Wage Adjustment and Productivity Shocks*

Mikael Carlsson[†], Julián Messina[‡] and Oskar Nordström Skans[§]

December 20, 2012

Abstract

We study how workers' wages respond to changes in physical productivity using unique Swedish data that combines firm-level price and quantity information with individual wage data. Wages respond three times as much to productivity shocks that are shared with outside firms within the same narrow sector compared to purely idiosyncratic (firm-level) productivity shocks. Further, the larger impact of sectoral productivity is explained by a high degree of within-sector labor mobility, suggesting that workers' outside options are fundamental for wage determination. In contrast, we show that the skill mix of firms, and their internal relative pay scale, remain largely unaffected by changes in physical productivity.

*We thank three anonymous referees, John Abowd, Per-Anders Edin, Peter Fredriksson, John Haltiwanger, Marcel Jansen, Philipp Kircher, Francis Kramarz, Rafael Lalive, Jean-Marc Robin, Nuria Rodríguez-Planas and seminar participants at SciencePo, the Helsinki Center of Economic Research (HECER), the Euro-System WDN network, the 6th ECB/CEPR Labor Market Workshop, IFAU, University of Oslo, Stockholm University, the Swedish National Conference in Economics 2010, the "Labor Market and the Macroeconomy" conference at Sveriges Riksbank, the LMDG workshop 2010, the UCLS Industrial Relations workshop 2011, and the 2011 ESSLE and 2010 EEA, EALE/SOLE, and LACEA meetings for their useful comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank or the World Bank.

[†]Uppsala University, UCLS and Sveriges Riksbank. E-mail: mikael.carlsson@nek.uu.se

[‡]Office of the Chief Economist for Latin America and the Caribbean, World Bank. E-mail: jmessina@worldbank.org

[§]Institute for Evaluation of Labour Market and Education Policy (IFAU), Uppsala University, UCLS and IZA. E-mail: oskar.nordstrom_skans@ifau.uu.se

1 Introduction

A large literature shows that real wages and labor productivity are intimately related at the aggregate level.[1] Although productivity growth originates at the firm level, much less is known about the relationship between wages and productivity at the micro level. A unifying result spanning across a large family of wage-bargaining models is that incumbent workers should reap rents both from changes in firm-specific productivity and from changes in the productivity of other firms that draw their labor from the same market segment. The hold-up problem associated with workers' extraction of rents has long been a concern for economists because it is likely to reduce the incentives for firms to innovate.[2] Moreover, the importance for wage determination of firm-level productivity versus productivity developments in competing firms remains a key issue within current labor economics. Understanding their relative weight is informative about the importance of worker's outside options during the bargaining process.

In this empirical paper we make two interrelated contributions to the literature. The first is to provide evidence of the distinct impact on incumbent wages of shocks that are purely idiosyncratic to the employing firm and shocks that are shared with the firms that represent the workers' relevant outside options. The second contribution is to identify these wage effects of productivity shocks from shifts in the production function using an empirical approach that handles confounding elements arising from endogenous adjustments along the production function as well as relative price adjustments.

Most bargaining models assume a role for outside options in shaping the wages of

[1] A recent example is Carneiro, Guimarães, and Portugal (2012), who finds an elasticity equal to unity of aggregate productivity on individual wages. See also Carneiro, Guimarães, and Portugal (2012) for references to this literature.

[2] See e.g. Card, Devicienti, and Maida (2011) and references therein.

incumbent workers, either through continuous renegotiations or through renegotiation in response to counteroffers from other firms.[3] But the magnitude of the impact of changes in outside options is poorly documented in the empirical literature. In this paper, therefore, we provide an analysis that separates idiosyncratic shocks from shocks that are shared with other firms in the same sector. We motivate the analysis using a stylized model building on the excess probability that workers remain within the same sector if separated from the employing firm. The model provides us with predictions regarding the relative impact of sectoral and idiosyncratic shocks that we can use as a benchmark for our empirical results. In the end, we show that the wage impact is indeed consistent with the observed worker mobility patterns. To the best of our knowledge, this is the first paper to document that shocks to sector-specific outside options are empirically relevant for incumbent workers' wages with a magnitude that is consistent with workers' observed cross-sectoral mobility patterns.[4]

The second contribution of the paper is motivated by what we perceive as a conceptual gap between much of the empirical literature and the related body of theoretical work. Whereas changes in physical total factor productivity, i.e. shifts in firm-level production functions, are the closest empirical counterpart of the exogenous productivity processes embedded in most relevant theoretical models, empirical studies instead relate wages to revenue-based measures of rents or productivity.[5] The concepts differ since the frequently used empirical measures (e.g. value added per

[3] In contrast, theories of internal labor markets (Doeringer and Piore, 1971) tend to assume that workers' wages are insulated from outside disturbances.

[4] Importantly, analyzing the direct wage impact of sectoral shocks without accounting for firms' idiosyncratic shocks does not allow the researcher to distinguish between the impact of sectoral outside options and the direct impact of idiosyncratic productivity.

[5] Some of these theoretical models are Smith (1999), Postel-Vinay and Turon (2010), and Bagger, Fontaine, Postel-Vinay, and Robin (2011) and Eeckhout and Kircher (2011).

worker or profits) in a real-world settings will include endogenous adjustments of the scale of production and relative prices in response to many types of shocks. A likely reason for this apparent shortcoming is that few data sets (see e.g. the survey by Syverson, 2011) allow the researcher to remove the impact of relative prices from measured productivity, which is needed in order to derive measures of physical total factor productivity (TFPQ) in contrast to revenue-based total factor productivity (TFPR). The distinction between TFPQ and TFPR was highlighted by Foster, Haltiwanger, and Syverson (2008), who show that firms' exit probabilities tend to respond differently to the different shocks that are embedded in revenue and physical productivity constructs. In our particular setting, it is important to note that firm-level prices tend to be a function of factor prices, including wages (see e.g. Carlsson and Nordstrom Skans, 2012). Firms with high costs (e.g. due to high wages) - and thus high prices - and firms with high physical productivity can therefore only be separated if we are able to remove relative-price differences between the firms. It follows that if a distinction is not made between firm and sectoral prices we may severely underestimate the relative importance of sectoral outside options.

Empirically, our paper is related to a large literature studying how various measures of firm performance affect the wages of incumbent workers (recent examples include Card, Devicienti and Maida, 2011, Guiso, Schivardi and Pistaferri, 2005, and Bell and Van Reenen, 2011; also, see Manning, 2011, for a thorough review). None of these studies are, however, able to isolate the impact of productivity shocks stemming from shifts in the firm-level production function.

A common feature in this literature is the frequent use of sector-level instruments as a tool to handle potentially endogenous adjustments by the firms. Using either sectoral instruments or relying on other shocks with an aggregate component for identification of firm-level effects is also common practice in other literatures where firm characteristics are linked to the wages of employees, including studies on the exporter wage premium,[6] studies of the wage effect of firm-level offshoring,[7] or studies of the links between firm performance and CEO pay.[8] Our stylized model highlights the fact that this practice is inappropriate if the instrument is correlated across firms within the same mobility segment, unless workers' wages are set in a perfectly insulated internal labor market. Although this concern has been acknowledged in the existing literature (also in several of the listed examples), our paper is the first to take the implied bias to an empirical test and explicitly show the link to observed mobility patterns. This link is important since it suggest a useful way to a priori assess the magnitude of the implied bias.

Our analysis draws on a very rich matched employer-employee panel data set from the Swedish manufacturing sector. Data include detailed information on worker and firm characteristics linked to firm-level price indices for the compound of goods that each of the firms sells.[9] This allows us to handle a number of key empirical challenges. Most notably, we are able to derive measures of TFPQ using a firm-level production-function approach.[10] We use the shifts in the firm-level physical production function represented by TFPQ for identification in order to purge the analysis from endogenous adjustments of the scale of production and form firms' relative prices.

Our data also allow us to deal with worker sorting and fixed firm-specific factors

[6] See Greenaway and Kneller (2007), for a survey, and Verhoogen (2008) for a recent example.

[7] E.g. Hummels, Jørgensen, Munch, and Xiang (2011).

[8] E.g. Bertrand and Mullainathan (2001).

[9] Our main sample is composed of single establishment firms. Hence, we use the terms establishment and firm interchangeably in the paper.

[10] In the paper, we will use the term TFPQ when we deflate the (nominal) firm-level output series with firm-level price indices, as opposed to TFPR, which uses sector-level price indices to deflate firm-level output. However, in a strict sense we are not measuring physical output units of a homogeneous good, as in Foster, Haltiwanger, and Syverson (2008).

and analyze the impact of idiosyncratic and sectoral shocks on skill sorting and returns to human capital. Throughout, we exploit the matched employer-employee nature of our panel and estimate the wage impact in models with employer-by-employee (match) fixed effects. This implies that inference is made from time-varying firm-level productivity for ongoing matched worker-firm pairs, effectively allowing us to abstract from both fixed firm-level wage policies, assortative matching, and endogenous match quality. In addition, we are able to measure the degree of complementarity between time-varying idiosyncratic and sectoral productivity and worker human capital by both analyzing to what extent the wage impact varies with the workers' human capital and by documenting how the shocks affect the firms' composition of employees.

To preview our results, we show that wages are affected by changes in technology-induced firm-level productivity. The elasticity of wages to shocks that are shared with other firms within the same sector is about three times larger than the elasticity with respect to purely idiosyncratic shocks, suggesting that the bias in empirical studies that rely on sectoral instruments may be substantial. We also show that the impact of sectoral shocks is likely to be driven by market forces stemming from sector-specific outside options, whereas bargaining institutions are a less likely explanation. Notably, the relative magnitude of the sectoral and idiosyncratic estimates is well in line with the prediction from a simple model that weighs the sectoral outside option according to the observed cross-sectoral mobility patterns of workers who change jobs. Sectoral shocks matter more for workers with a higher predicted probability of returning to the original sector, whereas the impact of idiosyncratic shocks is independent of predicted mobility patterns. In addition, we show that the relative importance of shocks that are shared within a sector is vastly understated when relying on revenue-based measures of TFP, or endogenous labor productivity.

Overall, our analysis thus documents that firms renegotiate wages of incumbent workers in response to both internal and sector-specific external market forces. In contrast, we show that other aspects of firms' human resource policies, such as the skill mix and the relative pay-scale associated with differences in human capital, remain largely unaffected when labor productivity changes due to either idiosyncratic or sectoral technology shocks. This suggests that key elements of firms' human resource policies, such as the within-firm wage dispersion and skill composition, may be considerably more rigid than firms' average wage levels in response to changes in both idiosyncratic factors and outside conditions.

The rest of the paper is organized as follows. First, we present our empirical strategy in Section 2. Details on the construction of the data are provided in Section 3. The main empirical results in the paper are presented in Section 4, and Section 5 discusses variations and robustness exercises. Finally, Section 6 concludes.

2 Method

This section is divided into two parts. First, we outline a model to motivate the wage equation we estimate and discuss the different identification challenges that arise in the attempt to interpret the impact of firm-level productivity on wages as a causal relationship. Secondly, we describe our strategy for deriving physical TFP, and discuss the importance of an appropriate definition of this series.

2.1 Estimating the Wage Impact of Productivity

In this section we outline a model of wage setting that allows us to separately assess the roles played by productivity shocks through both inside and outside forces. Our model is highly stylized and mainly aims to help us interpret the empirical results. Thus, it abstracts from many aspects that may be crucial for firm-level wage setting in general, but which are of second-order importance for our empirical analysis. Following most of the theoretical literature, we allow for an explicit role for productivity outside the firm in shaping workers' wages, by altering worker's outside options. In line with our empirical analysis, we distinguish between productivity in the sector and other aggregate conditions as the two main components in workers' outside option. The most natural justification for such a distinction is sector-specific human capital, which is likely to ease the transition of workers within sectors (Rogerson, 2005).

