IFAU
Institute for Evaluation of Labour
Market and Education Policy

Firm-level shocks and labor adjustments

Mikael Carlsson
Julián Messina
Oskar Nordström Skans

WORKING PAPER 2014:28

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala. IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

IFAU also provides funding for research projects within its areas of interest. The deadline for applications is October 1 each year. Since the researchers at IFAU are mainly economists, researchers from other disciplines are encouraged to apply for funding.

IFAU is run by a Director-General. The institute has a scientific council, consisting of a chairman, the Director-General and five other members. Among other things, the scientific council proposes a decision for the allocation of research grants. A reference group including representatives for employer organizations and trade unions, as well as the ministries and authorities concerned is also connected to the institute.

Postal address: P.O. Box 513, 751 20 Uppsala
Visiting address: Kyrkogårdsgatan 6, Uppsala
Phone: +46 18 471 70 70
Fax: +46 18 471 70 71
ifau@ifau.uu.se
www.ifau.se

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

ISSN 1651-1166

Firm-level shocks and labor adjustments*

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

December 2014

Abstract

We analyze how firms adjust their labor in response to idiosyncratic shifts in their production function and demand curves using a unique data-set of Swedish manufacturing firms. We show that permanent shocks to firm-level demand is a main driving force behind both job and worker reallocation. In contrast, shocks to physical productivity and temporary demand shocks have a very limited impact on firm-level employment despite being important determinants of other firm-level fundamentals. We also present evidence suggesting that the adjustment to permanent demand shocks is fairly unconstrained. Most notably, firms primarily downsize through increased separations of both short- and long-tenured workers even when they could have adjusted their employment through reduced hires.

Keywords: Technology, Demand, Job Creation, Rigidities, Worker Flows
JEL classifications: J23, J63, O33, C33

*We would like to thank Marcela Eslava, Richard Freeman, Nils Gottfries, Erica Groshen, John Haltiwanger, Francis Kramarz, Per Krusell, Claudio Michelacci, Frédèric Warzynski and Nikolaj Angelov for helpful comments. In addition, we thank seminar participants at ELE, 2012; IAE-MOVE, 2012; Greater Stockholm Macro Group, 2013; CAP, 2013; EALE 2014; SOLE 2014; U. Carlos III; Bank of Spain; U. of Warwick; World Bank; IFAU; UCLS; Sveriges Riksbank; Linnaeus U. and Gothenburg U. We are also grateful to Karl Harmenberg, Mats Levander, Matias Quiroz, and Juan Pablo Uribe for excellent research assistance. The data are confidential, but the authors' access is not exclusive. Carlsson and Skans are grateful for financial support from the Ragnar Söderberg Foundation. 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
‡World Bank and IZA. e-mail: jmessina@worldbank.org
§Uppsala University, IFAU, UCLS and IZA. e-mail: oskar.nordstrom_skans@nek.uu.se

Contents

1 Introduction

About one in every five jobs are either created or destroyed every year (Davis, Faberman and Haltiwanger, 2006). The bulk of this firm-level labor adjustment is truly idiosyncratic as firms operating in the same sector and area shrink and grow side-by-side. Hence, jobs are rapidly created and destroyed, even in sectors with stable net employment. Following the seminal work of Davis, Haltiwanger and Schuh (1996), the importance and magnitude of these labor flows has been documented for a large number of countries.[1] However, while the empirical regularities of job and worker flows have been abundantly documented, little is known about the structural determinants of these flows.[2]

This paper presents novel evidence on how firms adjust employment through hires and separations when their positions in the performance distribution change. We derive empirical measures of permanent idiosyncratic firm-level shocks to analyze how firms adjust their labor when their relative performance has been altered beyond simple transitory fluctuations.[3] By analyzing the direct impact of idiosyncratic shocks, we concentrate on employment adjustments in a stable market environment, effectively abstracting from feedback effects through changes in market wages or aggregate unemployment. Using the terminology of Foster, Haltiwanger, and Syverson (2008) we define technology shocks as shifts in the firm-level physical production function (i.e., the ability to produce at a given level of inputs) and demand shocks as shifts in the firm-level demand curve (i.e., the ability to sell at a given price).[4] We present direct evidence on how the magnitudes and signs of permanent changes in these key determinants of firm performance affect firms' labor adjustments along different margins. Our empirical analysis relies on a unique data base that links measures of firm-level input, output, and prices to individual worker-flow data for Swedish manufacturing firms.

The analysis adds to a vibrant empirical literature, surveyed in Syverson (2011), that documents the distinct impacts of firm-level technology and demand shocks on productivity

[1] See Davis, Faberman, and Haltiwanger (2012) for an overview. For evidence from Sweden, which is the empirical subject of this paper, see Andersson (2003).

