B), suggesting that the sorting of skilled workers into high-productivity firms intensified somewhat more in high-exposure than in low-exposure industries. Panels C to E show that, while the relationship between establishment productivity and establishment wages, overall and conditional on worker skill, increased also in low-exposure industries, the increase was stronger in industries with higher levels of technology exposure. Finally, the association between establishment productivity and the log skill premium has remained roughly stable over time in both low- and high-exposure industries (Panel F).

The findings in Figure 7 provide strong empirical support for the predictions of the model. Specifically, the evidence indicates that the key channels identified by our model—namely, the differential employment and wage growth of high-productivity establishments compared to low-productivity establishments, and the increased clustering of high-skilled workers into high-productivity establishments—are more pronounced in industries with greater exposure to SBTC.

To summarize the changes over time in the strength of the relationships between different establishment outcomes and our model-consistent proxy of establishment productivity ($\hat{\theta}_{f(k)t}$) in high-exposure versus low-exposure industries, we estimate a modified version of Equation

(9) using data for 1990 and 2010 only. Specifically, we estimate:

$$x_{f(k)t} = \beta_0 \widehat{\theta}_{f(k)t} + \beta_1 \mathbf{1}_{(k \in H)} \widehat{\theta}_{f(k)t} + \beta_2 \mathbf{1}_{(t=2010)} \widehat{\theta}_{f(k)t} \\ + \beta_3 \mathbf{1}_{(t=2010)} \mathbf{1}_{(k \in H)} \widehat{\theta}_{f(k)t} + d_{kt} + \epsilon_{f(k)t}, \tag{10}$$

where $\mathbf{1}_{()}$ is an indicator function and H refers to the high-exposure industry group. Here, β_0 captures the slope of the relationship between productivity and the outcome of interest for the low-exposure industry group in 1990; β_1 captures the slope differential for the high-exposure industry group (relative to the low-exposure group) in 1990; β_2 captures the change in the slope between 1990 and 2010 for the low-exposure industry group (corresponding to the gray marker for 2010 in Figure 7); and β_3 captures the additional change in the slope between 1990 and 2010 for the high-exposure group (corresponding to the gap between the gray and the black markers for 2010 in Figure 7).

In Table 2 we report estimates for β_2 and β_3 from Equation (10). The coefficients in the first row of each panel are the estimates for β_2, reflecting the change over time in the strength of the relationships between productivity and the various outcomes of interest for firms in low-exposure industries. The consistently positive coefficients in Panel A imply that the associations between all relevant establishment outcomes—including employment, raw and residualized wages (overall and by skill), and the log skill ratio—and establishment productivity have strengthened between 1990 and 2010 in low-exposure industries. Hence, the gaps in employment, wages, and skill intensities between low-productivity and high-productivity establishments widened also in low-exposure industries, likely because these industries were also exposed somewhat to SBTC.

More importantly, however, the coefficients in the second row show the estimates for β_3, i.e. the *differential* change in high-exposure industries beyond what is observed for low-exposure industries. In Panel A, all coefficients (with the exception of the log skill premium in Columns 6 and 9) are positive and statistically significant. Therefore, the employment and wage gaps between low- and high-productivity establishments widened more in industries more exposed to SBTC, in line with the patterns observed in Figure 7. The sorting of skilled workers into high-productivity establishments also increased more in high-exposure industries.

The industries that we have classified as "high exposure" may have also been more exposed to other shocks, such as trade or offshoring, which may have impacted these associations. To rule out this possibility as much as possible, we re-estimate the regressions in Equation (10), but add controls for industry-level measures of trade exposure and offshorability, each interacted with establishment productivity and year fixed effects. Our measure of trade exposure captures changes in industry-level exports and imports per worker to and from China and Eastern Europe between 1990 and 2010.[25] Our offshorability measure draws on data provided by Goos et al. (2014) on occupation-level offshorability and aggregates this up to the industry level using each industry's occupational structure in 1990.

The results in Panel B show that adding these controls has little impact on our coefficients of interest. Therefore, these findings corroborate the importance of SBTC as a driver of the observed patterns, even conditional on other shocks at the industry level.

In Panels C and D, we adopt an alternative approach to classify industries according to their exposure to SBTC. Specifically, given the standard prediction that SBTC leads to a higher skill premium, we infer exposure to SBTC indirectly for each industry based on the change over time in their sector-wide (average) skill premium. Industries with increases in the skill premium above the median between 1990 and 2010 are categorized as "high exposure". The results based on this alternative measure consistently point in the same direction as those based on ICT adoption.

In summary, these findings align with the model predictions and support the hypothesis that SBTC is an important driver of the observed increase in between-establishment wage inequality. Specifically, we find that the channels highlighted in our model through which SBTC affects wage inequality between establishments—namely, larger increases in employment, skill intensities, and skill-specific wages in high-productivity establishments relative to low-productivity establishments—are indeed more prevalent in industries more exposed to SBTC.

6 Conclusions

In this paper, we link skill-biased technological change both theoretically and empirically to the rise in between-establishment wage inequality. Although an extensive literature has explored the role of skill-biased technological change for wage inequality, it has focused on frameworks with aggregate production functions and perfectly competitive labor markets. Consequently, this framework has had implications only for wage differences *between* workers from different skill groups. However, empirically, a major component of the increase in wage inequality is observed *within* skill groups, between establishments within industries.

By embedding a skill-biased technological change shock within a rich yet tractable heterogeneous firm framework, we show that this type of shock leads to heterogeneous responses at the firm level, thereby generating a rise in between-firm wage inequality. Using detailed

[25] We use the trade data from Dauth et al. (2017), who crosswalk the product-level information from the UN Comtrade Database to the industry level.

administrative social security data from Germany from 1990 to 2010, we document several novel empirical patterns at the establishment level and show that they are consistent with the model's predictions.

The model highlights that an industry-wide skill-biased technological change shock will increase between-firm wage inequality through three main channels: differential employment growth, whereby workers relocate towards more productive and more skill-intensive firms; increased sorting and segregation, whereby skilled workers increasingly cluster together in the same firms and increasingly sort into more productive firms that pay higher wages and are larger; and differential wage growth, whereby establishment wage premia conditional on worker skill increase more for more productive, larger, more skill-intensive firms. We find that all of these channels are empirically relevant.

We also provide evidence that the key establishment-level patterns we identify are more pronounced in industries more exposed to technological change, even after controlling for trade shocks and industry-level offshorability.

Our results point to the importance of moving beyond the traditional framework with an aggregate production function and competitive labor markets when assessing the impact of aggregate shocks such as technological change. While the literature has generally focused on how the effects of technological change vary according to the skills that an individual possesses or the tasks that they perform (e.g. Autor & Dorn, 2013; Cortes, 2016), our findings indicate that the type of firm that an individual is matched to is at least as important. Unskilled workers employed in low-productivity firms lose out not only relative to skilled workers in these firms, but also relative to unskilled workers in high-productivity firms. Understanding which types of policies can mitigate the negative impacts of technological change on some groups of workers, within a more realistic environment with heterogeneous firms and various market frictions, such as the one in this paper, remains a crucial avenue for future research.

References

Acemoglu, D. (1999). Changes in unemployment and wage inequality: An alternative theory and some evidence. *American Economic Review*, 89(5), 1259–1278.

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.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2020). What are the labor and product market effects of automation? New evidence from France. *Working Paper*.

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

Altonji, J. G. & Pierret, C. R. (2001). Employer learning and statistical discrimination*. *The Quarterly Journal of Economics*, 116(1), 313–350.

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. (2013). The growth of low skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, 103(5), 1553–1597.

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., Katz, L. F., & Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics*, 113(4), 1169–1214.

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.

Autor, D. H. & Scarborough, D. (2008). Does job testing harm minority workers? evidence from retail establishments*. *The Quarterly Journal of Economics*, 123(1), 219–277.

Azar, J., Marinescu, I., & Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, 57(S), S167–S199.

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

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

Barron, J. M., Bishop, J., & Dunkelberg, W. C. (1985). Employer search: The interviewing and hiring of new employees. *The Review of Economics and Statistics*, 67(1), 43–52.

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.

Barth, E., Davis, J., & Freeman, R. B. (2018). Augmenting the human capital earnings equation with measures of where people work. *Journal of Labor Economics*, 36(S1), S71–S97.

Bessen, J., Goos, M., Salomons, A., & van den Berge, W. (2020). Firm-level automation: Evidence from the Netherlands. *AEA Papers and Proceedings*, 110, 389–393.

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

Blossfeld, H.-P. (1987). Labor-market entry and the sexual segregation of careers in the federal republic of Germany. *American Journal of Sociology*, 93(1), 89–118.

Bonfiglioli, A., Crinò, R., Fadinger, H., & Gancia, G. (2020). Robot imports and firm-level outcomes. *Working Paper*.

Cahuc, P., Postel-Vinay, F., & Robin, J.-M. (2006). Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2), 323–364.

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.

Carranza, E., Garlick, R., Orkin, K., & Rankin, N. (2022). Job search and hiring with limited information about workseekers' skills. *American Economic Review*, 112(11), 3547–83.

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. (2023). Rising concentration and wage inequality. *The Scandinavian Journal of Economics*, Forthcoming.

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. (2017). Trade and manufacturing jobs in germany. *American Economic Review: Papers and Proceedings*, 107(5), 337–42.

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6), 3104–3153.

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.

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.

Ferrer, A. & Lluis, S. (2008). Should workers care about firm size? *Industrial and Labor Relations Review*, 62(1), 104–125.

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.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5), 874–904.

Ghironi, F. & Melitz, M. J. (2005). International trade and macroeconomic dynamics with heterogeneous firms. *Quarterly Journal of Economics*, 120(3), 865–915.

Gibbons, R. & Katz, L. F. (1991). Layoffs and lemons. *Journal of Labor Economics*, 9(4), 351–380.

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.

Hoffman, M., Kahn, L. B., & Li, D. (2017). Discretion in hiring. *The Quarterly Journal of Economics*, 133(2), 765–800.

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.

Koch, M., Manuylov, I., & Smolka, M. (2021). Robots and firms. *The Economic Journal*, 131, 2553–2584.

Kramarz, F., Lollivier, S., & Pele, L.-P. (1996). Wage inequalities and firm-specific compensation policies in France. *Annales d'Economie et de Statistique*, 41/42, 369–386.

Lindenlaub, I. (2017). Sorting multidimensional types: Theory and application. *The Review of Economic Studies*, 84(2), 718–789.

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

Lucas, R. E. (1978). On the size distribution of business firms. *The Bell Journal of Economics*, 9(2), 508–523.

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.

Manning, A. & Petrongolo, B. (2017). How local are labor markets? evidence from a spatial job search model. *American Economic Review*, 107(10), 2877–2907.

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.

Moretti, E. (2004). Workers' education, spillovers, and productivity: Evidence from plant-level production functions. *American Economic Review*, 94(3), 656–690.

Mortensen, D. T. (1970). A theory of wage and employment dynamics. In E. S. Phelps & others (Eds.), *The Microeconomic Foundations of Employment and Inflation Theory*. New York: W. W. Norton & Company.

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.

OECD (2021). The role of firms in wage inequality: Policy lessons from a large scale cross-country study. *OECD Publishing, Paris*.

Pissarides, C. A. (1974). Risk, job search, and income distribution. *Journal of Political Economy*, 82(6), 1255–1267.

Rinz, K. (2022). Labor market concentration, earnings, and inequality. *Journal of Human Resources*, 57(S), S251–S283.

Rosen, S. (1982). Authority, control, and the distribution of earnings. *The Bell Journal of Economics*, 13(2), 311–323.

Schönberg, U. (2007). Testing for asymmetric employer learning. *Journal of Labor Economics*, 25(4), 651–691.

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.

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. *American Sociological Review*, 86(6), 1100–1130.