For simplicity we start of by following the conventions in the empirical rent-sharing literature (see e.g. Manning, 2011) by outlining a model of wage setting where workers are identical and where the total worker surplus is shared equally across workers. Since there are good theoretical reasons for why this may not be a realistic description if firm and worker productivity are complements (see e.g. Eeckhout and Kircher, 2011, for a discussion), we will return to the possibility of heterogeneous effects in the empirical analysis.[11]

We further assume that wages of workers at firm j are set as a weighted average over firm-specific productivity and an outside option, again following the conventions of the canonical rent-sharing model discussed in, e.g., Manning (2011). We enrich this model by specifying the outside option as a mix of a sector-specific component and a general component. With probability ϕ the worker remains in the sector receiving the average sectoral wage W_{st}^{S} in expectation. The worker is severed from the sector with probability $(1-\phi)$, in which case he enters the general labor market,

[11] Given that our empirical analysis suggests that the heterogeneity is of limited importance in this context, we choose to outline the most parsimonious model possible in this section.

receiving expected monetary value Ω_t.[12] Thus, wages are set according to

$$W_{jt} = \beta LP_{jt} + (1-\beta)[\phi W^S_{st} + (1-\phi)\Omega_t], \tag{1}$$

where W_{jt} is the wage of workers employed by firm j at time t, β captures workers' ability to extract rents from firm-specific productivity advancements and LP_{jt} denotes the firm's productivity level (to be defined in more detail below).

Letting LP^S_{st} denote the average productivity in sector s, and abstracting from issues related to firm size, we get the average wage in the sector as the average wage in the relevant firms

$$W^S_{st} = \beta LP^S_{st} + (1-\beta)[\phi W^S_{st} + (1-\phi)\Omega_t]. \tag{2}$$

Thus, combining (1) and (2), and defining the firms' truly idiosyncratic productivity component as the deviation from the sector average ($LP^I_{jt} = LP_{jt} - LP^S_{st}$) we get

$$W_{jt} = \eta_1 LP^S_{st} + \eta_2 LP^I_{jt} + \gamma\Omega_t, \tag{3}$$

where $\eta_1 = \left[\frac{\beta}{1-(1-\beta)\phi}\right]$ measures the wage impact of internal labor productivity and $\eta_2 = \beta$ measures the wage impact of productivity at the sectoral level. The wage impact of aggregate shocks are captured by $\gamma = \frac{(1-\phi)(1-\beta)}{1-(1-\beta)\phi}$. The model thus naturally decomposes productivity movements into a sectoral and an idiosyncratic component with potentially different effects on wages. Clearly, shocks that are shared within the sector have a larger impact on wages than idiosyncratic shocks as long as $\phi > 0$ and the difference is strictly increasing in the value of ϕ. In the empirical section we

[12] Notably, Ω_t will comprise of all elements of the outside option that are shared across workers in the economy, including the (unconditional) transition-probability weighted average of sectoral wages. Thus, ϕ should be interpreted as the difference between the conditional transition probability of re-entering into the original sector and the unconditional transition probability to enter into this sector.

investigate how the impact of idiosyncratic and sectoral shocks relates to workers' predicted values of ϕ.

When taking the model to the data, we first need to specify sectors which, according to the logic of the model, should capture workers' mobility segments. Arguably, any choice of sector will be somewhat arbitrary in this context, and we will therefore assess the robustness of the results using alternative definitions. However, our baseline choice is to define sectors according to the firms' employer organizations since bargaining sectors are constructed in order to cover homogenous labor markets. This also brings to the forefront the potential role played by bargaining institutions, an issue which we return to when estimating variations of the model.[13]

In order to account for movements in the parts of the outside option that are shared across different sectors, Ω_t, we include time dummies, ρ_t, and local labor market tightness, θ_{lt}, in our empirical specifications, where l index local labor markets.

Our stylized model abstracts from worker and firm heterogeneity. In order to handle the potential direct role played by heterogeneity at the individual or firm level (as already noted, we return to the issue of heterogeneous effects), we include a set of fixed effects in the model. Firm fixed effects eliminate any firm-specific characteristic that remains constant over the period of observation, thus removing possible omitted variable biases stemming from time-constant working conditions or wage-setting policies which may affect both firm-level productivity and wages.[14] By also including worker fixed effects we eliminate possible composition biases associated with systematic changes in unobserved characteristics of the employees. Our most stringent specifications interacts the firm and worker fixed effects into a set of

[13] We refrain from modelling sectoral bargaining explicitly as our results suggest a limited role for this mechanism as compared to market forces in the firm-level wage-formation process.

[14] See Daniel and Sofer (1998) for a discussion.

match-specific fixed effects. This allows us to estimate the effect of interest even if, e.g., poor matches are the first to be dissolved in response to negative productivity shocks. In addition, we add controls for time-varying individual characteristics, χ_{it}.

We estimate the model in logarithms due to the well-known log-normal distribution of wages and the fact that estimated elasticities are much more intuitive to interpret than level-estimates.[15] The covariates of interest are the log of the idiosyncratic and sectoral productivity in year t.

All in all, using lower-case letters to denote logs, and allowing for the match-specific effects, ν_{ij}, our empirical model for the log wage of worker i in firm j, w_{ijt}, can be written as

$$w_{ijt} = \eta_1 lp^S_{st} + \eta_2 lp^I_{jt} + \rho_t + \alpha\theta_{lt} + \boldsymbol{\lambda}\chi_{it} + \nu_{ij} + \epsilon_{ijt}, \tag{4}$$

where α and $\boldsymbol{\lambda}$ are coefficients (a vector in the latter case) and ϵ_{ijt} is a worker-firm-year specific disturbance.

It is noteworthy that standard search models in the Mortensen and Pissarides (1994) tradition suppresses the firm and explicitly model employment relationships between workers and jobs, creating an identity between marginal and average productivity. Recently, the literature has instead seen a surge of large-firm models where firms are allowed to have non-constant returns to scale (see e.g. Kaas and Kircher, 2011). A key discussion in the literature concerns whether wages are set according to marginal or average productivity (or a mixture of these as in Stole and Zwiebel, 1996). For obvious reasons, we need to work with average productivity in the empirical analysis, although the estimated elasticities are equal to the impact of marginal productivity under a first-order (Cobb-Douglas) approximation. Im-

[15] Considering the year-specific intercepts we include in the empirical model, it is straightforward to show that the estimated equation (4) provides the log-linear counterpart of the linear equation (3).

portantly, though, our identifying variation comes from measures of TFPQ, which are derived from a production function approach (described below) that allows for non-constant returns to scale. Thus, we isolate the impact of shifts in the production function (hence unaffected by the firms' choice of inputs) that affect marginal and average productivity alike, regardless of returns to scale. Our analysis is therefore less sensitive to endogenous adjustments (potentially altering the firms' relationship between average and marginal productivity) than studies based on, e.g., value added approximations.

We rely on variation in physical TFP (TFPQ in the terminology of Foster, Haltiwanger and Syverson, 2008) for identification since estimation of equation 4 relying on variation in revenue-based TFP (TFPR) or labor productivity may generate biased estimates if firms are subject to shocks to product demand or wages, respectively. Demand shocks in combination with non-constant returns would simultaneously alter productivity and wages if the labor supply curve is upward sloping. Wage shocks are likely to change the capital labor ratio, and hence labor productivity, inducing reverse causality. In contrast, technological progress, if appropriately mapped by TFPQ, shifts the production function providing a source of fluctuations in labor productivity that is unaffected by changes in input usage or the scale of production induced by e.g. shocks to product demand or factor prices.

Importantly, though, our maintained assumption is that wage shocks have a negligible impact on TFPQ, conditional on time-varying worker and firm observable characteristics and the match-specific (employer-by-employee) fixed effects. This implies that we assume that within-firms changes in the usage of efficiency wages over the sample period have a negligible impact on the estimates.[16]

[16] Although we are unable to test this assumption, we take some indirect support from results suggesting that the within-firm variation in our TFPQ series is virtually uncorrelated with other aspects of firms' human resource policies.

We estimate equation 4 by instrumental variables (IV), using TFPQ as an instrument for labor productivity following Haefke, Sonntag, and Van Rens (2008). An alternative strategy would be to use TFPQ directly as the regressor. Although we also present results using this strategy for completeness, we focus on the IV estimates for two reasons. The first is that labor productivity rather than TFP tends to be the focus of existing wage-setting models. The second reason is purely practical. TFPQ needs to be estimated, and using it as an instrument instead of as a regressor is convenient since the inference of the IV estimator is not affected by the use of generated instruments; see Wooldridge (2002) for a discussion.[17] In the end, however, neither qualitative nor quantitative conclusions hinge on this particular choice.

2.2 Measuring TFPQ

The technology series we use is derived using a production-function approach. The underlying idea is that technology can be measured as the residual from a production function once changes in both stocks and variable utilization of the production factors are accounted for.

2.2.1 Measuring TFPQ: Derivation

Let firm j's technology be described by

$$Y_{jt} = f(Z_{jt}K_{jt}, H_{jt}N_{jt}, V_{jt}, M_{jt}, TFP_{jt}), \tag{5}$$

where gross output, Y_{jt}, is produced combining the stock of capital, K_{jt}, labor, N_{jt}, energy, V_{jt}, and intermediate materials, M_{jt}. The firm may also adjust the level of utilization of capital, Z_{jt}, and labor, H_{jt}. Finally, TFP_{jt} is the index of technology that we want to capture.

[17] Using TFP directly implies that the estimator suffers from the well-known generated-regressor problem.

Using small letters to denote logs, taking the total differential of the log of (5) and invoking cost minimization, we arrive at:[18]

$$\Delta y_{jt} = \psi_j [\Delta x_{jt} + \Delta u_{jt}] + \Delta tfp_{jt}, \tag{6}$$

where Δy_{jt} is the growth rate of gross output and ψ_j the overall returns to scale. Here, C_j^F denotes the cost share of factor F in total costs, where $F \in \{K, N, V, M\}$. Then, Δx_{jt} is a cost-share weighted input index defined as $C_j^K \Delta k_{jt} + C_j^N \Delta n_{jt} + C_j^V \Delta v_{jt} + C_j^M \Delta m_{jt}$. Similarly, the change in utilization of capital and labor is denoted by $\Delta u_{jt} = C_j^K \Delta z_{jt} + C_j^N \Delta h_{jt}$.[19] Thus, given data on factor compensation, changes in output, input and utilization, and an estimate of the returns to scale ψ_j, the resulting residual Δtfp_{jt} provides a times series of technology growth for the firm. Note that Δtfp_{jt} reduces to a gross-output Solow residual if $\psi_j = 1$, $\Delta u_{jt} = 0$, $\forall\ j$, and there are no economic profits.[20] Hence, Δtfp_{jt} is a Solow residual purged of the effects of non-constant returns, imperfect competition, and varying factor utilization.

In order to properly identify the contribution of technology, it is also important to distinguish between employees with different levels of education (see e.g. Jorgenson, Gollop and Fraumeni, 1987). Hence, using the same logic as above, we define Δn_{jt} as

$$\Delta n_{jt} = C_j^{LHE} \Delta n_{jt}^{LHE} + C_j^{HE} \Delta n_{jt}^{HE} + C_j^{TE} \Delta n_{jt}^{TE}, \tag{7}$$

where superscripts LHE, HE and TE denote workers with less than high school education, high school education and tertiary education, respectively, and C_j^{EDU}

[18] To make the measurement of technology consistent with an imperfectly competitive labor market, we implicitly assume a timing sequence where the wage is determined first and then the firm takes the wage as given when making its input and production choices.

[19] Here, the cost shares are assumed to be constants. We will return to this assumption later.

[20] The zero-profit condition implies that the factor-cost shares in total costs equal the factor-cost shares in total revenues, which are used when computing the Solow residual.

denotes the cost share of category EDU workers in total labor costs, where $EDU \in \{LHE, HE, TE\}$. Hence, our labor input index will capture changes in the skill composition of the workforce of the firm.[21]

The main empirical problem associated with expression (6) is that capital and labor utilization are unobserved. A solution to this problem is to include proxies for factor utilization. As a baseline, we follow the approach taken by Burnside, Eichenbaum, and Rebelo (1995), who use energy consumption as a proxy for the flow of capital services. This procedure, which is well suited for our manufacturing sector data, can be legitimized by assuming that there is a very low elasticity of substitution between energy and the flow of capital services. This, in turn, implies energy and capital services to be highly correlated (see below for a robustness exercise using a direct capital-stock measure instead in the TFP calculation). Assuming that labor utilization is constant,[22] and including a set of time dummies to capture any aggregate trends in technology growth, τ_t, we arrive at the empirical specification used to estimate technology shocks

$$\Delta y_{jt} = \psi_j \Delta \widetilde{x}_{jt} + \tau_t + \Delta tfp_{jt}, \tag{8}$$

where input growth, $\Delta \widetilde{x}_{jt}$, is defined as $(C_j^K + C_j^V)\Delta v_{jt} + C_j^N \Delta n_{jt} + C_j^M \Delta m_{jt}$. Note

[21] We are, however, not accounting for the contribution to production of the unobservable skills of workers or match quality. Note, though, that although this will affect the technology measures and their estimated distributions, it is not a problem for the estimation of equation (4) as long as the specification includes match-specific fixed effects.

[22] In a related paper, Carlsson (2003) experiments with using various proxies for labor utilization (hours per employee, overtime per employee and the frequency of industrial accidents per hour worked) when estimating production functions like equation (8) on Swedish two-digit manufacturing industry data. Including these controls has no discernible impact on the results of that paper. Moreover, Carlsson (2003) reports that the growth rate of hours per employee is acyclical on the annual frequency. Thus, we are not likely to leave out any important variation in labor input by looking only at the growth rate of the extensive margin.

that Δtfp_{jt} encompasses any firm-specific constant.

2.2.2 Measuring TFPQ: The Importance of Getting the Measures Right

Our key identifying assumption in the IV estimation of equation (4) is that physical TFP is exogenous to individual wages and only affects wages through labor productivity conditional on time-varying worker and firm characteristics and employer by employee-specific fixed effects. Here we illustrate some of the details of the empirical implementation of TFP, and we discuss why some of these details are crucial for this condition to be met. Fundamentally, we show how using alternative measures of TFP, either based on sector-deflated output (TFPR) or on value added rather than gross output, would yield an invalid instrument.[23]

A first point is that it is crucial that nominal output is deflated by appropriate firm-level prices and not by sectoral price indices as is customary. We use firm-level prices aggregated from unit prices for each good the firm produces (see Section 3 for further details), allowing us to derive true volume measures from gross output at the firm-level. Following Klette and Griliches (1996), the problem with the usual approach, which uses a sectoral price index (P_{st}^S) instead of a firm-level price index (P_{jt}), can easily be seen by noting that the measure of real output deflated by sectoral prices would be $\ln\left(Y_{jt}P_{jt}/P_{st}^S\right) = \ln Y_{jt} + \ln(P_{jt}/P_{st}^S)$. Hence, real output deflated by sectoral prices would be a function of relative prices. Assume next that the firm faces a constant-elastic demand function and sets its price as a (constant) markup over marginal cost as in the standard monopolistic-competition model. Since marginal cost, under standard assumptions, is proportional to unit labor cost, the relative price will be a function of wages (see Carlsson and Nordström-Skans, 2012,

[23] Similar problems would, of course, emerge if we studied the direct effects of TFP on wages.

for direct empirical evidence). Importantly, this implies that sales deflated by sectoral prices $\ln\left(Y_{jt}P_{jt}/P_{st}^{S}\right)$, and consequently also the labor productivity and TFP measures derived from it, will respond to idiosyncratic wage shocks. The relationship between sector-deflated labor productivity (or TFP) and wages would then produce upwardly biased estimates of the causal impact of productivity on wages, even if firms are wage takers and produce according to a constant returns to scale technology (in which case marginal cost is independent of the scale of the production).

A second point is that gross output, as opposed to value-added, should be used as the output measure. TFP series derived from standard measures of value-added are only valid under perfect competition and constant returns. Instead, as shown in appendix B, in the case of decreasing returns to scale a TFP measure derived from value-added would be negatively correlated with the growth rate of primary inputs. The drawback from this negative correlation can be easily illustrated in an example. Suppose there is a positive demand shock and the firm has decreasing returns. Profit maximizing firms are likely to respond by increasing production, pushing up the demand for labor, electricity and other intermediate goods. As a consequence of decreasing returns, measured TFP based on VA will decline. If the demand shock has a positive impact on wages, for instance due to an upward sloping labor supply curve, we then expect a negative bias in the wage regressions.

Finally, we use electricity consumption to proxy for variations in the use of capital services as suggested by Burnside, Eichenbaum, and Rebelo (1995) in our main analysis. Although this may not be ideal in all settings, it should provide a good approximation for the manufacturing firms we study. An alternative (explored in section 5.1) is to use a direct measure of the capital stock, derived using the perpetual inventory method, when estimating TFP. A disadvantage of this alternative, in finite samples, is that it requires book values as starting values and these may be poor proxies for physical capital since they tend to be strategically constructed for tax purposes. Using electricity also has the advantage that it accounts for the actual use of the capital stock, i.e. the flow of productive services from capital, since electricity consumption responds to both changes in the stock of capital and changes in capital utilization. The latter (endogenous) term will then end up in the TFP-measure when taking a capital stock approach.

The empirical importance of these measurement issues are all thoroughly examined in Section 5.

2.2.3 Measuring TFPQ: Empirical Implementation

When empirically implementing specification (8), we take an approach akin to the strategy outlined by Basu, Fernald, and Shapiro (2001). First, the specification is regarded as a log-linear approximation around the steady state. Thus, the products $\psi_j C_j^F$ (i.e. the output elasticities) are treated as constants. Note that using constant cost shares (including the cost share(s) of labor) precludes variation in wages to spill into variation in the TFP measure if, for any reason, $\psi_j C_{jt}^N$ is an imperfect measure of the output elasticity of labor input. Second, the steady-state cost shares are estimated as the time average of the cost shares for the two-digit industry to which the firm belongs (SNI92/NACE). Third, to calculate the cost shares, we assume that firms make zero profit in the steady state.[24] Importantly, as noted by Basu and Fernald (1995), zero profits in equilibrium are consistent with a mark-up if the mark-up is equal to the returns to scale. Taking total costs as approximately equal to total revenues, we can infer the cost shares from factor shares in total revenues. The cost share of capital and energy is then given by one minus the sum of the cost

[24] Using the sectoral level data underlying Carlsson (2003) we find that the time average ($1968-1993$) for the share of economic profits in aggregate Swedish manufacturing revenues is about -0.001, thus supporting the assumption made here.

shares for all other factors.

Note that the estimation of equation (8) cannot be carried out by OLS, since the firm is likely to consider the current state of technology when making its input choices.[25] Here, we exploit the panel nature of our firm-level data to use internal instruments, as described in Section 4.

Once the series of technical change has been obtained following equation (8), the next step consists of integrating the growth rates in technology into a log level technology series using the following recursion

$$tfp_{jt} = tfp_{j0} + \sum_{r=1}^{r=t} \Delta tfp_{jr}. \tag{9}$$

Note that the initial level of technology (tfp_{j0}) is a firm-specific constant that is not observed, but will be captured by firm fixed effects in the second stage estimation.

2.3 Sector vs. Idiosyncratic Productivity

In the empirical specification of equation (4) we distinguish between sectoral, lp^S_{st}, and idiosyncratic, lp^I_{jt}, labor productivity. These measures can easily be obtained by running a regression of firm-level labor productivity, measured as gross output per worker, on sector-specific time dummies. The projection from the sector-specific time dummies in this regression is then a measure of lp^S_{st}, and the residuals are a measure of lp^I_{jt}. We use employee weights when running this decomposition, such that sector-specific productivity is the average employee-weighted productivity, and idiosyncratic productivity is the firm-level deviation from this average. In an

[25] This is the so-called transmission problem in the empirical production-function literature (Marshak and Andrews, 1944). Technology change (i.e. the residual) represents a change in a state variable for the firm and changes in the level of production inputs (the explanatory variables) are changes in the firm's control variables, which should react to changes in the state variable. In this case there will be a correlation between the error term and the explanatory variable, hence the need of IV methods.

analogous fashion, we decompose the TFP series derived from (9) into a sectoral, tfp_{st}^{S}, and an idiosyncratic TFP component, tfp_{jt}^{I}.

3 Data

We combine three data sources to construct our sample. The employer side of the data set is primarily drawn from the Statistics Sweden Industry Statistics Survey (IS) and contains annual information for the years 1990-1996 on inputs and output as well as geographical location for all Swedish industrial (manufacturing and mining) plants with 10 employees or more and a sample of smaller plants (see appendix A for details).[26] Our baseline specification focuses on continuing single-plant firms, which offers the cleanest identification of the issues at hand. That is, we are not forced to address potential differences in wage-setting practices across plants within the same firm or issues with firms consisting of plants pertaining to different sectors. Neither do we need to consider heterogeneous outcomes in terms of productivity across plants within the firm. The focus on continuing plants can also mitigate possible selection effects due to firm demographics associated with productivity shocks. In the robustness section, however, we show that our main results hold also in an unbalanced version of the data set where we include multi-plant firms as well.

A crucial feature of the IS data we use is that it includes a firm-specific producer price index constructed by Statistics Sweden.[27] The firm-specific price index is a chained index with Paasche links that combines firm-specific unit values and detailed disaggregate producer-price indices (either at the goods level, when available, or at the most disaggregate sectoral level available). Note that in the case in which a firm-specific unit-value price is missing (e.g., when the firm introduces

[26] The availability of detailed factor-input data limits the sample years to 1990-1996.

[27] For this Statistics Sweden use data from the Industrins Varuproduktion Survey.

a new good), Statistics Sweden uses a price index for similar goods defined at the minimal level of aggregation (starting at 4-digits goods code level). The disaggregate sectoral producer-price indices are only used when a plausible goods-price index is not available. Thus, the concern raised by Klette and Griliches (1996) regarding biased returns to scale estimates when sectoral price deflators are used in the computations of real gross output should not be an issue here. We use this price index to deflate output both when constructing labor productivity and when deriving the TFP series.

The employee side of the data is obtained from the Register Based Labor Market Statistics data base (RAMS) maintained by Statistics Sweden. This data contains information on annual labor earnings for all privately employed workers in Sweden. The raw data was compiled by the Swedish Tax Authority in order to calculate taxes. The data includes information on annual earnings, as well as the first and last remunerated month received by each employee from each firm. We use this information to construct a measure of monthly wages for each employee in each of the firms in our sample, closely following the procedures of Nordström Skans, Edin, and Holmlund (2009) and Carlsson and Nordström-Skans (2012). The data lacks information on actual hours, so in order to restrict attention to workers reasonably close to full time workers we only consider a person to be a full-time employee if the (monthly) wage exceeds 75 percent of the mean wage of janitors employed by municipalities.[28] We only include employment spells that cover November following the practice of Statistics Sweden. We focus on primary jobs and therefore only keep the job resulting in the highest wage for workers with multiple jobs. The data

[28]Using the same procedure with RAMS data, Nordström Skans, Edin, and Holmlund (2009) found that this gives rise to a computed wage distribution that is close to the direct measure of the wage distribution taken from the 3 percent random sample in the LINDA database, where hourly wages are the measure of pay.

also includes information on age, gender, education, and immigration status of the individual workers.

Unemployment and vacancy data at the local labor market level for November is collected from the National Labor Market Board (AMS). Here, we rely on the 1993 definition of homogenous local labor markets constructed by Statistics Sweden using commuting patterns, which divides Sweden into 109 geographic areas.

Note that we use the labor-input measure available in IS to compute labor productivity, whereas the labor-input measures used when estimating TFP are taken from RAMS. As mentioned above, the IS employment data is based on a survey collected by Statistics Sweden, whereas the RAMS employment data is based on the income statements that employers are, by law, required to send to the Swedish Tax Authority. Since the IS and RAMS measures of labor input are independently collected it is highly unlikely that any measurement errors are common in the two. This, in turn, is important for ruling out that any observed relationship between labor productivity and technology is only due to common measurement errors in the labor input measures, as well as ensuring that potential measurement errors will not be transferred to the projection of labor productivity from the first step of the IV-procedure.

Both RAMS and IS provide unique individual and firm identifiers that allow us to link the employees to each of the firms in the sample. Since the RAMS data covers the universe of workers, we observe every worker employed in each of the IS firms during the sample period. Given the restrictions mentioned above and after standard cleaning procedures (see appendix A for details), we are left with a balanced panel of $1,136$ firms observed over the years 1990-1996 and $472,555$ employee/year observations distributed over $106,050$ individuals. Our used data set covers 11 percent of the full-time employees in the manufacturing sector.