Figure 1: Aggregate Patterns over Time

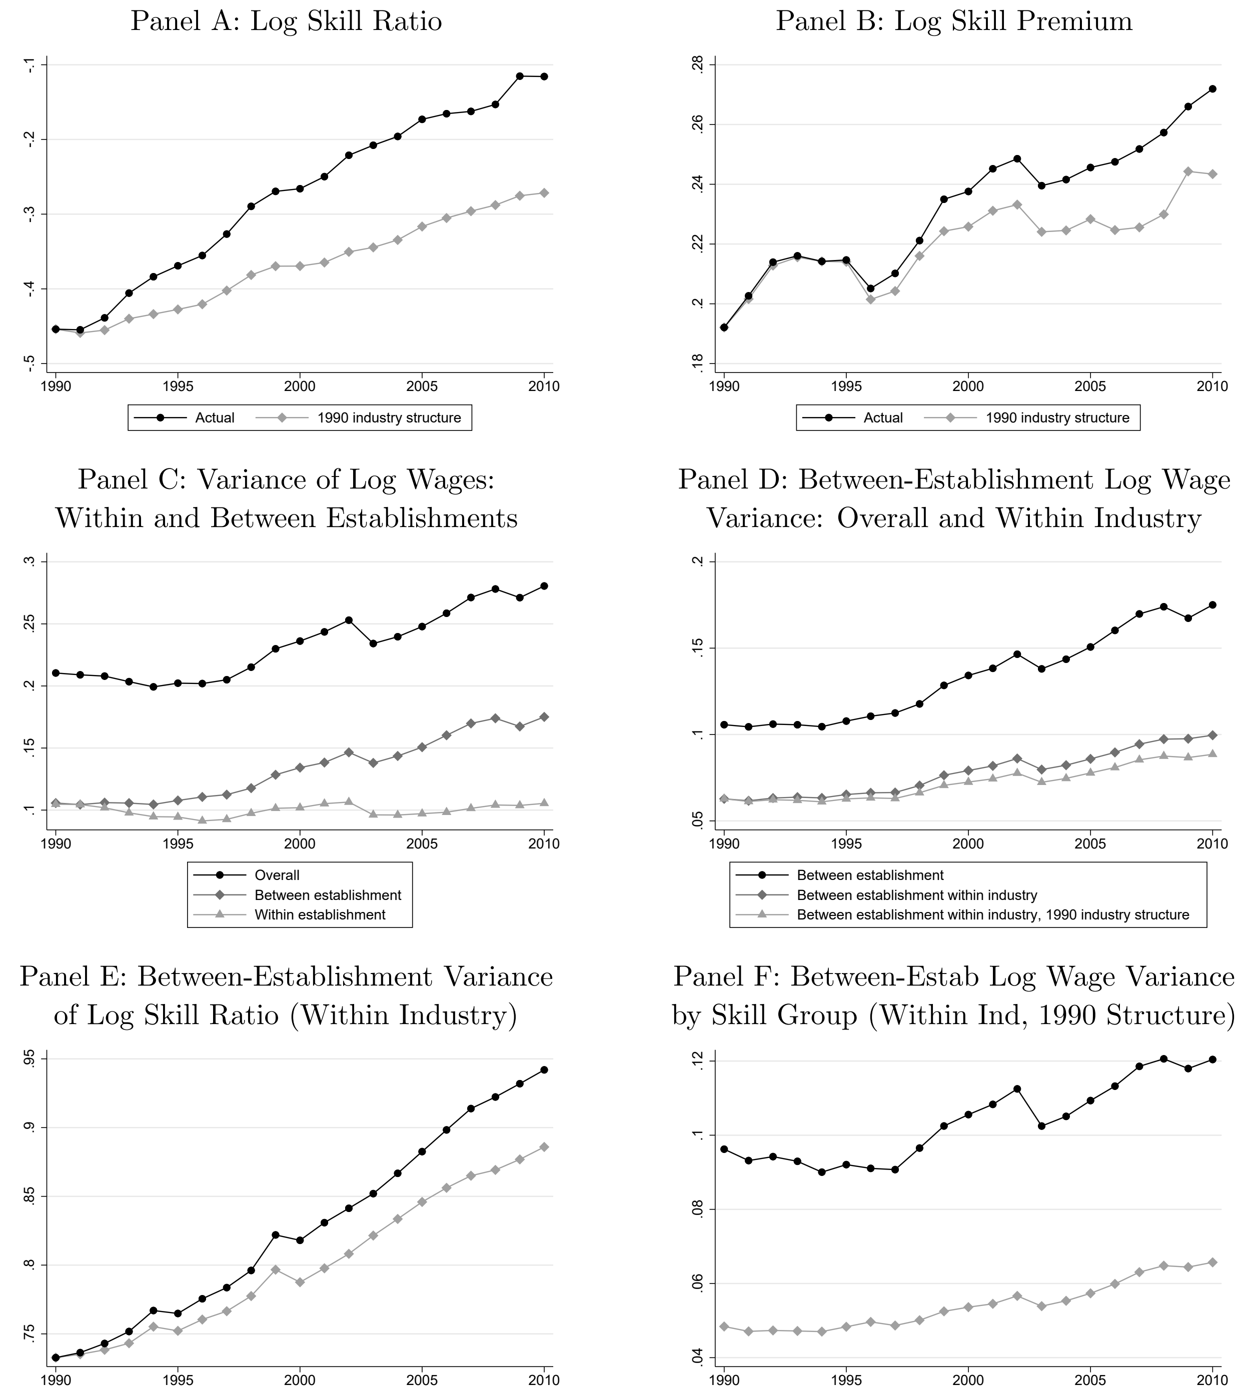

Note: Panel A shows the evolution of log relative employment of skilled workers in Germany between 1990 and 2010. Panel B plots the evolution of the skilled wage premium, calculated as the difference between the average log wage of full-time skilled and unskilled workers in each year. Panel C displays the overall variance of individual log wages, along with its between- and within-establishment components. Panel D displays the variance of establishment-level average log wages overall and within industries. Panel E shows the within-industry variance of establishments' log skill ratios (among establishments with at least one worker of each skill type). Panel F shows the variance in log wages across establishments within industries, aggregated using the 1990 industry structure, for skilled and unskilled workers, respectively. See Appendix B for details on how the series in this figure are computed.

Figure 2: Simulation Results: Firm-Level Outcomes at Baseline and after an SBTC Shock

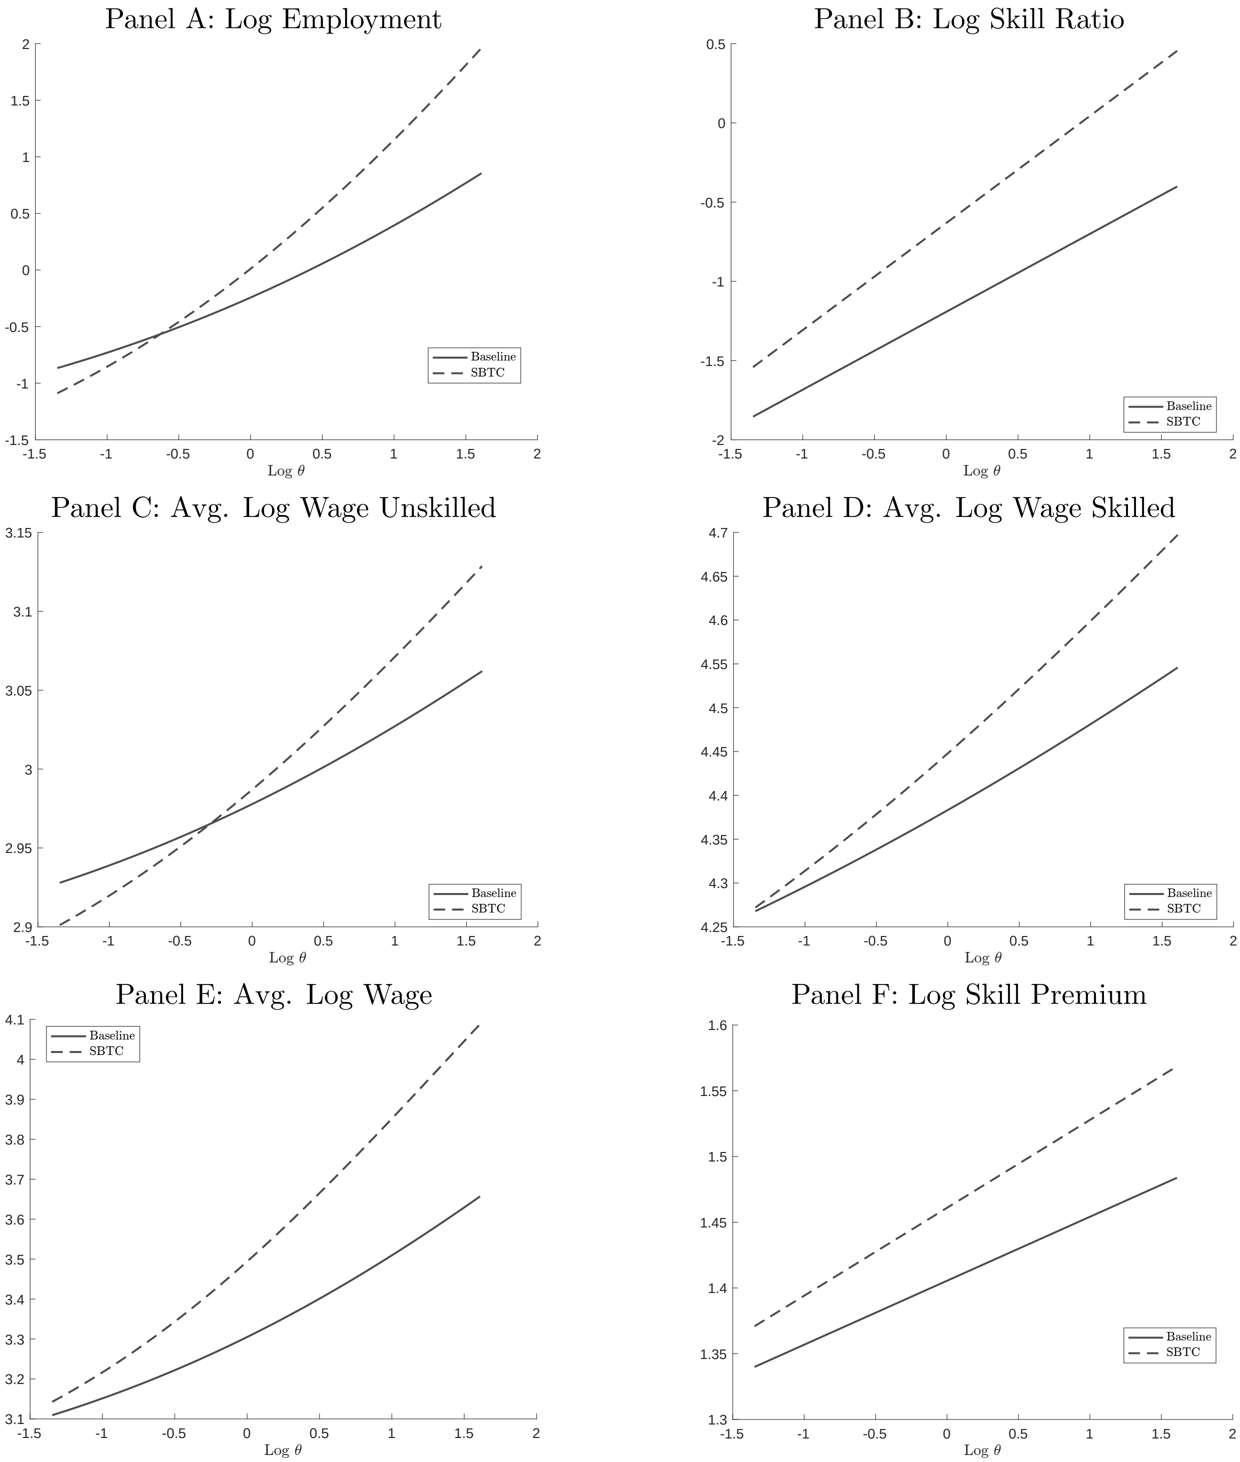

Note: The figure shows the simulated relationship between firm-level outcomes (as indicated in each panel title) and the firm log productivity parameter (log θ). The solid line corresponds to a low value of the aggregate technology parameter μ_s, representing a baseline relationship before an SBTC shock. The dashed line corresponds to a higher value of μ_s, capturing the effect of the shock. Simulation details are provided in Appendix D.

Figure 3: Simulated Regression Slopes of Log Firm Outcomes on Log Productivity across SBTC Levels