4 Estimation Results

4.1 Estimating TFP

We first estimate the technology disturbances relying on the empirical specification (8) outlined in section 2.2.1 above. Here, we allow the returns to scale parameter ψ_j to vary across durables and non-durables sectors as suggested by Basu, Fernald, and Shapiro (2001).[29] The models include firm fixed effects, which capture any systematic differences across firms in average technology growth. Since the firm is likely to consider the current state of technology when making its input choices, we need to resort to an IV technique. Following Carlsson and Smedsaas (2007) and Marchetti and Nucci (2005), we use the difference GMM estimator developed by Arellano and Bond (1991) and report robust, finite-sample corrected, standard errors following Windmeijer (2005). Here we use $\Delta\widetilde{x}_{jt-s}$, for $s \geq 3$, as instruments and collapse the instrument set in order to avoid overfitting (see Roodman, 2006).[30]

In table 1, we present the estimation results for equation (8). The estimate of the returns to scale for the durables sector equals 0.99, and 0.88 for the non-durables sector, but both are somewhat imprecisely estimated (s.e. of 0.19 and 0.22, respectively). It is reassuring, however, to see that the point estimates of the returns to scale are very similar to estimates reported by earlier studies. For example, Basu, Fernald, and Shapiro (2001) reports estimates of 1.03 and 0.78 for durables and non-durables, respectively, using U.S. sectoral data. Moreover, the Hansen test of over-identifying restrictions cannot reject the joint null hypothesis of

[29] The data does not allow us to identify the returns to scale parameter separately across (SNI92/NACE) two-digit industries since many sub-samples become too small.

[30] Given that we use a difference GMM estimator, the second and higher ordered lags of $\Delta\widetilde{x}$ should be valid instruments under the null hypothesis of no serial dependence in the residual. However, when including the second lag in the instrument set, the Hansen test of the over-identifying restrictions is significant at the five-percent level.

a valid instrument set and a correctly specified model.

Table 1 shows that the AR(2) test of the differenced residuals (see Arellano and Bond, 1991) shows no sign of any important serial dependence in the estimated technology change series.[31] This implies that the innovations have a highly persistent effect on the level of technology. When estimating an AR(1) process for the level of technology, as in e.g. Eslava, Haltiwanger, Kugler, and Kugler (2004), we find a persistence estimate of 0.88 (s.e. 0.012).[32] This estimate is in between the Colombian estimate of 0.92, presented in Eslava, Haltiwanger, Kugler, and Kugler (2004), and the U.S. estimate of 0.79, presented in Foster, Haltiwanger, and Syverson (2008).

4.2 The Impact of Productivity on Wages

Before moving into the main results of the paper, we provide a brief description of the distribution of wages, TFP and labor productivity. Summary statistics are available in table 2. First of all, note that the dispersion of productivity is much larger than the dispersion of wages, but that this relationship to a large extent is driven by large differences between firms. The variance (over time) within an employment spell (i.e. a match between a worker and a firm) is about equal for the two variables. In the analysis we distinguish between a sectoral and an idiosyncratic component as discussed above. The sectors are identified following the 16 employer federations that sign collective agreements in the manufacturing sector.[33] When decomposing productivity within and between sectors we see that the within-match

[31] This is important for the validity of the instruments used in the estimation of the returns to scale. Note also that we take an additional lag of the instrument set than would be required under the null of no serial dependence in the technology change series, which further safeguards from any deviations from the null.

[32] Note that this standard error is not appropriate to use to test the null of a unit root, since such an hypothesis would imply a different parameter distribution.

[33] In practice we allocate the firm to the most common employer federation among firms in the same five-digit industry according to the standard SNI92 (NACE) classification.

standard deviation of idiosyncratic firm-level productivity is more than three times larger than the variance of sectoral productivity.

We proceed by investigating the role of sectoral and firm idiosyncratic productivity on individual wages, following equation (4). The first column in table 3 shows the results from a simple OLS regression that relates labor productivity to individual wages controlling for firm-level fixed effects. Column 2 shows the same specification, but using idiosyncratic and sector-level TFPQ as instruments for the two labor productivity measures. Column 3 adds a third-order age polynomial and worker fixed effects. Column 4 presents our most stringent specification, including match-specific fixed effects. Column 5 repeats the last exercise but uses TFPQ (idiosyncratic and sectoral) directly as regressors instead of labor productivity. All standard errors are robust to intra-firm correlation.

The overall impression is that both firm-level idiosyncratic labor productivity movements and movements that are shared within a sector matter for wage determination, but the impact of the latter is much larger. In order to obtain the second conclusion it is fundamental to properly handle firms' endogenous adjustments of labor productivity. The point estimates of the OLS model in column 1 suggest that idiosyncratic productivity shocks have the same impact as shocks that are shared within a sector (0.033 versus 0.027), although sectoral shocks are not statistically different from zero, but when we isolate the effects of shifts in the production function by relying on identifying variation from TFPQ movements in column 2, we find an elasticity of wages to sectoral productivity that is substantially larger than the elasticity with respect to idiosyncratic productivity (0.123 compared to 0.032). In table 4, we show the first-stage regressions of the IV-specification.[34] Both estimated

[34] The values of the F-statistics are well above 10, suggesting that our instruments are not weak. More importantly, the Kleibergen-Paap rk LM statistic (see Kleibergen and Paap, 2006), which is presented at the bottom of table 3, clearly reject the null hypothesis of underidentification.

coefficients in column 2 are statistically different from zero at the 1 percent level and the p-value for a two sided test for equal coefficients is 1.6 percent.[35] In Section 5 below we return to potential explanations for why the OLS and IV results might differ for sectoral and idiosyncratic productivity.

The rest of table 3 shows that the results are somewhat larger when we account for worker and match quality. Column 3 accounts for individual observed and unobserved heterogeneity by means of an age polynomial and individual fixed effects. The idiosyncratic component increases to 0.050, while the sectoral component increases to 0.149. The results are virtually identical if worker and firm effects are replaced by worker-by-firm match fixed effects as shown in column 4. The latter may be a result of the fact that we are using a fairly short panel and only a subset of the economy, which means that the individual fixed effects in many cases are identified from single spells (i.e. that the match fixed effects are already captured in the model with worker and firm fixed effects). The p-values for two-sided tests of equal coefficients from sectoral and idiosyncratic shocks in columns 3 and 4 remain at 1 percent.

In the final column, we show the direct effect of TFPQ on wages, paralleling the specification in column 4 but using OLS with TFPQ as the regressor. We find very similar, but marginally smaller (0.124 and 0.042), effects compared to column 4. This is not surprising considering that the the first stages of the IV regressions (table 4) is close to, but somewhat smaller than, unity.[36] The differences between the coefficients remain statistically significant. Since TFP is a generated regressor, the

[35] In a similar vein, Fuss and Wintr (2009) find that firm-level average labor compensation is more reactive to sectoral than to firm-level changes in TFPR in Belgium.

[36] Note that the fact that the first stages are close to unity suggests that the endogenous response of labor productivity through changes in input usage is not very large. This is reassuring since we are using a balanced panel and, due to the relatively short period available, are unable to model exits of firms. Note though that Section 5 shows robustness checks without balancing the panel.

model which uses TFP directly as the regressor is, however, sensitive to potentially attenuating measurement errors and we therefore consider the IV specification as the preferred model.

In Section 5 below, we present a large number of alternative specifications using alternative measures and definitions. To preview our results, in all cases we find larger effects for shocks that are shared within a sector as long as we focus on shocks that stem from changes in TFPQ.

Regarding the magnitudes of the estimates, it is clear that elasticities for both idiosyncratic and sectoral shocks are far from unity. For the interpretation of the estimates it is, however, crucial to note that we are analyzing productivity shocks at the firm level, rather than match level as in the standard Mortensen-Pissarides model. This implies that the productivity advances remain with the firm if the worker is replaced, thus effectively reducing the workers ability to extract rents relative to the one-worker-per-firm case with match specific productivity (see e.g. Faberman and Nagypál, 2008, for a discussion). It is also important to bear in mind that the variance of the underlying productivity processes is relatively large. This is especially true in the case of idiosyncratic firm-level productivity. Removing variation between firms and using our preferred estimates in column 4 of table 3, we find that an increase of one standard deviation in either of the productivity measures (sector or idiosyncratic) raises wages by about one-quarter of the average real wage growth in our sample.[37]

[37] Note that our estimated elasticities net out aggregate productivity effects, since all regressions include time dummies and labor market tightness. Considering that average real wage growth within the manufacturing establishments included in the sample is 2.4%, the estimated impact of one s.d. idiosyncratic productivity on wages amounts to 28% ($0.051 * 0.130/0.024$) of this average real wage growth, while the impact of one s.d. sectoral productivity is 22% ($0.149 * 0.036/0.024$).

4.3 Sectoral Productivity, Bargaining Power and Outside Options

The difference in estimated sectoral and idiosyncratic effects implies that workers extract more rents when productivity advancements are shared with other firms within the sector. In the stylized theoretical model outlined in Section 2, we related the difference in impact to the excess probability of returning to the same sector (ϕ in equation 3). To obtain a quantitative mapping between our estimates and the model we calculate the empirical counterpart of ϕ relying on all full-time employment spells in RAMS who result in a change of employer between adjacent years within our 16 bargaining sectors during the sample period 1990-1996 (120,652 observations). We then calculate the difference between the conditional probability of staying within the same sector when changing employer and the unconditional probability of entering each sector when changing employer (regardless of the sector of origin), and arrive at an estimate of the excess probability of returning to the same bargaining sector of 0.54. Using equation (3) and the estimates of column 4 in table 3 to calculate the implied value ϕ instead gives an estimate of 0.70, with a 95 percent confidence interval ranging from 48 to 91 percent.[38] This suggests that the model, despite its very stylized nature, produces quantitative predictions that are in line with observed behavior in data.

Our baseline specification relies on information on the bargaining sector of the firm since we wish to group firms according to the markets from where the firms draw their labor. However, in Sweden, as in many other OECD countries, wage bargaining has a multitiered structure with negotiations at the industry level as well as at the firm level which opens up to alternative interpretations of the results. To test whether the sectoral effects stem from market forces or higher bargaining

[38] Defining a residual bargaining sector for all workers outside the 16 bargaining sectors of the sample and using all data on transitions (660,859 observations) yields a very similar estimate of $\phi = 0.52$.

power at the sectoral level, we constructed a transition-weighted sectoral productivity index for each bargaining sector. To get the value of the index for sector A, we take the fraction of workers moving from sector A to sector B as a weight for the productivity of sector B, and then repeat for all 16 sectors in the sample (including inbreeding within sector A itself) and finally calculate the weighted sum of the sectoral productivities.

By using these transition probabilities we get a weighted productivity index that allows us to check if the bargaining sector is more important than other sectors, conditional on empirical transition patterns. As can be seen in column 2 of table 5, adding this weighted sectoral labor productivity index, $lp_{st}^{S,Weighted}$, in the baseline regression (using a similarly weighted sectoral technology index as an instrument) brings down the point estimate of the sectoral effect to close to the point estimate of the idiosyncratic effect (from 0.149 to 0.029). Moreover, the point estimate of the weighted sectoral index is about four times larger than the point estimate of the idiosyncratic effect (0.199 vs. 0.051). Thus, the results are more in line with an outside option interpretation than with a sectoral bargaining story, although the statistical precision becomes an issue since the two measures are highly correlated due to the high frequency of intra-sectoral mobility (running a regression of one on the other yields a within-sector R^2 of 0.93).

In the third column of table 5, we drop the sectoral productivity index and obtain a statistically significant estimate of the effect from the weighted productivity index (0.244), which is larger than that of the baseline specification with only the bargaining sector. The fact that the point estimate is larger for the transition-weighted variable than for the bargaining sector variable lends additional support to the market forces interpretation of the results since shifting to the weighted sectoral index only adds measurement errors to the sectoral variable if interpreted as a proxy for the bargaining sector. Similarly, in the fourth column we instead rely on two-digit industry according to the SNI92/NACE classification to capture the sectors and find somewhat larger sectoral results than in the baseline specification. Our overall conclusion is therefore that the additional impact from shocks that are shared within a sector primarily is driven by effects through workers' outside options.

As is evident from equation (3), the relative impact of sectoral shocks should be larger, the higher the excess probability of returning to the same sector (ϕ) if market forces are driving the results. In order to test this prediction, we interact the variables of interest with measures of worker-specific predicted probabilities of returning to the same sector. To this end, we use the data on job changers described above and estimate a model where we predict the probability of returning to the same sector using detailed information about the field (three-digit ISCED 97) and level (two-digit ISCED 97) of schooling alongside gender, an age polynomial of order three, immigration background and year dummies.[39] Based on the parameters of this model we make an out-of-sample prediction of return mobility and interact it (after rescaling it as deviations from the mean) with the variables of interest. The results, displayed in the last column of table 5, show that the impact of idiosyncratic shocks is independent of predicted return mobility, whereas the impact of sectoral shocks has a strong and significant positive relationship with predicted return mobility. The magnitude of the estimated interaction effect (1.81) implies that the sectoral effect varies from 0.074 to 0.208 within the range of two standard deviations of predicted return mobility.[40]

[39]The model also includes dummies for the sector of origin, but the impact of these is not included in the prediction.

[40]The standard deviation of the prediction is 3.7 percent

4.4 Complementarities and Sorting

So far, we have followed the standard route of the empirical literature and assumed an homogenous impact of productivity across worker types. However, as pointed out by e.g. Eeckhout and Kircher (2011), production complementarities suggest that better workers should benefit more from working in more productive firms. To the extent that this argument holds for within-firm changes in productivity, we would also expect to see a larger wage impact among high-skilled workers. To test this hypothesis, we have reestimated the model allowing for heterogeneous impacts across the human-capital distribution. We first analyzed heterogeneity over different education levels finding no significant effects, although the precision is fairly poor; see column 1 of table 6.

To construct a more precise test, column 2 of table 6 interacts the variables of interest with predicted human capital ($\hat{w}_{ijt}$). To construct a measure of predicted human capital we estimate a Mincer-type wage equation including detailed information about the field (three-digit ISCED 97) and level (two-digit ISCED 97) of schooling alongside gender, an age polynomial of order three and immigration background. We use the prediction from this model, which explains 40 percent of the within-firm wage dispersion, as our indicator of human capital. The model also includes a firm fixed effect and year dummies but the impact of these variables are not included in the human capital prediction.

The results, displayed in the second column of table 6 again suggest no statistically significant heterogeneity. The general impression is therefore that the wage impact of both idiosyncratic and sectoral TFPQ shocks is homogenous.[41]

[41] Splitting the data according to the median or quartiles of predicted human capital instead gives the same impression. We have also explored models that define human capital according to workers' portable earnings capacity, as measured from person effects from an estimated Abowd, Kramarz, and Margolis (1999) model, using data from a pre-analysis period 1985-1989, again finding no

Turning to sorting, it is worth noting that our main estimates displayed in table 3 are only marginally affected by the inclusion of individual-specific fixed effects. This suggests that compositional effects through firm recruitment and firing policies as a response to technology-induced changes in firm-level productivity should be minor. In order to investigate this issue further, we relate our measures of firm- and sector-level productivity to measures of the employed workers' earnings capacity. We use the predicted human capital (constructed as above) of the workers as the outcome in regressions corresponding to column 3 of table 3.[42] The estimates, presented in column 3 of table 6, show no signs of assortative matching between observed human capital and time-varying idiosyncratic or sectoral productivity movements. An important aspect of the analysis is that the metric is comparable to that of the main analysis. Thus, given the confidence interval of the estimates we are able to conclude that any effects on the sorting of workers must be substantially smaller than the endogenous wage responses which we focus on in the main analysis.[43]

The apparent lack of assortative matching responses may appear at odds with a number of studies documenting assortative matching between workers and jobs (e.g., Eeckhout and Pinheiro, 2012). Two distinctions relative to parts of the existing literature should be noted here. First, the productivity measures we use are at the firm level and not at the job level. Hence, we cannot attribute productivity differences to each individual job, which is fundamental to assess the importance of sorting in search models where the commonly held assumption is one-firm-is-one-job

indications of heterogenous effects related to measured human capital.

[42] Under the null hypothesis of no sorting, we can estimate the model using predictions from a model without time-varying labor productivity. We have reestimated the model including time varying productivity as well, finding identical results of interest.

[43] We have also explored this the impact on sorting relative to workers' portable earnings capacity as measured from person effects from an estimated Abowd, Kramarz, and Margolis (1999) model using data from a pre-analysis period, finding similar results.

(see Eeckhout and Kircher (2011) for a discussion). Our analysis is more closely related to the part of the literature which attempts to provide a causal estimate of the impact of firm-level productivity on the skill mix of organizations such as Haltiwanger, Lane, and Spletzer (1999) and Mendes, Van Den Berg, and Lindeboom (2010).[44] The second distinction is that we rely on changes in productivity and not productivity levels. An alternative interpretation is therefore that firms' personnel policies change infrequently, and that changes in TFPQ therefore either are transmitted into changes in the skill mix with a lag which is beyond our empirical horizon or remain deeply embedded in technology choices made when the firm was initially created. In line with this interpretation Haltiwanger, Lane, and Spletzer (1999), who use matched employer-employee data for the US, find that while the skill distribution within establishments is tightly linked to the average sales per worker, there is virtually no relationship between changes in labor productivity and changes in the worker mix.

5 Robustness, Measurement and Variations

5.1 Robustness

In this sub-section we discuss a number of robustness checks in different dimensions. First, for reasons discussed in Section 3, our main analysis focuses on single-plant continuing firms which gives us a non-representative distribution of plants relative to the economy. To evaluate the empirical significance of this particular choice of sample, we have redone the analysis on an unbalanced panel of plants that also include plants that belongs to multi-plant firms.[45] This results in a sample of four

[44]Theoretically, Eeckhout and Pinheiro (2012) present a model with multi-worker firms and analyze the relationship between the skill distribution within the firm and TFP.

[45]Note, though, that since we redo the whole analysis on the broader data set, the econometric procedure requires that the plants are in the data long enough to identify the (minimum) number

times as many worker/time observations (2,048,555) and twice as many firms (2,444). As can be seen in column 2 of table 7, the results are qualitatively unchanged.[46]

Our analysis keeps the cost shares and the overall returns to scale parameter constant, which implies that we rely on a log-linear approximation of the production technology. A second order approximation could be accomplished by constructing a Törnqvist index by using the average of observed cost shares of adjacent observations (current and lagged) combined with an assumption that the overall returns to scale parameter is time-invariant.[47] We have therefore redone the analysis using a Törnqvist index and find very similar results; see column 3 of table 7. Note, though, that, as pointed out by Basu and Fernald (2001), to legitimize the use of cost shares at all, factor prices must equal the shadow values of the factors to the firm. In a world with adjustment costs or long-term relationships this will only be true in the long run, or on average, which motivates the use of average cost shares as the baseline specification.

Another issue is the sensitivity of the results to the estimated returns to scale parameter. Table 7 present results from models which perturb the estimates up or down with 0.1 (Columns 4 and 6) and when imposing constant returns to scale in both the durables and the non-durables sectors (column 5). None of these exercises affect our conclusions.

Finally, we have experimented with alternative ways to estimate the productive contribution of capital by relying on explicit estimates of the capital stock instead of using electricity to proxy the flow of capital services (appendix C provides the details

of lags required for the Arellano-Bond estimator.

[46] As an additional robustness check we have redone the analysis without any trimming of the data and the results are, again, qualitatively unchanged.

[47] Diewert (1976) shows that a Törnqvist approximation is "superlative", in other words it is exact if the underlying production function is translog; otherwise it provides a second-order approximation to any functional form.

of the construction of the alternative TFP measure). Estimating capital relaxes the assumption of perfect complementarity between the flow of capital services and electricity use we made in our preferred TFP series. Nevertheless, measuring the flow of capital services with the stock of capital has its own problems. First, this approach ignores variation in capital utilization and secondly, measuring the stock of capital is problematic with an imprecise measure of the initial stock combined with a short sample in the time dimension. However, it is reassuring that the two approaches deliver a very similar message with the estimated elasticities only slightly smaller than those reported above. We find an elasticity of 0.043 (s.e. 0.010) with TFP based on capital for the idiosyncratic productivity component and 0.115 (s.e. 0.031) for the sectoral component.

Our overall impression is that the results are robust to fairly large variations in the estimated models. In particular, we find that the measured impact for sectoral shocks remain larger than the impact of idiosyncratic shocks throughout a wide set of variations (p-values for two-sided tests of equality between sectoral and idiosyncratic effects remain well below 5 percent in all columns).

5.2 Measurement

In Section 2.2.2, we argued that wage shocks will transmit into measured real output series if sectoral prices are used when deflating sales to obtain real output, generating a positive bias in the estimated impact of idiosyncratic productivity. This conjecture is confirmed by results presented in table 8 where we use three-digit PPI deflators instead of firm-level prices to derive gross output resulting in an estimated elasticity of idiosyncratic productivity of almost twice the size (0.092) of the benchmark. Thus, as when relying on endogenous labor productivity without instruments, failing to deflate by a proper price index (i.e. relying on TFPR) would have led us to underestimate the relative importance of outside options.

Appendix B also shows the impact of using a measure of value added instead of gross output to derive the productivity and technology series. As discussed in Section 2.2.2, the main problem with value added-based measures of TFP is that they will be negatively correlated with the intensity of the use of primary inputs, including labor, if there are decreasing returns to scale. Column 3 of table 8 presents the results of using a value-added Solow residual to instrument for value-added labor productivity.[48] The value-added estimates are considerably smaller than those based on TFPQ, as expected if demand shocks have a positive impact on wages. The negative bias is somewhat larger for the sectoral elasticity, suggesting that the wage effect of labor demand is more likely to be seen when demand shocks are shared within sectors. Moreover, the bias is larger in the non-durables sector as compared to the durables sector (columns 4 and 5), as expected, since the point estimates of the returns to scale points towards the returns to scale being lower in this sector.

5.3 Returns to scale, OLS and IV

The estimated impact of sectoral productivity developments on wages in the IV specifications is much larger than in OLS regressions with measures of labor productivity as the regressors of interest. A simple yet plausible explanation for such differences is attenuation bias due to measurement errors, but in that case it would be expected that there is a similar gap between IV and OLS estimates for the idiosyncratic productivity shocks. As this bias is not found, our results are most likely indicative of an endogenous negative association between labor productivity and wages at the sectoral level.

One straightforward explanation is a combination of decreasing returns to scale

[48] Note that the sample is slightly smaller due to negative value-added observations.

and an upward sloping labor supply curve, the intuition being that when firms choose to scale up production (e.g., in response to demand shocks) they will endogenously lower labor productivity if returns to scale are decreasing. The resulting increase in demand for labor will lead to higher wages if the supply curve facing the sector (or firm) is upward sloping. This may explain why instrumentation matters specifically at the sectoral level and not at the idiosyncratic firm level, since wages may be pushed up more in response to sectoral adjustments if firms within a sector compete over a restricted set of workers. Hence, increased demand for labor within a sector may raise wages while single firms may be allowed to hire freely without affecting wages in the market.

While the combination of decreasing returns with an upward sloping wage-setting curve is consistent with our main results, we also try to provide a piece of somewhat more direct evidence by tentatively investigating the role of returns to scale. As shown previously, estimated returns to scale vary between the manufacturing plants producing durable goods (decreasing returns) and those producing non-durables (almost constant returns). Although our estimates of the returns to scale are imprecise, similar differences between durables and non-durables have been previously found in the literature (e.g. Basu, Fernald and Shapiro, 2001). Hence, we expect the gap between IV and OLS estimates for sectoral productivity to be larger in firms operating in durable goods sectors than in firms operating in the non-durables sectors.

In table 9 we estimate our preferred model (i.e. with match-specific fixed effects) using OLS and IV separately for firms with decreasing and constant returns. The results are consistent with the proposed hypothesis. The entire difference between the sectoral OLS and IV estimates stems from the firms facing decreasing returns in our sample, while differences within firms facing constant returns are negligible and non-statistically significant. In the case of firms with decreasing returns, we see that the elasticity of wages to sectoral shocks becomes highly significant and more than four times larger in the IV specification (0.140 vs. 0.027 in OLS). Interestingly, we also see that instrumentation leads to a non-negligible increase in the estimate of the idiosyncratic productivity effects on wages (from an elasticity of 0.033 in the OLS specification in column 1 to 0.052 in column 2).[49] This suggests that aggregation over the sectors also blurred an important role for scale adjustment in response to idiosyncratic productivity in the sectors where returns are decreasing.