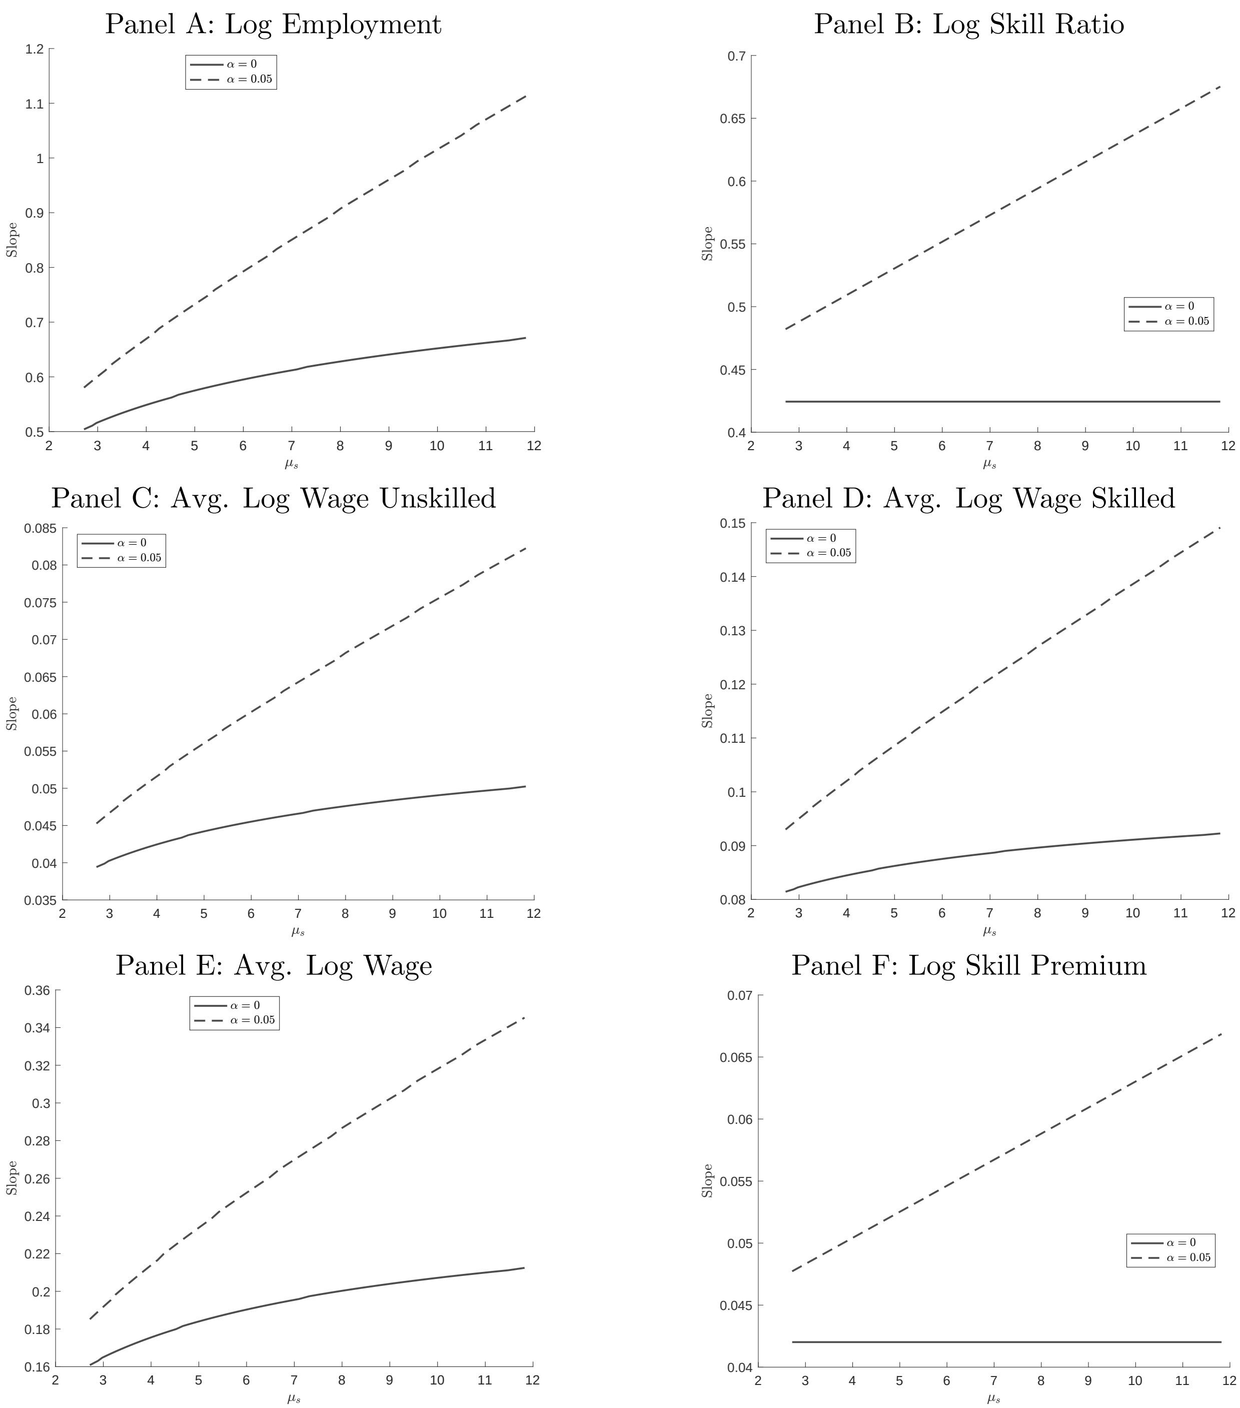

Note: The figure displays the simulated relationship between firm outcomes (as indicated in each panel title) and the firm log productivity parameter (log θ) across different values of the SBTC parameter, μ_s. Each point on the lines corresponds to the slope from a regression of the firm outcome on log θ, at the associated value of μ_s. The dashed line shows results from the baseline model with $\alpha = 0.05$ in the production function, while the solid line corresponds to results with $\alpha = 0$. Simulation details are provided in Appendix D.

Figure 4: Year-by-Year Associations between $\hat{\theta}_{f(k)t}$ and Other Establishment Characteristics

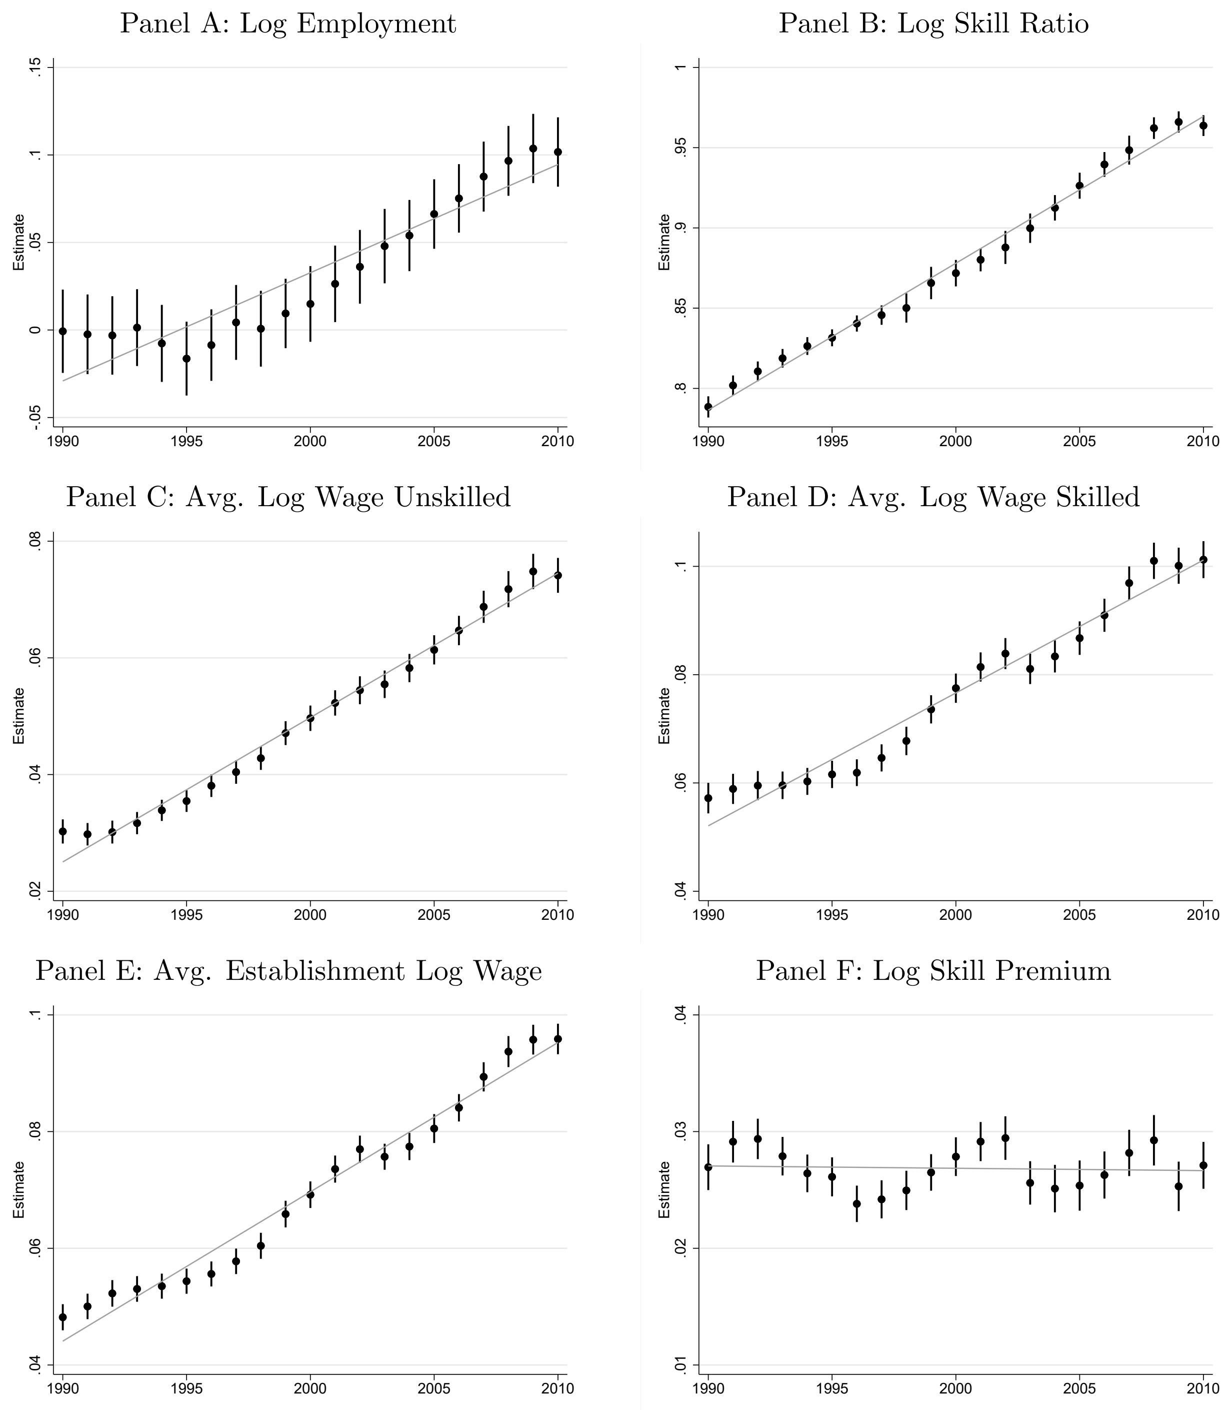

Note: The figure shows estimated coefficients and 95% confidence intervals from regressions of the outcome indicated in each panel title on $\hat{\theta}_{f(k)t}$, our model-consistent proxy for establishments' underlying productivity $\log \theta_f$ (see Section 5.1 for construction details). Coefficients are allowed to vary by year, with all regressions controlling for fully interacted 3-digit industry×year fixed effects. Estimates are based on BeH data, and observations are weighted by establishment size. Standard errors are clustered at the establishment level.

Figure 5: Associations between Establishment Productivity Proxies and Residualized Wages