5.4 The Role of Dynamics

The specifications we have presented so far are static, i.e., they assume that the wage impact of technology-driven innovations in productivity is immediate. In reality, highly persistent technology shocks may require some time to be absorbed by wages. In order to assess the importance of potential delays in the impact of productivity, we have estimated models with lagged productivity.

Estimates from specifications with lagged productivity are presented in table 10. We concentrate on our preferred specification (including match-specific fixed effects) and proceed parsimoniously, first introducing one lag in column 2 and two lags in column 3.[50] The bottom of the table shows the long-run accumulated effect (for a permanent change in productivity) and its associated level of significance. The results show that there is a role for lagged productivity in shaping current wages. The precision does, however, seem to deteriorate fairly rapidly for the sectoral productivity and we never find the individual lags to be statistically significant. Although the individual lags are estimated with poor precision, the long-run elasticity remains

[49] The differences between the IV and OLS elasticities in Columns (1) and (2) are statistically significant at the 5% level. The p-values of one-sided tests are 0.044 in the case of idiosyncratic productivity and 0.004 in the case of sectoral productivity.

[50] Given the short nature of our panel we were not able to estimate models with more than two lags to any precision.

statistically significant in all cases. The magnitude of the long-run impact is about twice as large as the contemporaneous impact for both the idiosyncratic effect (0.091 vs. 0.051) and the sectoral effect (0.303 vs. 0.149) when two lags are considered.

5.5 Additional Variations

Our data does not allow us to properly control for part-time work, but since part-time work in Sweden is very rare among males in the manufacturing sector we have reestimated the model using only males. The estimates (available on demand) show that the response of male wages to changes in productivity is very similar to that obtained in the overall sample. The elasticity of wages to sectoral productivity is almost three times as large as the elasticity to idiosyncratic movements in productivity, both estimates being statistically significant at standard levels of testing.

A very active literature (Shimer, 2005, Hall, 2003 and Pissarides, 2009) discusses whether search models can be reconciled with the large fluctuations we see in unemployment over the cycle. A key element in this debate is the exact modeling of how wages for new hires react when productivity changes (see e.g. Haefke, Sonntag and Van Rens, 2008 and Pissarides, 2009). We have analyzed the impact on incumbents and new hires of sectoral and idiosyncratic productivity, but found no significant differences in their impact. However, it should be acknowledged that the interacted estimates were quite imprecise.

Finally, we have analyzed whether productivity has a differential impact depending on whether the shocks are positive or negative, where one might suspect that negative shocks have a smaller effect due to downward nominal wage rigidity. We find no evidence of such asymmetries. Although this may seem surprising, it should be noted that the magnitudes of the estimated elasticities are such that the wage impact of any "normal" shock is smaller than the average nominal wage increase among incumbent workers. This implies that there is indeed scope for a symmetric impact of positive and negative productivity shocks, even if nominal wages never fall.

6 Conclusions

We have studied how individual wages are affected by the changes in productivity of the firms where the workers are employed and how those impacts vary when technology is shared within firms operating in homogeneous sectors. By relying on a carefully constructed measure of physical total factor productivity (TFPQ), we isolate changes in firm-level productivity that are caused by shifts in the firms' production functions and analyze how these changes affect individual wages. In addition, we use matched employer-employee data to purge the analysis of sorting on both the supply and demand side.

Our findings suggest that firm-level productivity has an impact on workers' wages, which contrasts to simple frictionless competitive models where individual wages only depend on aggregate labor market conditions and individual skills. Importantly, we show that changes in productivity that are shared within a sector have a three times larger impact on wages than purely idiosyncratic innovations in the short run and in the long run. The results therefore suggest that both workers and firms benefit from firm-level technological advancements, but workers are able to extract a substantially larger fraction of the rent generated by productivity advancements that are shared with other firms in the sector. Since the standard deviation of idiosyncratic (within-match) productivity is about three times larger than that of sectoral productivity, they play a similar role in shaping workers' wage increases: a one standard deviation increase in either sector-specific or idiosyncratic productivity has a wage impact amounting to about one quarter (half) of the average yearly wage growth of incumbent workers in the short (long) run.

We provide a stylized model of mobility and wages that shows how differences in the relative magnitudes of idiosyncratic and sectoral productivity shocks on wages may be explained through worker's outside options. The predictions of the model are well matched by the observed degree of excess mobility within sectors in the data. A worker who changes employer has a 54 percentage points higher probability to find a new job in the original, narrowly defined sector. This strong tendency for intra-sector mobility provides a straightforward rationale for sector-specific outside options to affect the wages in ongoing employment relationships. First, we show that a mobility-weighted productivity index has a larger impact on incumbent wages than a sectoral productivity measure that disregards predicted mobility patterns. Second, we show that sectoral shocks matter more for workers with a higher predicted probability of returning to the original sector, while idiosyncratic firm-level shocks are unaffected by typical mobility patterns.

In spite of having a sizable impact on wages, we find that idiosyncratic and sectoral technology shocks appear the be shared fairly equally across skill groups, regardless of how human capital is measured. In addition, we provide an analysis of how technology-induced changes in labor productivity affect workers' sorting patterns, which reveals that systematic sorting of workers in response to both idiosyncratic and sectoral factors is of minor importance in this context. This suggests that studying the conditions under which new firms choose their initial human capital distributions, and their human resource policies more generally, may be key for understanding the role of firms in the joint evolution of wage dispersion and skill sorting. Once the firm is established, we observe little changes in the skill mix and pay scale as a consequence of movements in technology.

The importance of workers' outside options for wage determination documented here further highlights an issue for consideration in future empirical research. On the one hand, these findings forcefully illustrate that it may be inappropriate to use industry-level instruments when analyzing firm-level processes, even if narrowly defined. These instruments are likely to capture internal responses to market forces alongside the firm-level processes under study, and therefore generate upward biased estimates. Notably, reliance on industry-level instruments is common practice within several strands of the current empirical literature. On a more constructive side, our results suggest that the size of the bias is related to the mobility patterns of workers in the economy, which implies that the magnitude of the bias could be assessed *a priori* by exploring the extent to which the instrument is correlated with workers' mobility patterns.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High-Wage Workers and High-Wage Firms," *Econometrica*, 67, 251–333.

Arellano, M., and S. Bond (1991): "Some Tests of Specifications for Panel Data: Monte Carlo Evidence and Application to Employment Equations," *Review of Economic Studies*, 58, 277–297.

Bagger, J., F. Fontaine, F. Postel-Vinay, and J.-M. Robin (2011): "Tenure, Experience, Human Capital and Wages: A Tractable Equilibrium Search Model of Wage Dynamics," Mimeo, Siences Po.

Basu, S., and J. Fernald (1995): "Are Apparent Productive Spillovers a Figment of Specification Errors?," *Journal of Monetary Economics*, 36, 165–188.

——— (2001): *Why Is Productivity Procyclical? Why Do We Care?* In C. Hulten, D. Edwin and M. Harper, ed., New Developments in Productivity Analysis, The University of Chicago Press, Chicago, IL.

Basu, S., J. Fernald, and M. Shapiro (2001): "Productivity Growth in the 1990s: Technology, Utilization, or Adjustment?," *Carnegie-Rochester Conference Series on Public Policy*, 55, 117–165.

Bell, B., and J. Van Reenen (2011): "Firm Performance and Wages: Evidence from Across the Corporate Hierarchy," CEP Discussion Papers 1088.

Bertrand, M., and S. Mullainathan (2001): "Are CEOs Rewarded for Luck? The Ones without Principles Are," *Quarterly Journal of Economics*, 116, 901–932.

Burnside, C., M. Eichenbaum, and S. Rebelo (1995): *Capital Utilization and Returns to Scale*. In B. Bernanke and J. Rotemberg, ed., NBER Macroeconomics Annual, MIT Press, Cambridge, MA.

Card, D., F. Devicienti, and A. Maida (2011): "Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data," IZA Discussion Paper 6086.

Carlsson, M. (2003): "Measures of Technology and the Short-Run Response to Technology Shocks," *Scandinavian Journal of Economics*, 105, 555–579.

Carlsson, M., and O. Nordström-Skans (2012): "Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost," *American Economic Review*, 102, 1571–1595.

Carlsson, M., and J. Smedsaas (2007): "Technology Shocks and the Labor Input Response: Evidence from Firm-Level Data," *Journal of Money, Credit and Banking*, 39, 1509–1520.

Carneiro, A., P. Guimarães, and P. Portugal (2012): "Real Wages and the Business Cycle: Accounting for Worker, Firm, and Job Title Heterogeneity," *American Economic Journal: Macroeconomics*, 4, 133–152.

Daniel, C., and C. Sofer (1998): "Bargaining, Compensating Wage Differentials, and Dualism of the Labor Market: Theory and Evidence for France," *Journal of Labor Economics*, 16, 546–575.

Diewert, E. (1976): "Exact and Superlative Index Numbers," *Journal of Econometrics*, 4, 115–146.

Doeringer, P., and M. Piore (1971): *Internal Labor Markets and Manpower Analysis*. Heath Lexington Books, Lexington, MA.

Eeckhout, J., and P. Kircher (2011): "Identifying Sorting - In Theory," *Review of Economic Studies*, 78, 872–906.

Eeckhout, J., and R. Pinheiro (2012): "Diverse Organizations and the Competition for Talent," SSRN Working Paper 1689753.

Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler (2004): "The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence From Colombia," *Journal of Development Economics*, 75, 333–371.

Faberman, R. J., and E. Nagypál (2008): "Quits, Worker Recruitment, and Firm Growth: Theory and Evidence," Federal Reserve Bank of Philadelphia Working Paper 08-13.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 98, 394–425.

Fuss, C., and L. Wintr (2009): "Rigid Wages and Flexible Employment? Firm-Level Evidence Based on Productivity for Belgium," ECB Working Paper 1021.

Greenaway, D., and R. Kneller (2007): "Firm Heterogeneity, Exporting and Foreign Direct Investment," *Economic Journal*, 117, F134–F161.

Guiso, L., F. Schivardi, and L. Pistaferri (2005): "Insurance within the Firm," *Journal of Political Economy*, 113, 1054–1087.

Haefke, C., M. Sonntag, and T. Van Rens (2008): "Wage Rigidity and Job Creation," IZA Discussion Papers 3714.

Hall, R. E. (2003): "Modern Theory of Unemployment Fluctuations: Empirics and Policy Applications," *American Economic Review Papers and Proceedings*, 93, 145–150.

Haltiwanger, J., J. Lane, and J. Spletzer (1999): "Productivity Differences Across Employers: The Roles of Employer Size, Age, and Human Capital," *American Economic Review*, 89, 94–98.

Hummels, D., R. Jørgensen, J. R. Munch, and C. Xiang (2011): "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," NBER Working Paper 17496.

Jorgenson, D., F. Gollop, and B. Fraumeni (1987): *Productivity and U.S. Economic Growth*. Harvard University Press, Cambridge, MA.

Kaas, L., and P. Kircher (2011): "Efficient Firm Dynamics in a Frictional Labor Market," IZA Discussion Paper 5452.

Kleibergen, F., and R. Paap (2006): "Generalized Reduced Rank Tests Using the Singular Value Decomposition," *Journal of Econometrics*, 127, 97–126.

Klette, T., and Z. Griliches (1996): "The Inconsistency of Common Scale Estimators when Output Prices are Unobserved and Endogenous," *Journal of Applied Econometrics*, 11, 343–361.

Manning, A. (2011): "Imperfect Competition in the Labor Market," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, chap. 11, pp. 973–1041. Elsevier, Amsterdam.

Marchetti, D., and F. Nucci (2005): "Price Stickiness and the Contractionary Effect of Technology Shocks," *European Economic Review*, 49, 1137–1163.

Marshak, J., and W. Andrews (1944): "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 50, 649–670.

Melander, O. (2009): "The Effect of Cash Flow on Investment: An Empirical Test of the Balance Sheet Channel," Sveriges Riksbank Working Paper Series No. 228.

Mendes, R., G. J. Van Den Berg, and M. Lindeboom (2010): "An Empirical Assessment of Assortative Matching in the Labor Market," *Labour Economics*, 17, 919–929.

Mortensen, D., and C. Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397–415.

Nordström Skans, O., P.-A. Edin, and B. Holmlund (2009): "Wage Dispersion Between and Within Plants: Sweden 1985-2000," in *Wage Structure, Raises and Mobility: International Comparisons of the Structure of Wages Within and Across Firms*, ed. by E. Lazear, and K. Shaw. University of Chicago Press, Chicago, IL.

Pissarides, C. (2009): "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?," *Econometrica*, 77, 1339–1369.

Postel-Vinay, F., and H. Turon (2010): "On-the-job Search, Productivity Shocks and the Individual Earnings Process," *International Economic Review*, 51, 599–629.

Rogerson, R. (2005): "Sectoral Shocks, Human Capital, and Displaced Workers," *Review of Economic Dynamics*, 8, 89–105.

Roodman, D. (2006): "How to Do xtabond2: An introduction to "Difference" and "System" GMM in Stata," Working Paper Number 103, Center for Global Development.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

Smith, E. (1999): "Search, Concave Production and Optimal Firm Size," *Review of Economic Dynamics*, 2, 456–471.

Stole, L., and J. Zwiebel (1996): "Intra-Firm Bargaining under Non-Binding Contracts," *Review of Economic Studies*, 63, 375–410.

Syverson, C. (2011): "What Determines Productivity?," *Journal of Economic Literature*, 49, 326–365.

Verhoogen, E. (2008): "Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics*, 123, 489–530.

Windmeijer, F. (2005): "A Finite Sample Correction for the Variance of Linear Efficient Two-step GMM Estimators," *Journal of Econometrics*, 126, 25–51.

Wooldridge, J. (2002): *Econometric Analysis of Cross Section and Panel Data.* MIT Press, Cambridge, MA.

Appendices

A Data Construction

The firm data set we use is primarily drawn from the Industry Statistics Survey (IS) and contains annual information for the years 1990-1996 on inputs and output for all Swedish manufacturing plants with 10 employees or more and a sample of smaller plants. The data is matched to RAMS, which adds individual wages and worker characteristics of each employee of the manufacturing plants included in the sample. In the baseline sample we focus on continuing plants that are also a firm.

When computing labor productivity, $LP_{jt}(= Y_{jt}/N_{jt}^{IS})$, labor input, N_{jt}^{IS}, is measured as the average number of employees during the year and is taken from the IS. To compute the input index, $\Delta\widetilde{x}_{jt}$, which is used to estimate the returns to scale and change in technology, real intermediate inputs, M_{jt}, are measured as the sum of costs for intermediate goods and services collected from the IS deflated by a three-digit (SNI92/NACE) producer price index collected by Statistics Sweden. Moreover, energy, V_{jt}, is measured as the plants' electricity consumption in MWh taken from the IS.

When computing the (overall) cost shares, we also need a measure of the firms' labor cost, which is defined as total labor cost including e.g. payroll taxes available in the IS. Also, to calculate the cost shares by education in expression (7) as well as the growth rate for respective category of labor input, we use the RAMS data (see discussion in section 3 above). Here we define LHE (less then high school education) as individuals with a one-digit ISCED 97 level code smaller than or equal to two, HE (high school education) as individuals with a one-digit ISCED 97 level code equal to three and TE (tertiary education) as individuals with a one-digit ISCED 97 level code larger than or equal to four.[51] Moreover, since Sweden experienced a boom-bust cycle in the late 80s and early 90s we do not use observations from firms experiencing large losses when calculating the two-digit (SNI92/NACE) cost shares. In the calculations we drop observations for firms where the (residual) capital share is below -10 percent of sales. This procedure gives rise to aggregate manufacturing cost shares that

[51] We exclude individuals with missing information on education from the calculations.

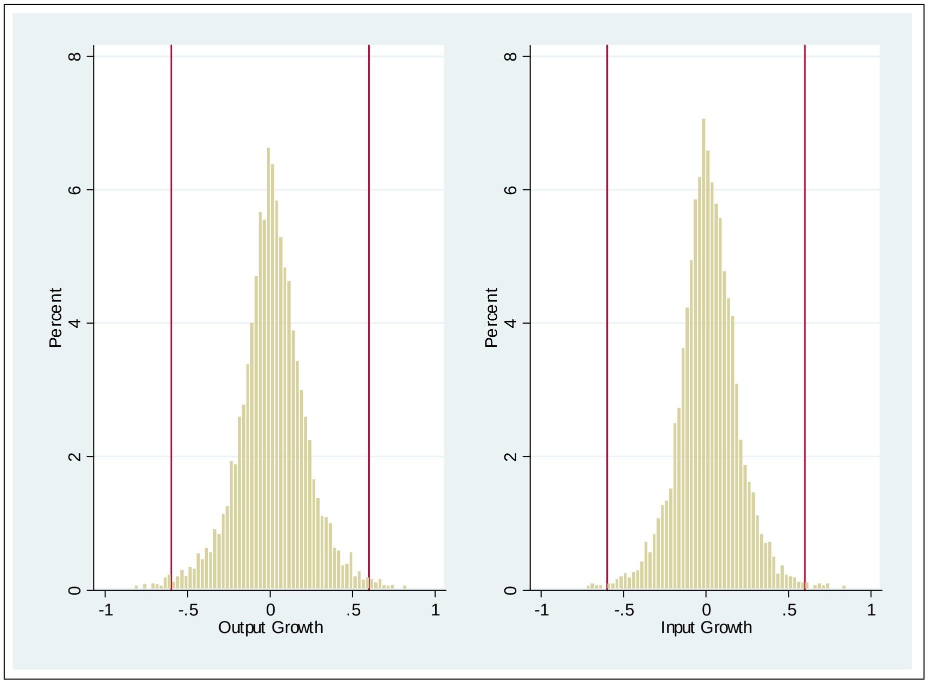

Figure 1: Distribution of output and input growth rates. Vertical lines indicate truncation limits.

are similar to those obtained using the data underlying Carlsson (2003).[52]

Although we have removed obviously erroneous observations, the firm data set still contains very large observations in Δy_{jt} and $\Delta\widetilde{x}_{jt}$. To avoid our returns to scale estimates being affected by firms subject to episodes of extreme conditions, these observations are removed (see below). In figure 1, the data distributions are plotted for the relevant variables for estimating returns to scale and technology change (truncated at ± 1 in log-difference space).

Since the main mass of the data seems to be well captured in the interval ± 0.6 for all variables, we limit the data set to contain firms with observations only within this interval. Note that e.g. $\Delta y_{jt} = 0.6$ corresponds to an annual increase of 82 percent in real output.

[52] The aggregate manufacturing shares in Carlsson (2003) (our sample) equals $C^M = 0.65$ $(0.66), C^N = 0.25$ $(0.20), C^K = 0.07$ (0.12) and $C^V = 0.03$ (0.03).

This procedure removes 160 firms from the sample, leaving us with 1,138 firms.[53,54] In order to decompose the technology series into a sectoral and an idiosyncratic part we need to drop two additional firms since they are the only firms in the sample pertaining to a particular sectoral agreement. This then leaves 1,136 firms in the data set we then use to estimate the specification (8).

After merging the final firm-level data with the employee data in RAMS we arrive at 474,528 employee observations across 106,815 individuals. Removing observations where education information is missing we have 472,555 observation across 106,050 individuals left. This data set covers 11 percent of the full-time employees in the manufacturing sector.

Finally, unemployment and vacancy data on the local labor market level is collected from the National Labor Market Board (AMS). The data contains information on the number of registered vacancies and the number of individuals registered as openly unemployed at an unemployment office in November. We use the (1993) definition of homogenous local labor markets constructed by Statistics Sweden using commuting patterns, which divides Sweden into 109 areas.

B The problems with using VA to derive TFP

A standard approach to derive TFP is to rely on value added as a measure of production. As we illustrate in this appendix, the use of value added in combination with deviations from non-constant returns will result in a measure of TFP that is not independent from the use of intermediate inputs and factor input growth.

[53] We do not remove observations with large movements in labor productivity since this variable will be instrumented in the econometric procedure.

[54] Robustness exercises show that the baseline results, presented in the main text, are not sensitive to this trimming of the data.

Using the implicit definition of the divisia index of value added, we arrive at[55]

$$\Delta va_{jt} = \psi_j \Delta x_{jt}^{VA} + \frac{\Delta tfp_{jt}}{1 - C_j^M - C_j^V} + \frac{(\psi_j - 1)}{1 - C_j^M - C_j^V} \left(C_j^V \Delta v_{jt} + C_j^M \Delta m_{jt} \right), \tag{B1}$$

where $\Delta x_{jt}^{VA} = \left[\left(C_j^K / (C_j^K + C_j^N) \right) \Delta k_{jt} + \left(C_j^N / (C_j^K + C_j^N) \right) \Delta n_{jt} \right]$ is the weighted growth rate of primary factors and Δva_{jt} is the growth rate of real value-added.[56] As can be seen in equation (B1), real value added will not only depend on primary factors, but also on materials and energy growth, unless there are constant returns. To see why, one can think of real value added as a partial TFP measure subtracting the productive contribution of materials and energy from real gross output under the assumption of perfect competition and constant returns. Hence, when constructing a value-added Solow Residual, Δtfp_{jt}^{VA},

$$\Delta tfp_{jt}^{VA} = (\psi_j - 1) \Delta x_{jt}^{VA} + \frac{\Delta tfp_{jt}}{1 - C_j^M - C_j^V} + \frac{(\psi_j - 1)}{1 - C_j^M - C_j^V} \left(C_j^V \Delta v_{jt} + C_j^M \Delta m_{jt} \right), \tag{B2}$$

by subtracting Δx_{jt}^{VA} from Δva_{jt}, the resulting measure will also depend on materials and energy use, unless there is constant returns. We also see that there will be an effect on the value-added Solow residual working via primary inputs growth through the implied $(\psi_j - 1)\Delta x_{jt}^{VA}$ term in Δtfp_{jt}^{VA}. Note, though, that this particular effect (but not the effect working through intermediate materials and energy growth) would vanish if we allowed for non-constant returns when computing TFP from value-added data.

Comparing expressions (B1) and (B2), it is easy to see that unless there are constant returns to scale, or a very special covariance structure across the different production factors, there will be a component in the correlation between (B1) and (B2) that is driven by input factor growth and not technology growth. To the extent that e.g. demand shocks are

[55] See Basu and Fernald (1995) for a full derivation. Note that Basu and Fernald (1995) does not separate between intermediates and energy as is done here.

[56] For clarity, we do not substitute Δk_{jt} with Δv_{jt} here. This is, however, done in the empirical work.

correlated with factor input growth, this type of shock will affect both the instrument as well as the instrumented variable, giving rise to a bias in the coefficient of labor productivity on wages.

C Constructing a Measure of the Capital Stock

We calculate the capital stock using investment data and book values (for the starting values). When using a measure of the capital stock, the input index is defined as $\Delta\widetilde{x}_{jt}^{C} = C_j^K \Delta k_{jt} + C_j^V \Delta v_{jt} + C_j^N \Delta n_{jt} + C_j^M \Delta m_{jt}$. The capital stock, K_{jt}, is computed using a variation of the perpetual inventory method which utilizes all the information we have available in the data.

We calculate the capital stock in two steps. In the first step we calculate the forward recursion

$$K_{jt} = \max\left\{(1-\delta_s)K_{jt-1} + I_{jt}, BookValue_{jt}\right\}, \tag{C3}$$

where δ_s is a sector-specific depreciation rate (two-digit SNI92/NACE) and is computed as an asset share-weighted average between the depreciation rates of machinery and buildings (collected from Melander, 2009, table 2), I_{jt} is real net investments in fixed tangible assets (deflated using a two-digit (SNI92/NACE) sector-specific investment deflator collected from Statistics Sweden) and $BookValue_{jt}$ is the real book value of fixed tangible assets (computed using the same deflator as for investment) and

$$K_{j0} = \begin{cases} 0 \text{ if } BookValue_{j0} \text{ is missing,} \\ BookValue_{jt} \text{ otherwise.} \end{cases} \tag{C4}$$

Since the firm has an incentive to keep the book values low for tax reasons, we use the book values as a lower bound of the capital stock. In a second step, we calculate the backward recursion

$$K_{jt-1} = \frac{K_{jt} - I_{jt}}{(1-\delta_s)}, \tag{C5}$$

where the ending point of the first recursion, K_{jT}, is used as the starting point for the backward recursion. This is done in order to maximize the quality of the capital-stock series given that we do not have a very reliable starting point and the time-series dimension is short. Taking account for missing data when calculating the capital stock, we can project the technology levels for 944 firms using $\Delta\widetilde{x}_{jt}^{C}$ instead of $\Delta\widetilde{x}_{jt}$.