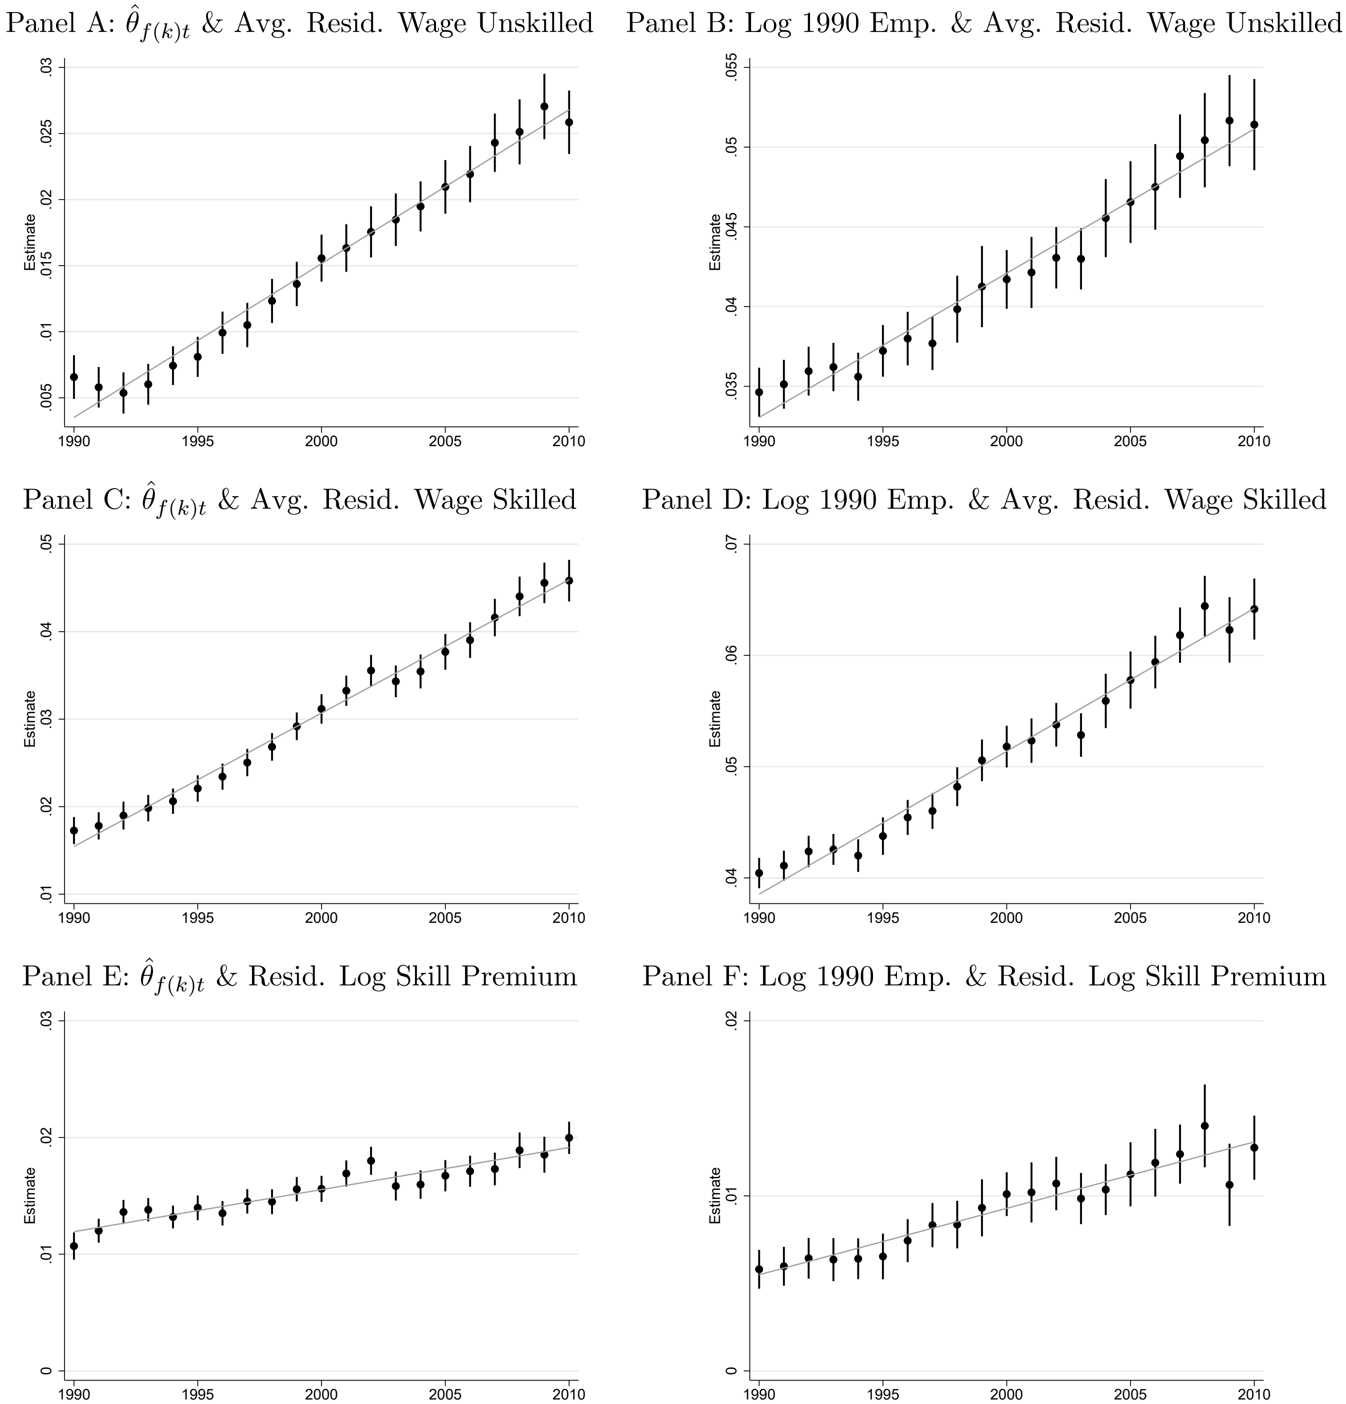

Note: Panels A, C and E show estimated coefficients and 95% confidence intervals from regressions of the outcome indicated in each panel title on $\hat{\theta}_{f(k)t}$, our model-consistent proxy for establishments' underlying productivity $\log \theta_f$ (see Section 5.1 for construction details). Panels B, D and F show analogous estimates using baseline establishment log employment in 1990 as the key regressor of interest. Coefficients are allowed to vary by year, with all regressions controlling for fully interacted 3-digit industry×year fixed effects. Estimates are based on BeH data—all establishments in Panels A, C and E, and only establishments that remain in operation throughout the entire 1990–2010 period in Panels B, D and F. Observations are weighted by establishment size—current size in Panels A, C and E, and baseline (1990) size in Panels B, D and F. Standard errors are clustered at the establishment level.

Figure 6: Average Change in Industry-Level Outcomes for Low and High Technology Exposure Industries

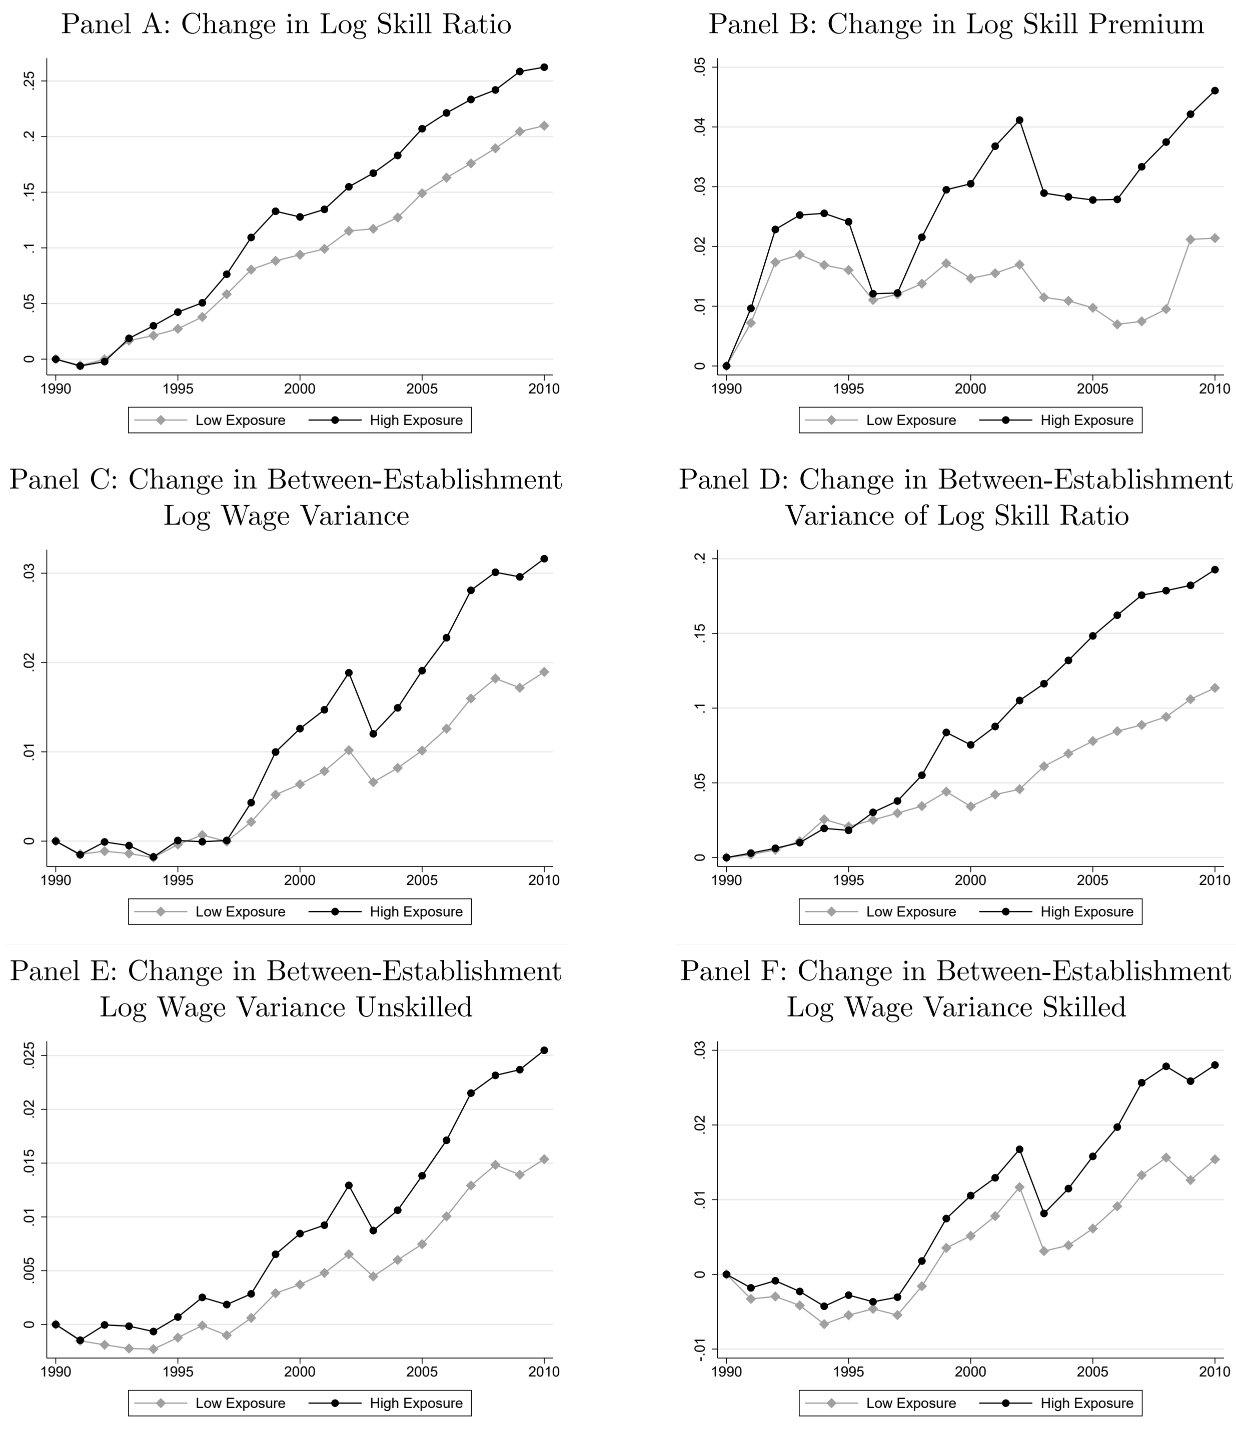

Note: The figure displays the evolution of average industry-level outcomes (as indicated in each panel title) over time relative to 1990, separately for low-exposure and high-exposure industries. High-exposure industries are those with above-median changes in their ICT capital stock per worker between 1991 and 2007, as measured in the EUKLEMS data. Industry-level outcomes are aggregated using 1990 industry employment shares as weights.

Figure 7: Changes over Time in the Associations between $\hat{\theta}_{f(k)t}$ and Other Establishment Characteristics for Low and High Technology Exposure Industries