Tables

Table 1: Returns to Scale

Industry	RTS
Durables	0.986
	(0.194)
Non-Durables	0.882
	(0.224)
Observations	5,680
Firms	1,136
AR(2)	[0.210]
AR(3)	[0.886]
Hansen	[0.296]

Note: Sample 1991-1996 with 1,136 firms. Difference GMM second-step estimates with robust Windmeijer (2005) finite-sample corrected standard errors in parenthesis. See main text for instruments used. Regression includes time dummies and firm fixed effects. P-values for diagnostic tests inside brackets.

Table 2: Summary Statistics

	All		Male	
	Mean	S.D.	Mean	S.D.
Wages:				
w_{ijt}	9.615	0.313	9.662	0.308
w_{ijt} (Within Match)	-	0.146	-	0.146
Productivity:				
lp_{jt}	6.835	0.667	6.865	0.677
lp_{jt} (Within Match)	-	0.155	-	0.157
lp^S_{st}	-	0.049	-	0.050
lp^S_{st} (Within Match)	-	0.036	-	0.037
lp^I_{jt}	-	0.669	-	0.680
lp^I_{jt} (Within Match)	-	0.130	-	0.131
TFP				
tfp^S_{st}(Within Match)	-	0.022	-	0.022
tfp^I_{jt}(Within Match)	-	0.092	-	0.092
Worker characteristics:				
Age_{ijt}	39.8	11.8	39.7	11.9
Share of Male	0.794		1	
Share of HE	0.511		0.519	
Share of TE	0.122		0.127	
Share of Non-Immigrants	0.895		0.904	
Firm-Size	212.6		213.8	
Observations	472,555		374,975	

Note: The "Within match" rows show the dispersion within a combination of person and firm. All statistics are weighted according to the number of employees.

Table 3: The Impact of Productivity on Individual Wages

	(1)	(2)	(3)	(4)	(5)
Estimation Method:	OLS	IV	IV	IV	OLS
Regressor	Prod	Prod	Prod	Prod	TFP
lp^S_{st}	0.027	0.123**	0.149**	0.149**	
	(0.022)	(0.036)	(0.037)	(0.038)	
lp^I_{jt}	0.033**	0.032**	0.050**	0.051**	
	(0.008)	(0.011)	(0.010)	(0.010)	
tfp^S_{st}					0.124**
					(0.031)
tfp^I_{jt}					0.042**
					(0.008)
Firm FE	Yes	Yes	Yes	–	–
Worker FE	No	No	Yes	–	–
Worker Characteristics	No	No	Yes	Yes	Yes
Worker by Firm FE	No	No	No	Yes	Yes
Observations	472,555	472,555	472,555	472,555	472,555
Firms	1,136	1,136	1,136	1,136	1,136
Worker by Firm Matches	-	-	-	107,086	107,086
Kleibergen-Paap rk LM statistic	-	46.83	NA	44.52	-
P-value	-	0	NA	0	-

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Individual controls include age, age squared and age cubed (columns 3-5). K-P denotes the Kleibergen-Paap (2006) rk LM statistic for testing the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.

Table 4: First-Stage Regressions

	(1)	(2)	(3)	(4)
	OLS	OLS	OLS	OLS
Dependent Variable:	lp_{st}^S	lp_{jt}^I	lp_{st}^S	lp_{jt}^I
tfp_{st}^S	0.854**	0.000	0.842**	-0.020
	(0.087)	(0.148)	(0.088)	(0.141)
tfp_{jt}^I	-0.000	0.846**	0.000	0.838**
	(0.008)	(0.031)	(0.009)	(0.031)
Firm FE	Yes	Yes	–	–
Worker FE	No	No	–	–
Worker Characteristics	No	No	Yes	Yes
Worker by Firm FE	No	No	Yes	Yes
Observations	472,555	472,555	472,555	472,555
Firms	1,136	1,136	1,136	1,136
Worker by Firm Matches	-	-	107,086	107,086
F-Stat($tfp_{st}^S = tfp_{jt}^I = 0$)	49.00**	382.8**	46.82**	378.8**

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. Regressions also include time effects and labor market tightness. Individual controls (columns 3-4) include age, age squared and age cubed. F denotes the F statistic for excluded instruments.

Table 5: Market Forces Vs. Bargaining Power

	(1)	(2)	(3)	(4)	(5)
Estimation Method:	IV	IV	IV	IV	IV
lp_{st}^{S}	0.149**	0.029			0.141**
	(0.038)	(0.128)			(0.037)
$lp_{st}^{S,Weighted}$		0.199	0.244**		
		(0.214)	(0.061)		
$lp_{st}^{S,NACE}$				0.203**	
				(0.045)	
$lp_{st}^{S}\hat{\phi}_{ijt}$					1.807**
					(0.693)
lp_{jt}^{I}	0.051**	0.051**	0.051**		0.050**
	(0.010)	(0.010)	(0.010)		(0.010)
$lp_{jt}^{I,NACE}$				0.043**	
				(0.011)	
$lp_{jt}^{I}\hat{\phi}_{ijt}$					-0.016
					(0.153)
Worker Characteristics	Yes	Yes	Yes	Yes	Yes
Worker by Firm FE	Yes	Yes	Yes	Yes	Yes
Observations	472,555	472,555	472,555	472,555	472,555
Firms	1,136	1,136	1,136	1,136	1,136
Worker by Firm Matches	107,086	107,086	107,086	107,086	107,086
Kleibergen-Paap rk LM statistic	44.52	41.74	40.29	71.72	186.62
P-value	0	0	0	0	0

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Individual controls include age, age squared and age cubed. The regression in column (5) also includes the main effect from the predicted return probability ($\hat{\phi}$). The Kleibergen-Paap (2006) rk LM statistic is a test statistic for the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.

Table 6: Complementarities and Sorting

	(1)	(2)	(3)
Estimation Method:	IV	IV	IV
Dependent Variable:	w_{ijt}	w_{ijt}	$\hat{w}_{ijt}$
lp^S_{st}	0.170**	0.147**	-0.009
	(0.039)	(0.038)	(0.009)
$lp^S_{st} * HE_{it}$	-0.053		
	(0.030)		
$lp^S_{st} * TE_{it}$	0.060		
	(0.076)		
$lp^S_{st} * \hat{w}_{ij}$		0.106	
		(0.150)	
lp^I_{jt}	0.047**	0.050**	-0.004
	(0.011)	(0.010)	(0.002)
$lp^I_{jt} * HE_{it}$	0.005		
	(0.009)		
$lp^I_{jt} * TE_{it}$	0.010		
	(0.025)		
$lp^I_{jt} * \hat{w}_{ij}$		0.046	
		(0.046)	
Firm FE	–	–	Yes
Worker FE	–	–	No
Worker Characteristics	Yes	Yes	Yes
Worker by Firm FE	Yes	Yes	No
Observations	472,555	472,555	472,555
Firms	1,136	1,136	1,136
Worker by Firm Matches	107,086	107,086	-
Kleibergen-Paap rk LM statistic	44.47	55.08	46.85
P-value	0	0	0

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Individual controls include age, age squared and age cubed. Interactions in column 2, and the dependent variable in colum 3, are based on predictions from human capital regressions (see main text for details). The regression in column 2 also includes the main effect of predicted human capital. The Kleibergen-Paap (2006) rk LM statistic is a test statistic for the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.

Table 7: Robustness Variations

Variation:	(1) Baseline	(2) Unbalanced	(3) Time Varying C_J	(4) RTS-0.1	(5) CRS	(6) RTS+0.1
lp^S_{st}	0.149**	0.097**	0.170**	0.125**	0.168**	0.184**
	(0.038)	(0.039)	(0.046)	(0.031)	(0.045)	(0.055)
lp^I_{jt}	0.051**	0.018**	0.044**	0.061**	0.042**	0.038**
	(0.010)	(0.008)	(0.011)	(0.009)	(0.011)	(0.011)
Worker Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Worker by Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	472,555	2,048,555	472,682	472,555	472,555	472,555
Firms/Plants	1,136	2,444	1,137	1,136	1,136	1,136
Worker by Firm Matches	107,086	465,760	107,140	107,086	107,086	107,086
Kleibergen-Paap rk LM statistic	44.52	23.70	39.87	55.43	42.68	32.56
P-value	0	0	0	0	0	0

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Worker characteristics include age, age squared and age cubed. The Kleibergen-Paap (2006) rk LM statistic is a test statistic for the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.

Table 8: The Impact of Different Deflators and Output Measures

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	IV
Prices (Productivity)	Firm-level	3-digit PPI	Firm-level	Firm-level	Firm-level
Prices (TFP)	Firm-level	3-digit PPI	Firm-level	Firm-level	Firm-level
Output Measure	Gross Output	Gross Output	Value-added	Value-added	Value-added
Sector	All	All	All	Non-Durables	Durables
lp^S_{st}	0.149**	0.112**	0.040*	0.035	0.086**
	(0.038)	(0.035)	(0.017)	(0.019)	(0.020)
lp^I_{jt}	0.051**	0.092**	0.023**	0.022**	0.027**
	(0.010)	(0.016)	(0.005)	(0.006)	(0.010)
Worker Characteristics	Yes	Yes	Yes	Yes	Yes
Worker by Firm FE	Yes	Yes	Yes	Yes	Yes
Observations	472,555	472,555	463,539	281,100	182,439
Number of Firms	1,136	1,136	1,125	712	413
Worker by Firm Matches	107,086	107,086	105,784	63,230	42,554
Kleibergen-Paap	44.52	121.61	77.56	45.61	35.42
P-value	0	0	0	0	0

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Worker characteristics include age, age squared and age cubed. K-P denotes the Kleibergen-Paap (2006) rk LM statistic for testing the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.

Table 9: OLS and IV with Decreasing and Constant Returns to Scale

	(1)	(2)	(2)	(4)
Estimation Method:	OLS	IV	OLS	IV
Returns to scale:	Decreasing returns (Non-Durables)		Constant returns (Durables)	
lp_{st}^{S}	0.027	0.140**	0.177**	0.169**
	(0.028)	(0.042)	(0.038)	(0.055)
lp_{jt}^{I}	0.033**	0.052**	0.050**	0.049**
	(0.007)	(0.012)	(0.011)	(0.017)
Worker Characteristics	Yes	Yes	Yes	Yes
Worker by Firm FE	Yes	Yes	Yes	Yes
Observations	286,907	286,907	185,648	185,648
Firms	720	720	416	416
Worker by Firm Matches	64,084	64,084	43,002	43,002
Kleibergen-Paap rk LM statistic	NA	42.89	NA	63.50
P-value	NA	0	NA	0

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Worker characteristics include age, age squared and age cubed. K-P denotes the Kleibergen-Paap (2006) rk LM statistic for testing the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.

Table 10: Dynamic Effects

	(1)	(2)	(3)
Estimation method:	IV	IV	IV
lp^S_{st}	0.149**	0.052	0.226
	(0.038)	(0.095)	(0.130)
lp^S_{st-1}		0.104	0.198
		(0.088)	(0.211)
lp^S_{st-2}			-0.122
			(0.147)
lp^I_{jt}	0.051**	0.035**	0.033*
	(0.010)	(0.012)	(0.014)
lp^I_{jt-1}		0.034**	0.032**
		(0.013)	(0.012)
lp^I_{jt-2}			0.026*
			(0.010)
Total Sector Effect	0.149**	0.157**	0.303**
	(0.038)	(0.045)	(0.110)
Total Idiosyncratic Effect	0.051**	0.068**	0.091**
	(0.010)	(0.015)	(0.020)
Worker Characteristics	Yes	Yes	Yes
Worker by Firm FE	Yes	Yes	Yes
Observations	472,555	402,058	335,291
Firms	1,136	1,136	1,136
Worker by Firm Matches	107,086	99,473	93,316
Kleibergen-Paap rk LM statistic	44.52	25.79	4.78
P-value	0	0	0.029

Note: * (**) denotes significance at the 5 (1) percent level. Standard errors clustered on firms reported inside parentheses. All specifications include time effects and labor market tightness. Worker characteristics include age, age squared and age cubed. K-P denotes the Kleibergen-Paap (2006) rk LM statistic for testing the null hypothesis that the equation is underidentified. P-value denotes the associated p-value for the test.