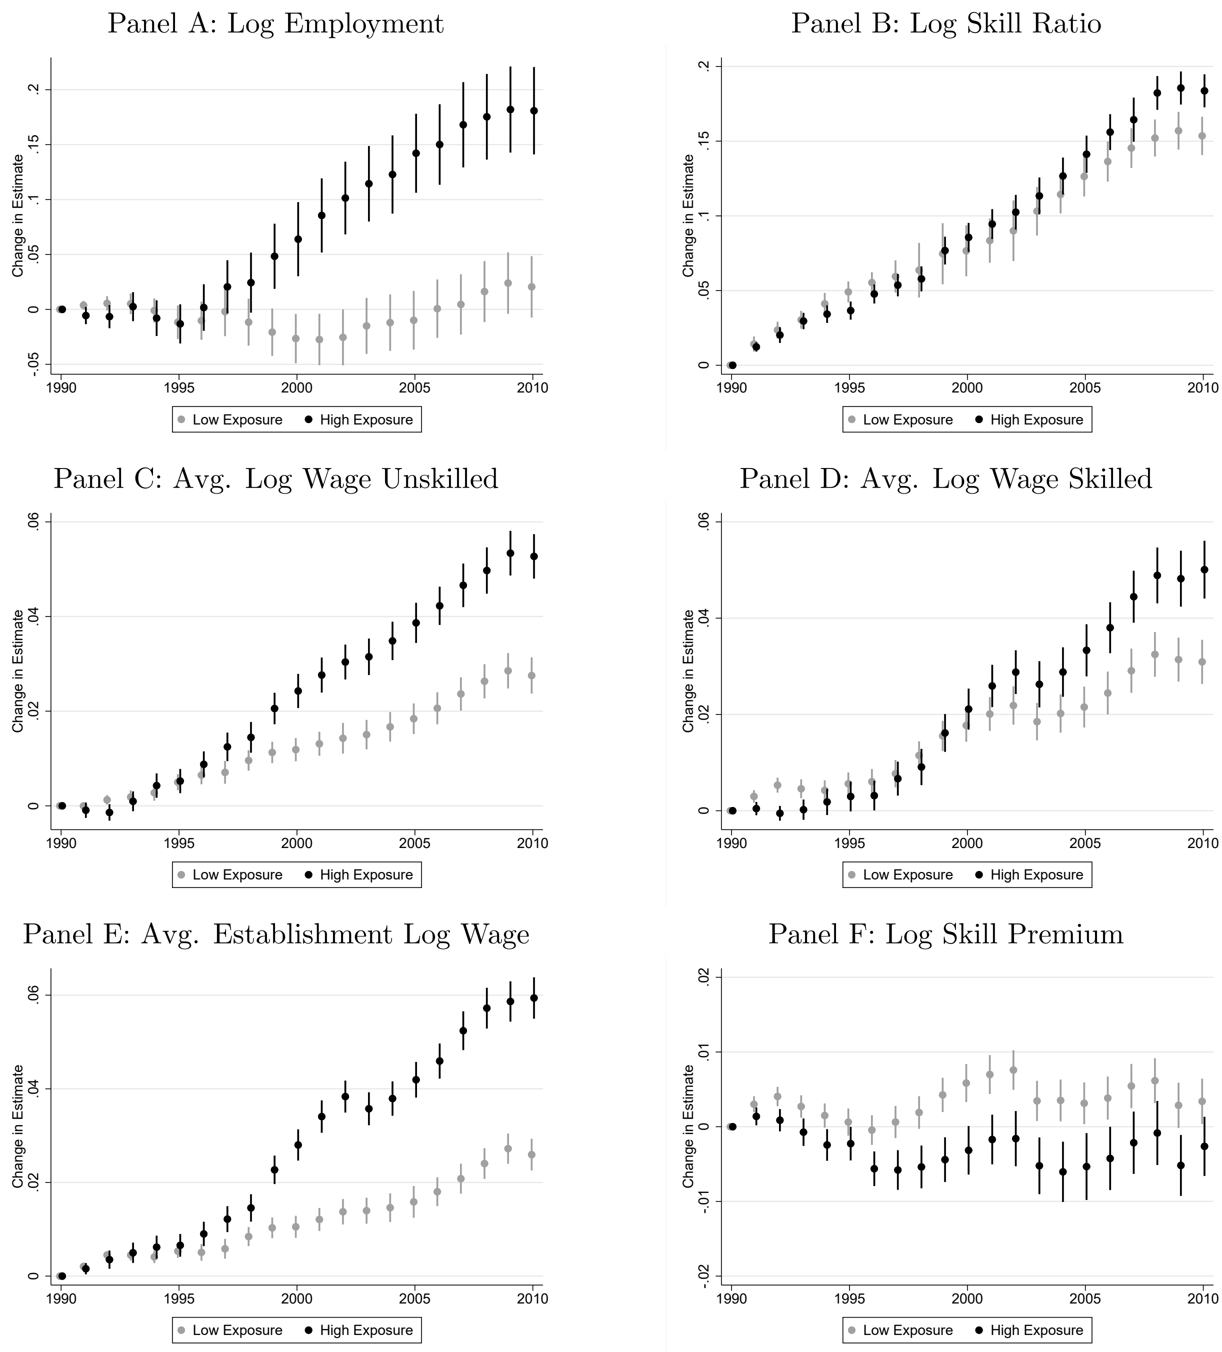

Note: The figure displays the change in the estimated coefficients relative to 1990 (and 95% confidence intervals) based on regressions of the outcome indicated in each panel title on $\hat{\theta}_{f(k)t}$, our model-consistent proxy for establishments' underlying productivity types (see Section 5.1 for details). The changes are plotted separately for low- and high-exposure industries. High-exposure industries are those with above-median changes in their ICT capital stock per worker between 1991 and 2007, as measured in the EUKLEMS data. All regressions control for fully interacted 3-digit industry-by-year fixed effects. Results are based on establishments in the BeH, and observations are weighted by establishment size. Standard errors are clustered at the establishment level.

Table 1: Cross-Sectional Relationships Between Establishment Outcomes and Productivity Proxies

	Log Employment	Log Skill Ratio	Avg. Log Wage	Avg. Log Wage Unskilled	Avg. Log Wage Skilled	Log Skill Premium	Avg. Res. Log Wage Unskilled	Avg. Res. Log Wage Skilled	Res. Log Skill Premium
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Relationship with Establishment Productivity Proxy									
$\hat{\theta}_{f(k)t}$	0.030***	0.87***	0.069***	0.049***	0.076***	0.027***	0.015***	0.030***	0.015***
	(0.0092)	(0.0027)	(0.00098)	(0.00089)	(0.0011)	(0.00066)	(0.00074)	(0.00071)	(0.00042)
N	26,895,503	26,895,503	26,895,503	10,199,928	10,199,928	10,199,928	10,199,928	10,199,928	10,199,928
Panel B: Relationship with 1990 Employment									
Log Employment 1990	–	0.045***	0.070***	0.051***	0.086***	0.034***	0.042***	0.051***	0.0093***
		(0.0061)	(0.0013)	(0.00093)	(0.0014)	(0.00093)	(0.00088)	(0.00087)	(0.00059)
N		6,998,796	6,998,796	4,018,948	4,018,948	4,018,948	4,018,948	4,018,948	4,018,948

Note: The table reports regression results where the dependent variable, indicated at the top of each column, is regressed on either our model-consistent establishment-level productivity proxy $\widehat{\theta}_{f(k)t}$ (Panel A), constructed as described in Section 5.1, or baseline (1990) establishment log employment (Panel B). All regressions include fully interacted 3-digit industry×year fixed effects. In Column (2), establishments with zero 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 allow for the computation of the log skill ratio for all establishments. Columns (4)–(9) are restricted to establishments that employ at least one full-time worker of each skill type. Estimates are based on BeH data—all establishments in Panel A, and only establishments that remain in operation throughout the entire 1990–2010 period in Panel B. Observations are weighted by establishment size—current size in Panel A and baseline (1990) size in Panel B. Standard errors are clustered at the establishment level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Changes in the Associations between $\hat{\theta}_{f(k)t}$ and Establishment Outcomes, 1990–2010, by Technology Exposure

	Log Employment	Log Skill Ratio	Avg. Log Wage	Avg. Log Wage Unskilled	Avg. Log Wage Skilled	Log Skill Premium	Avg. Res. Log Wage Unskilled	Avg. Res. Log Wage Skilled	Res. Log Skill Premium
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A. ICT Capital per Worker: Baseline									
Avg. Change	0.021	0.15***	0.026***	0.028***	0.031***	0.0034**	0.011***	0.024***	0.013***
Low Tech Exposure	(0.014)	(0.0065)	(0.0017)	(0.0020)	(0.0023)	(0.0015)	(0.0015)	(0.0015)	(0.0011)
Additional Change	0.16***	0.030***	0.033***	0.025***	0.019***	-0.0060**	0.013***	0.0054**	-0.0072***
High Tech Exposure	(0.025)	(0.0086)	(0.0028)	(0.0031)	(0.0039)	(0.0025)	(0.0024)	(0.0024)	(0.0016)
Panel B. ICT Capital per Worker: With Additional Controls									
Avg. Change	0.060***	0.091***	0.026***	0.024***	0.033***	0.0091***	0.0079***	0.021***	0.013***
Low Tech Exposure	(0.017)	(0.0069)	(0.0020)	(0.0023)	(0.0029)	(0.0018)	(0.0017)	(0.0019)	(0.0012)
Additional Change	0.17***	0.021***	0.035***	0.027***	0.021***	-0.0058**	0.014***	0.0060**	-0.0081***
High Tech Exposure	(0.024)	(0.0080)	(0.0029)	(0.0030)	(0.0039)	(0.0025)	(0.0024)	(0.0024)	(0.0016)
Panel C. Industry Skill Premium: Baseline									
Avg. Change	0.068***	0.13***	0.028***	0.023***	0.025***	0.0023	0.0090***	0.015***	0.0056***
Low Tech Exposure	(0.018)	(0.0063)	(0.0020)	(0.0021)	(0.0024)	(0.0016)	(0.0017)	(0.0016)	(0.0010)
Additional Change	0.070***	0.084***	0.040***	0.041***	0.036***	-0.0043	0.020***	0.027***	0.0071***
High Tech Exposure	(0.025)	(0.0085)	(0.0030)	(0.0032)	(0.0041)	(0.0027)	(0.0025)	(0.0026)	(0.0016)
Panel D. Industry Skill Premium: With Additional Controls									
Avg. Change	0.11***	0.081***	0.033***	0.025***	0.032***	0.0073***	0.011***	0.016***	0.0057***
Low Tech Exposure	(0.017)	(0.0064)	(0.0021)	(0.0023)	(0.0027)	(0.0017)	(0.0018)	(0.0018)	(0.0011)
Additional Change	0.074***	0.063***	0.035***	0.039***	0.036***	-0.0027	0.016***	0.024***	0.0073***
High Tech Exposure	(0.025)	(0.0079)	(0.0030)	(0.0032)	(0.0043)	(0.0028)	(0.0026)	(0.0026)	(0.0016)
N	2,490,678	2,490,678	2,490,678	952,899	952,899	952,899	952,899	952,899	952,899

Note: The table reports the estimates of β_2 (first row in each panel) and β_3 (second row in each panel) from Equation (10). β_2 represents the estimated change in the slope of the relationship between productivity and the outcome indicated in the title of each column, between 1990 and 2010, within low-exposure industries; β_3 represents the estimated differential change over time for high-exposure industries (over and above that of low-exposure industries). In Panels A and B, industries are classified as low- or high-exposure depending on whether the industry-level change in the ICT capital stock per worker between 1991 and 2007 is below or above the employment-weighted median. In Panels C and D, industries are classified as low- or high-exposure depending on whether their industry skill premium changed by more than the employment-weighted median over 1990–2010. All specifications include 3-digit industry×year fixed effects as well as the additional controls indicated in Equation (10). Panels A and C present baseline results. Panels B and D additionally control for changes in industry import and export exposure between 1990 and 2010, as well as a measure of the risk of offshorability of occupations (aggregated to the industry level using the 1990 occupation structure), each interacted with $\hat{\theta}_{f(k)t}$ and with year fixed effects. The estimation uses data for 1990 and 2010 only. All observations are weighted by establishment size. Observation numbers reported at the bottom of the table are identical across panels. 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 (IAB)
Uta Schönberg (University of Hong Kong, CReAM and IAB)
Jeanne Tschopp (University of Bern)

Appendix A Data

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; establishment size (quadratic, and a dummy for establishment 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 establishment'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.

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 Motivating Evidence

B.1 Relative Employment of Skilled Workers (Figure 1, Panel A)

The log relative employment of skilled workers at time t, $LogSkillRatio_t$ (the black line in Figure 1, Panel A) is given by:

$$LogSkillRatio_t = \ln\left(\frac{n_t^s}{n_t^u}\right)$$

where n_t^s and n_t^u represent the total number of skilled and unskilled workers at time t, respectively.

Note that this can also be expressed as:

$$LogSkillRatio_t = \ln\left[\frac{\sum_k n_{kt}^s}{\sum_k n_{kt}^u}\right] = \ln\left[\frac{\sum_k (n_{kt}^s/n_{kt})(n_{kt}/n_t)}{\sum_k (n_{kt}^u/n_{kt})(n_{kt}/n_t)}\right]$$

where n_{kt}^s (n_{kt}^u) denotes the total number of skilled (unskilled) workers in industry k at time t, n_{kt} denotes total employment in industry k at time t, and n_t is total employment at time t. This expression makes clear that we can think of the skill ratio at time t as the sum across industries of the share of skilled workers in each industry, weighting each industry with its share of total employment, divided by the sum across industries of the share of unskilled workers in each industry, again weighting each industry with its share of total employment.

We can then express the counterfactual log skill ratio in year t holding the industry structure constant at its 1990 employment level (the gray line in Figure 1, Panel A) as:

$$LogSkillRatio_t^{1990} = \ln\left[\frac{\sum_k (n_{kt}^s/n_{kt})(n_{k1990}/n_{1990})}{\sum_k (n_{kt}^u/n_{kt})(n_{k1990}/n_{1990})}\right]$$

B.2 Skill Wage Premium (Figure 1, Panel B)

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

$$LogSkillPrem_t = \overline{\ln w_t}^s - \overline{\ln w_t}^u,$$

where $\overline{\ln w_t}^s$ ($\overline{\ln w_t}^u$) is the average log wage of skilled (unskilled) workers at time t.

Note that this can also be expressed as:

$$LogSkillPrem_t = \sum_k \frac{n_{kt}^s}{n_t^s}\overline{\ln w_{kt}}^s - \sum_k \frac{n_{kt}^u}{n_t^u}\overline{\ln w_{kt}}^u,$$

where $\overline{\ln w}_{kt}^{s}$ ($\overline{\ln w}_{kt}^{u}$) denotes the average log wage of skilled (unskilled) workers in industry k at time t. In other words, the log skilled wage premium can be written as the difference between a weighted average of the average log wage of skilled and unskilled workers by industry and year, with the industrial employment shares by worker type as weights.

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

$$LogSkillPrem_t^{1990} = \sum_k \frac{n_{k1990}^s}{n_{1990}^s} \overline{\ln w}_{kt}^{s} - \sum_k \frac{n_{k1990}^u}{n_{1990}^u} \overline{\ln w}_{kt}^{u}.$$

B.3 Between-Establishment Inequality (Figure 1, Panels C and D)

We decompose the variance of individual log wages, denoted Var_t, into a within-establishment and a between-establishment component as follows:

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

where i denotes an individual, f indexes an establishment and i_{ft} denotes the set of workers in establishment f at time t. $\ln w_{it}$ 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_{ft} is the total number of workers at establishment f in year t. The evolution over time of Var_t and of each component in Equation (B.1) is displayed in Panel C of Figure 1.

We further decompose between-establishment wage differentials, Var_t^{BE} in Equation (B.1), as follows:

$$Var_t^{BE} = \underbrace{\frac{1}{n_t}\sum_k \sum_{f \in f_{kt}} n_{ft}(\overline{\ln w}_{ft} - \overline{\ln w}_{kt})^2}_{\text{between estab., within industries } (Var_t^{WIBE})} + \underbrace{\frac{1}{n_t}\sum_k n_{kt}(\overline{\ln w}_{kt} - \overline{\ln w_t})^2}_{\text{between industries}}, \tag{B.2}$$

where f_{kt} is the set of establishments in 3-digit industry k in year t and $\overline{\ln w}_{kt}$ is the average log wage in industry k at time t. The dark gray line with diamonds in Panel D of Figure 1 displays the evolution over time of Var_t^{WIBE} in Equation (B.2).

We also compute counterfactual within-industry between-establishment wage inequality using the 1990 industry structure (the light gray line with triangles in Panel D of Figure 1) as follows:

$$Var_t^{WIBE,1990} = \sum_k \frac{n_{k1990}}{n_{1990}} \underbrace{\sum_{f \in f_{kt}} \frac{n_{ft}}{n_{kt}} (\overline{\ln w}_{ft} - \overline{\ln w}_{kt})^2}_{Var_{kt}(\overline{\ln w}_{ft})}, \tag{B.3}$$

where $Var_{kt}(\overline{\ln w}_{ft})$ is the variance of establishment log wages in industry k at time t.

Appendix C Theoretical Framework

This section provides additional details related to our theoretical framework, which builds on the Helpman et al. (2010) model (hereafter HIR) with a CES production function and two types of workers (see Section 5.4 of the HIR Online Appendix).

C.1 Setup

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 $0 < \beta < 1$.

Given the aggregate consumption index, the demand function for variety j is given by $q(j) = A^{\frac{1}{(1-\beta)}} p(j)^{-\frac{1}{(1-\beta)}}$, where $p(j)$ denotes the price of variety j and $A = E^{1-\beta} P^{\beta}$ is a sectoral demand shifter. $E = PQ$ denotes the total expenditures on varieties and P is the aggregate price index within a sector. Using this demand function we can express revenues from variety j as follows:

$$r(j) = Aq(j)^{\beta}. \tag{C.1}$$

As in Melitz (2003), there is a continuum of potential entrants, each of which produces a single variety upon entering the market. Entry requires paying a fixed cost $f_e > 0$, after which each firm f draws an idiosyncratic productivity level θ_f from a common Pareto distribution $G_{\theta}(\theta)$ with scale parameter $\theta_{\min}$ and shape parameter z, i.e. $G_{\theta}(\theta) = 1 - (\theta_{\min}/\theta)^z$ for $\theta \geq \theta_{\min} > 0$ and $z > 2$.[1] Producing a variety requires a fixed cost $f_d > 0$. After observing their productivity draw, firms decide whether to exit or produce.

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

The production function of a firm with productivity θ_f is given by:

$$y_f = \left[\left(\theta_f^{1+\alpha\mu_s}\mu_s\overline{a}_{s,f}h_{s,f}^{\gamma}\right)^{\nu} + \left(\overline{a}_{u,f}h_{u,f}^{\gamma}\right)^{\nu}\right]^{1/\nu}, \tag{C.2}$$

where $h_{s,f}$ and $h_{u,f}$ are the firm's input of skilled and unskilled labor, respectively, $\overline{a}_{s,f}$ and $\overline{a}_{u,f}$ represent the average match-specific ability of skilled and unskilled workers hired by the firm, μ_s is the aggregate SBTC term, and $0 < \nu < \beta$. The parameter $\alpha \geq 0$ governs the strength of the interaction between SBTC and firm productivity. Setting $\alpha = 0$ yields a production function identical to that in HIR (Section 5.4 of their Online Appendix).

The labor market features standard search and matching frictions. The firm must pay a search cost of $b_\ell n_{\ell,f}$ in order to be matched with $n_{\ell,f}$ workers, $\ell = \{s, u\}$.[2] Skilled workers are relatively scarce and hence command a higher search cost, i.e. $b_s > b_u$. Workers of a given skill type are ex-ante identical. However, upon matching with a firm, worker i draws match-specific ability a_i 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 $k > 1$.[3] Ability a_i is not observable by the firm or the worker, but a screening technology is available. By paying a screening cost $c\widetilde{a}_{\ell,f}^{\delta}/\delta$, the firm can determine whether a worker's ability exceeds a chosen threshold $\widetilde{a}_{\ell,f}$, where $\ell = \{s, u\}$, $c > 0$, and $\delta > k$.[4] In equilibrium, firms hire only the matched workers whose ability exceeds their threshold $\widetilde{a}_{\ell,f}$. Since the screening process reveals only whether ability exceeds the threshold—not its exact value—all hired workers of type ℓ in a given firm are treated as having the same expected ability, equal to $\overline{a}_{\ell,f} = k\widetilde{a}_{\ell,f}/(k-1)$.

Given the screening threshold $\widetilde{a}_{\ell,f}$, the firm hires $h_{\ell,f} = n_{\ell,f}\left(a_{\min}/\widetilde{a}_{\ell,f}\right)^k$ workers of type ℓ. Substituting this expression into Equation (C.2) and combining with (C.1), firm revenue can be expressed as:

$$r_f = A\kappa_y^{\beta}\left[(\theta_f^{1+\alpha\mu_s}\mu_s\widetilde{a}_{s,f}^{1-\gamma k}n_{s,f}^{\gamma})^{\nu} + (\widetilde{a}_{u,f}^{1-\gamma k}n_{u,f}^{\gamma})^{\nu}\right]^{\frac{\beta}{\nu}} \tag{C.3}$$

where $\kappa_y = \frac{k a_{\min}^{\gamma k}}{k-1}$.

C.2 Firm's Problem

After observing its productivity, the firm decides whether to produce, how many workers of each type to match with, and the screening thresholds. It then bargains with its hired

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

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

[4] The assumption $\delta > k$ ensures that employment and wages of both worker types increase in θ_f, consistent with empirical evidence.

workers over revenue, following the Stole & Zwiebel (1996a,b) framework.

The resulting wage schedules are:[5]

$$w_{\ell,f} = \frac{\beta\gamma}{1+\beta\gamma}\frac{\phi_{\ell,f} r_f}{h_{\ell,f}}, \tag{C.4}$$

where $\phi_{s,f} \equiv \frac{\varphi_f}{1+\varphi_f}$, $\phi_{u,f} \equiv \frac{1}{1+\varphi_f}$ and $\varphi_f \equiv \frac{\left(\theta_f^{1+\alpha\mu_s}\mu_s \bar{a}_{s,f} h_{s,f}^{\gamma}\right)^{\nu}}{\left(\bar{a}_{u,f} h_{u,f}^{\gamma}\right)^{\nu}}$.

Anticipating the outcome of wage bargaining and using Equation (C.3), the firm chooses $n_{\ell,f}$ and $\widetilde{a}_{\ell,f}$ to maximize profits, i.e.:

$$\begin{aligned} \pi_f \;\; = & \max_{n_{\ell,f},\widetilde{a}_{\ell,f}} \left\{ \frac{A\kappa_y^{\beta}}{1+\beta\gamma}\left[(\theta_f^{1+\alpha\mu_s}\mu_s \widetilde{a}_{s,f}^{1-\gamma k} n_{s,f}^{\gamma})^{\nu} + (\widetilde{a}_{u,f}^{1-\gamma k} n_{u,f}^{\gamma})^{\nu} \right]^{\frac{\beta}{\nu}} \right. \\ & - \left. b_s n_{s,f} - b_u n_{u,f} - \frac{c}{\delta}\widetilde{a}_{s,f}^{\delta} - \frac{c}{\delta}\widetilde{a}_{u,f}^{\delta} - f_d \right\}, \end{aligned}$$

where b_ℓ denotes the search cost to be matched with a worker of type ℓ and $\frac{c}{\delta}\widetilde{a}_{\ell,f}^{\delta}$ captures the screening costs the firm must pay to identify type-ℓ workers above the threshold $\widetilde{a}_{\ell,f}$.

The first-order conditions that result from this profit maximization problem are:

$$\frac{\beta\gamma}{1+\beta\gamma}\phi_{\ell,f} r_f = b_\ell n_{\ell,f} \tag{C.5}$$

$$\frac{\beta(1-\gamma k)}{1+\beta\gamma}\phi_{\ell,f} r_f = c\widetilde{a}_{\ell,f}^{\delta}. \tag{C.6}$$

Taking the ratio of each condition across skill types, we obtain:

$$\frac{n_{s,f}}{n_{u,f}} = \varphi_f \frac{b_u}{b_s} \tag{C.7}$$

$$\frac{\widetilde{a}_{s,f}}{\widetilde{a}_{u,f}} = \varphi_f^{1/\delta}. \tag{C.8}$$

C.3 Key Equilibrium Relationships

In this section, we focus on the equilibrium relationships that are central for understanding the effects of SBTC on between-firm wage inequality. We begin by deriving an expression for the equilibrium value of φ_f as a function of firm productivity θ_f, model parameters, and the search costs b_ℓ. We then express all firm-level equilibrium outcomes as functions of φ_f, the search costs b_ℓ, the sectoral shifter A, and parameters. Finally, we discuss how the sector-level outcomes are determined.

Note that while in this section we discuss all of the equilibrium relationships in *levels*

[5] See Section 7.3 of the Online Appendix of HIR.

(as these are easy to infer from the model), in the body of the paper we focus on the log-log relationships between the various firm-level outcomes and the firm-level productivity parameter, in order to ensure a clear connection between our theoretical predictions and our empirical analysis.

Equilibrium $\boldsymbol{\varphi_f}$ First, using the expressions for $h_{\ell,f}$ and $\overline{a}_{\ell,f}$, we can write φ_f as:

$$\varphi_f = \left(\theta_f^{1+\alpha\mu_s}\mu_s\right)^{\nu}\left(\frac{n_{s,f}}{n_{u,f}}\right)^{\gamma\nu}\left(\frac{\widetilde{a}_{s,f}}{\widetilde{a}_{u,f}}\right)^{(1-k\gamma)\nu} \tag{C.9}$$

Substituting the firm-level ratios from Equations (C.7) and (C.8) into (C.9) yields:

$$\varphi_f = \left(\theta_f^{1+\alpha\mu_s}\mu_s\right)^{\nu/\Lambda}\left(\frac{b_s}{b_u}\right)^{-\gamma\nu/\Lambda}, \tag{C.10}$$

where $\Lambda \equiv 1-\nu\gamma-\nu(1-\gamma k)/\delta > 0$ and $\Gamma \equiv 1-\beta\gamma-\beta(1-\gamma k)/\delta > 0$. Given the assumption that $\nu < \beta$, we have that $\Lambda > \Gamma$. Note that in the body of the paper we express the equilibrium value of φ_f as $\varphi(\theta_f;\mu_s)$, to highlight its dependence on the aggregate and the firm-specific skill-biased technology parameters.

Equation (C.10) shows that the equilibrium φ_f is strictly increasing in θ_f. We now derive the remaining firm-level equilibrium outcomes as functions of φ_f, as given by Equation (C.10).

Firm-Level Revenue Using Equation (C.9) together with (C.3), firm revenue can be expressed as:

$$r_f = \kappa_y^{\beta}A\left(1+\varphi_f\right)^{\beta/\nu}\left(\widetilde{a}_{u,f}^{1-k\gamma}n_{u,f}^{\gamma}\right)^{\beta}. \tag{C.11}$$

Substituting the first-order conditions (C.5) and (C.6) into (C.11), we obtain the equilibrium expression for firm-level revenue:

$$r_f = \kappa_r\left(1+\varphi_f\right)^{\frac{\beta\Lambda}{\nu\Gamma}}, \tag{C.12}$$

where κ_r is defined as:

$$\kappa_r \equiv A^{1/\Gamma}\left[\kappa_y\left(\frac{\beta}{1+\beta\gamma}\right)^{\frac{1-k\gamma}{\delta}+\gamma}\left(\frac{1-\gamma k}{c}\right)^{\frac{1-k\gamma}{\delta}}\left(\frac{\gamma}{b_u}\right)^{\gamma}\right]^{\beta/\Gamma}. \tag{C.13}$$

From (C.12), it is straightforward to see that revenues will be increasing in θ_f.

Firm-Level Employment We first derive an expression for the employment of unskilled workers. Combining Equations (C.5) and (C.12), we obtain:

$$n_{u,f} = \left(\frac{\beta\gamma}{1+\beta\gamma}\right) b_u^{-1} \kappa_r \left(1+\varphi_f\right)^{\frac{\beta-\nu}{\nu\Gamma}}, \tag{C.14}$$

where $\frac{\beta-\nu}{\nu\Gamma} = \frac{\beta\Lambda}{\nu\Gamma} - 1 > 0$. Next, using Equations (C.6) and (C.12), we express the ability threshold for unskilled workers as:

$$\widetilde{a}_{u,f} = \left[\frac{\beta(1-\gamma k)}{1+\beta\gamma}\right]^{1/\delta} c^{-1/\delta} \kappa_r^{1/\delta} \left(1+\varphi_f\right)^{\frac{\beta-\nu}{\delta\nu\Gamma}}. \tag{C.15}$$

Equilibrium employment of unskilled workers is then obtained by substituting (C.14) and (C.15) into $h_{\ell,f} = n_{\ell,f}\left(\frac{a_{min}}{\widetilde{a}_{\ell,f}}\right)^k$, which yields:

$$h_{u,f} = h_{dr} \left(1+\varphi_f\right)^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\left(1-\frac{k}{\delta}\right)}, \tag{C.16}$$

where:

$$h_{dr} \equiv \left(\frac{\beta\kappa_r}{1+\beta\gamma}\right)^{1-k/\delta} \left(\frac{c}{1-\gamma k}\right)^{k/\delta} b_u^{-1} a_{min}^k. \tag{C.17}$$

Given that φ_f is increasing in θ_f, employment of unskilled workers will be increasing in θ_f as long as $\beta > \nu$ and $\delta > k$.

Proceeding similarly for skilled workers, we can express their equilibrium employment as:

$$h_{s,f} = \frac{b_u}{b_s} \varphi_f^{1-k/\delta} h_{u,f}. \tag{C.18}$$

This is also increasing in θ_f under the parameter restrictions discussed above, given that $h_{u,f}$ is increasing in θ_f and $(1 - k/\delta) > 0$.

It follows that total firm-level employment will be increasing in θ_f, and the firm-level skill ratio will be given by:

$$\frac{h_{s,f}}{h_{u,f}} = \frac{b_u}{b_s} \varphi_f^{1-k/\delta}, \tag{C.19}$$

which is increasing in θ_f as well.

Firm-Level Wages The equilibrium wage for unskilled workers is obtained by substituting Equations (C.11) and (C.16) into Equation (C.4), yielding:

$$w_{u,f} = w_{dr} \left(1+\varphi_f\right)^{\left(\frac{\beta-\nu}{\nu\Gamma}\right)\frac{k}{\delta}}, \tag{C.20}$$

where:

$$w_{dr} \equiv \left(\frac{\beta\gamma}{1+\beta\gamma}\right)\left(\frac{\kappa_r}{h_{dr}}\right). \tag{C.21}$$

Given $\beta > \nu$, the wages of unskilled workers will be increasing in θ_f.

Proceeding similarly for skilled workers, we obtain the equilibrium wage:

$$w_{s,f} = \frac{b_s}{b_u}\varphi_f^{k/\delta} w_{u,f}, \tag{C.22}$$

which will also be increasing in θ_f given the positive relationship between this parameter and both $w_{r,f}$ and φ_f.

Finally, it follows that the firm-level skill premium is:

$$\frac{w_{s,f}}{w_{u,f}} = \frac{b_s}{b_u}\varphi_f^{k/\delta}, \tag{C.23}$$

which is also increasing in θ_f.

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

The ZCP condition requires that the firm with productivity θ_d makes zero profits. This implies:[6]

$$\frac{\Gamma}{1+\beta\gamma} r_d = f_d, \tag{C.24}$$

where r_d denotes revenue for the marginal (zero-profit) firm with productivity θ_d.

Using Equation (C.12), relative revenues for any firm with productivity θ_f and the marginal firm with productivity θ_d are:

$$\frac{r_f}{r_d} = \left(\frac{1+\varphi_f}{1+\varphi_d}\right)^{\frac{\beta\Lambda}{\nu\Gamma}}, \tag{C.25}$$

[6]This follows from the profit expression:

$$\pi_f = \frac{\Gamma}{1+\beta\gamma} r_f - f_d.$$

where φ_d denotes $\varphi(\theta_d)$. Substituting Equation (C.25) into the ZCP condition (C.24) gives:

$$r_f = f_d \left(\frac{\Gamma}{1+\beta\gamma} \right)^{-1} \left(\frac{1+\varphi_f}{1+\varphi_d} \right)^{\frac{\beta\Lambda}{\nu\Gamma}}. \tag{C.26}$$

The FE condition requires that the expected profits of a potential entrant equal the fixed entry cost f_e. Since firms draw productivity θ_f from a common distribution $G(\theta)$, we have:

$$\int_{\theta_d}^{\infty} \pi_f dG(\theta) = f_e, \tag{C.27}$$

where π_f denotes the profit of a firm with productivity θ_f.

Combining Equation (C.26) and (C.27), we obtain:

$$f_d \int_{\theta_d}^{\infty} \left[\left(\frac{1+\varphi_f}{1+\varphi_d} \right)^{\frac{\beta\Lambda}{\nu\Gamma}} - 1 \right] dG(\theta) = f_e. \tag{C.28}$$

Equation (C.28) pins down the equilibrium threshold θ_d as a function of the model's parameters and the search costs b_s and b_u.

Other Sectoral Outcomes Following HIR, we assume that search costs are proportional to workers' expected income outside the sector. When the outside good is homogeneous and produced without search frictions—a standard assumption in international trade—expected worker income is pinned down by the wage in the outside sector and, as long as the economy remains incompletely specialized, does not respond to sectoral shocks. This implies that search costs can be treated as constant terms that are unaffected by shocks to a given sector (see Sections 2.3.1 and 6.2 of HIR for a full discussion). We adopt this approach here as well, and treat b_ℓ as constants.

The sectoral shifter A (which is part of κ_r in equation (C.13), h_{dr} in equation (C.17), and w_{dr} in equation (C.21)) can be pinned down from equation (C.24) (the ZCP condition), which defines r_d, together with equation (C.28), which determines the productivity threshold θ_d and therefore φ_d. With these two objects in hand, we can recover κ_r from equation (C.12) and, in turn, obtain the sectoral shifter from equation (C.13).

We refer the reader to HIR Section 2.3 for further details on how other sectoral variables are determined in equilibrium, and HIR Section 6 for a discussion of how the model is closed in general equilibrium.

Appendix D Simulation

Following Helpman et al. (2017), we set $\beta = 0.75$, $\gamma = 0.5$. The substitution elasticity between worker types is set to $\nu = 0.5\beta$. The shape parameter of the worker ability distribution is set to $k = 4/3$, also following Helpman et al. (2017), and we choose the scale parameter to be $a_{\min} = 0.1$. For the productivity distribution, we set the shape parameter to $z = 3$ to ensure finite variance, and the scale parameter to $\theta_{\min} = 0.1$. The screening function parameter δ is set to $\delta = 11.1k$, consistent with Helpman et al. (2017), and the parameter c is chosen to be small, as suggested in HIR, and set to $c = 0.01$. Fixed entry and operating costs are set to $f_e = 0.5$ and $f_d = 15$, values commonly used in the trade literature (e.g. Ghironi & Melitz, 2005). As discussed above, and consistent with HIR, search costs are assumed to be fixed. We set them to $b_s = 4$ and $b_u = 1$, reflecting the higher cost of recruiting skilled workers. Finally, we choose a small value of $\alpha = 0.05$ to avoid overstating the strength of the relationship between SBTC and firm productivity. We also show results using $\alpha = 0$.

We simulate the model for different values of μ_s. For each value of μ_s, we use the following sequence of steps:

1. Use equation (C.28) (which combines the ZCP and FE conditions) together with (C.10) to solve for the productivity threshold θ_d and, hence, φ_d;
2. Solve for revenue at the threshold, r_d, using equation (C.24);
3. Use equation (C.12) to determine κ_r (and hence, the sectoral demand shifter, A) from φ_d and r_d;
4. Determine h_{dr} using equation (C.17);
5. Determine w_{dr} using equation (C.21);
6. Use these quantities to recover the remaining relevant firm-level variables (i.e. equations (C.16)-(C.19) for employment variables, (C.20)-(C.23) for wage variables).

Figure 2 and Figure A.2 illustrate the relationship between simulated firm-level outcomes and firm log productivity for both a lower and a higher value of μ_s; the lower value reflects the baseline relationship prior to the SBTC shock, while the higher value captures the relationship following the shock. We set $\mu_s = e$ (so that $\ln(\mu_s) = 1$) for the baseline and $\mu_s = e^{1.15}$ for the post-shock scenario. These values are chosen arbitrarily but do not affect the qualitative nature of the relationships presented, as shown in Figure 3 and Figure A.1, which illustrate patterns for a wide range of potential values of μ_s.

Figure A.1: Aggregate Predictions of SBTC from the Model Simulation

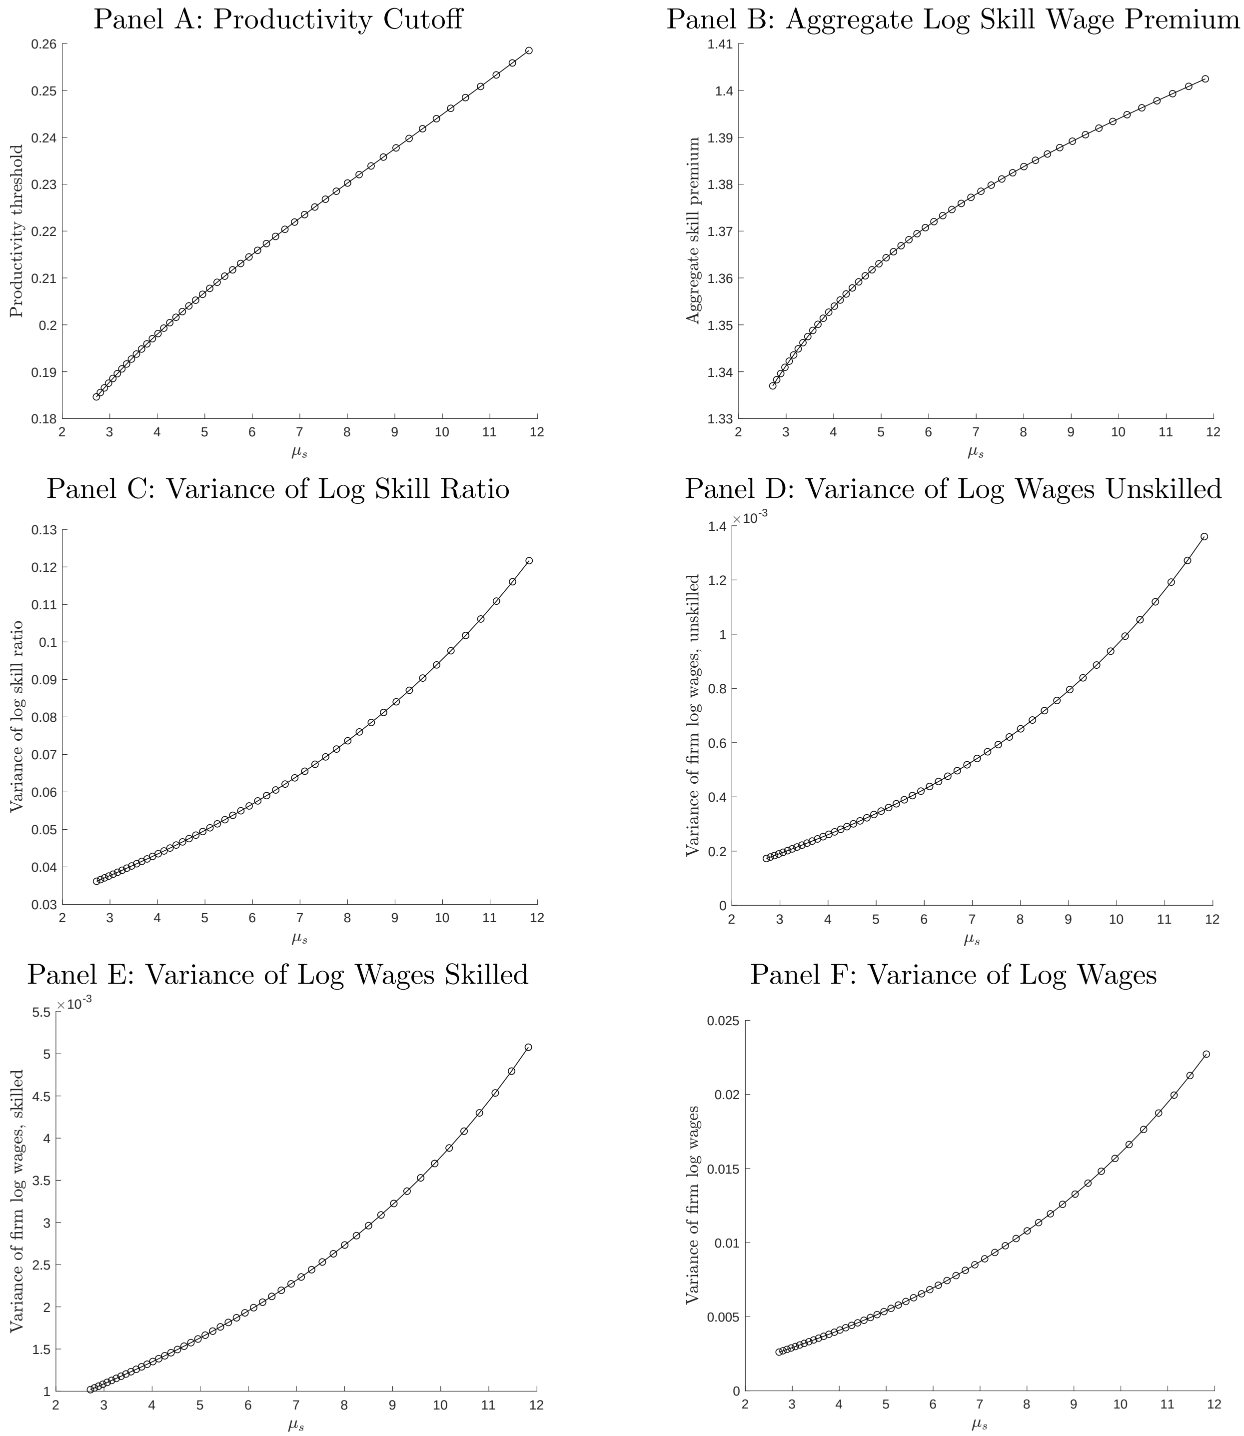

Note: The figure shows simulated aggregate outcomes for various levels of SBTC (μ_s). Panel A shows the productivity threshold below which firms optimally choose not to operate. Panel B reports the aggregate log skill premium, defined as the employment-weighted difference between average log skilled and unskilled wages. Panel C shows the employment-weighted variance of the firm log skill ratio. Panels D–F report the employment-weighted variances of firm log wages, by skill type and overall. The measures in Panels B-F take into account both the within-firm changes and the change in the composition of operating firms as μ_s changes. Simulation details are provided in Appendix D.

Figure A.2: Firm-Level Outcomes and Baseline Size, at Baseline and after an SBTC Shock

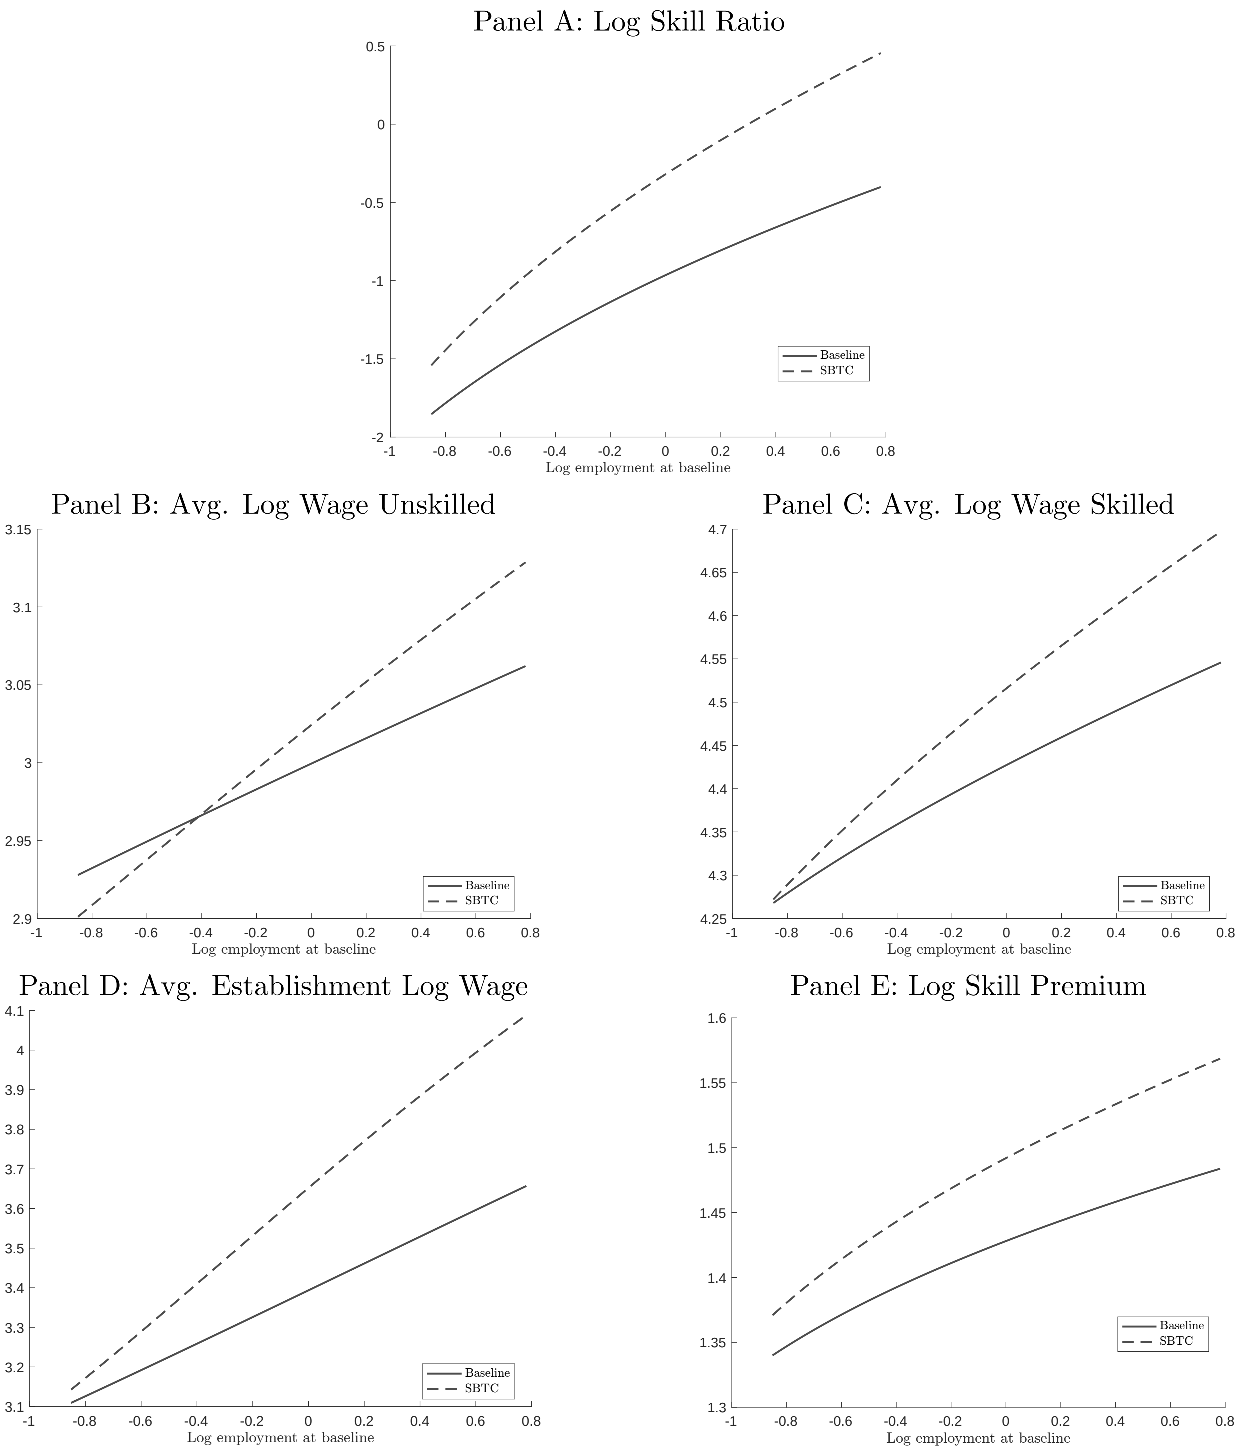

Note: The figure shows the simulated relationship between firm-level outcomes (as indicated in each panel title) and firm log baseline size. The solid line corresponds to a low value of μ_s, representing the baseline relationship prior to an SBTC shock. The dashed line corresponds to a higher value of μ_s, capturing the effect of the shock. Simulation details are provided in Appendix D.

Figure A.3: Year-by-Year Associations between Log Establishment Employment in 1990 and Other Establishment Characteristics

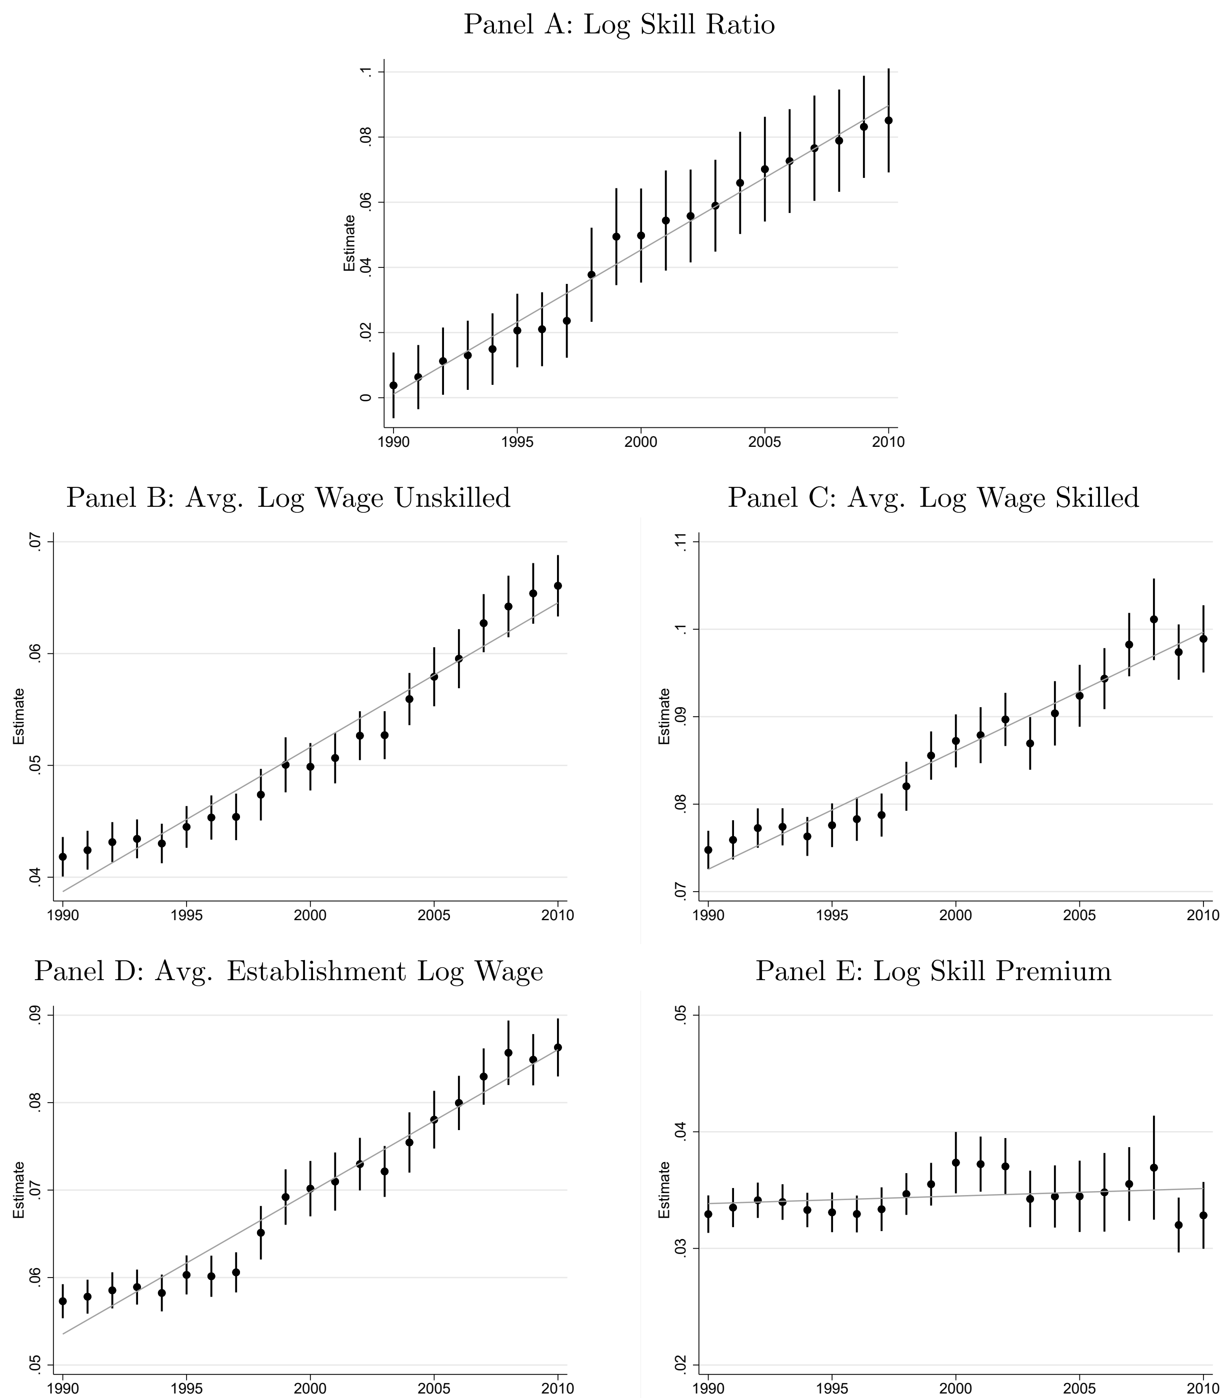

Note: The figure shows estimated coefficients and 95% confidence intervals from regressions of the outcome indicated in each panel title on log establishment employment in 1990—the alternative proxy for establishments' underlying productivity $\log \theta_f$. Coefficients are allowed to vary by year, with all regressions controlling for fully interacted 3-digit industry$\times$year fixed effects. Estimates are based on establishments in the BeH data that remain in operation throughout the entire 1990–2010 period. Observations are weighted by baseline (1990) establishment size. Standard errors are clustered at the establishment level.

Table A.1: Classification of Vocational Occupations

Vocational category	Occupation group	Occupation codes (KldB88)
Skilled Vocational	Technician, Engineer, Skilled Service, Skilled Administrative, Semiprofessions, Professions, Managers	11–22, 41–51, 53–549 excl. 303 & 304; 682, 685–688, 706, 713–716, 723–725, 731, 732, 734–744, 773, 782–794, 805, 838, 911–913, 923–937
Unskilled Vocational	Agricultural, Unskilled Service, Unskilled Administrative, Manual Occupations	31, 32, 52, 303, 304, 601–684 excl. 682; 683, 684, 691–712 excl. 706; 721, 722, 726, 733, 751–781 excl. 773; 801–902 excl. 805 & 838; 921, 922

Note: To classify individuals with an apprenticeship or vocational training as either skilled or unskilled, we rely on the vocational occupation of their training (or if the individual is not observed during their training, we rely on their occupation when they are first observed in the data). The classification of occupations into broad occupation groups is from Blossfeld (1987); we then divide these occupation groups into skilled and unskilled vocational occupations.