[2] A small macro-oriented literature aims to identify the employment responses to technology-driven changes in firm-level productivity, see e.g., Carlsson and Smedsaas (2007) and Marchetti and Nucci (2005). The macro literature also contains a number of related studies, e.g., Galí (1999) and Michelacci and Lopez-Salido (2007), the latter of which distinguished between neutral technology shocks and investment-specific technology shocks and derived the consequences for job reallocation.

[3] In an extension, we also provide an analysis of firms' responses to transitory shocks.

[4] Importantly, these shocks are defined according to their effects on firm-level optimization, not according to their origins.

and other firm-level outcomes. Most notably, Foster et al. (2008) shows that firm closures are driven primarily by changes in idiosyncratic demand and only to a lesser extent by changes in idiosyncratic physical productivity. Recent evidence in Foster, Haltiwanger, and Syverson (2012) suggests that the growth of young firms in the US is due to a shrinking product-demand gap relative to incumbents. Pozzi and Schivardi (2012) shows similar results using Italian data. In addition, Carlsson, Messina, and Nordström-Skans (2014) shows that firm-level technology shocks affect workers' wages, using Swedish data.

This paper is, however, the first to detail how firm-level technology and demand shocks affect firms' labor adjustments through hires and separations, a question that speaks to a huge body of theoretical research regarding the relationship between firm-level revenue productivity and labor adjustments (Bentolila and Bertola, 1990; Davis and Haltiwanger, 1992; Hopenhayn and Rogerson, 1993; Mortensen and Pissarides, 1994; and more recently Cahuc, Postel-Vinay, and Robin, 2006; and Lise, Meghir, and Robin, 2013). The focus of this paper is to disentangle the separate roles in labor adjustments of two fundamental drivers of firm-level revenue productivity fluctuations: demand and technology shocks.

Our analysis departs from a simple model of monopolistic competition that motivates a set of restrictions on the long-run relationship between firms' fundamentals and shocks. In the spirit of Franco and Philippon (2007), we then impose these long-run restrictions in a structural vector autoregression (SVAR) setting to filter out our empirical measures of permanent idiosyncratic demand and technology shocks. This allows us to derive the shocks without imposing any restrictions on the firms' short-run behavior.

The most important restriction we rely on is the notion that the physical gross Solow residual is independent of all shocks except the technology shock in the long run. To translate this prediction to the data, we benefit from a firm-specific price index. Using a strategy similar to Eslava et al. (2004) and Smeets and Warzynski (2013), we deflate the (nominal) firm-level output series with firm-level price indices to derive measures of firm-level real output volumes. Importantly, the fact that we filter out the technology shocks using long-run restrictions implies that other shocks, or changes in factor utilization or inventories, are allowed to have a transitory impact on the physical Solow residual without affecting the measured technology shocks.

Our model also allows us to derive sufficient restrictions to identify permanent demand shocks, again without imposing any restrictions on the nature of short-run shocks or dynamics. Since we use data from a small, open economy, our system also explicitly allows for permanent sector-level shocks to factor prices.[5]

When implementing the SVAR, we use a strategy that differs from standard time-series applications such as Blanchard and Quah (1989) and Franco and Philippon (2007). Because we depart from a broad cross-sectional panel of firms, we estimate the reduced-form equations using dynamic panel data methods building on Arellano and Bond (1991). This allows us to estimate both the parameters and the covariance matrix of the error terms with considerable precision, thereby avoiding standard macro-data concerns regarding the practical implementation of SVARs.

Before moving to the main analysis we corroborate the interpretation of the idiosyncratic firm-level technology and demand shocks by showing that prices and output respond as theorized: Idiosyncratic output increases in response to positive shocks to both demand and technology. But firm-level prices decrease only in response to firm-level technology shocks, whereas they remain largely unaffected when demand shifts.

Turning to our key research questions, we start by showing that, despite being crucial for both firm-level prices and output, firm-level technology shocks have a relatively limited effect on labor inputs. In contrast, product demand is a key driving force behind firm-level labor adjustments. An idiosyncratic demand shock of 1 standard deviation increases employment by 6 percentage points. The employment change is fast and appears to be relatively unconstrained. Most of this adjustment takes place within a year. These results are robust to a wide range of variations in measures and specifications.

The strong impact of product demand changes on employment offers a natural instrument to analyze the dynamics of hirings and separations. Our analysis proceeds in the spirit of Abowd, Corbel, and Kramarz (1999) and Davis, Faberman, and Haltiwanger (2012), which decompose positive and negative changes in net employment into job and worker flows in France and the United States, respectively. Interestingly, these studies show that French firms reduce employment primarily by reducing hirings, whereas U.S. firms behave much more symmetrically by adjusting hirings and separations in similar proportions. In contrast to these decomposition exercises, we analyze hirings and separations in response to net employment changes pushed by a well-identified shock: permanent shifts in firms' product demand schedules. This allows us to abstract from other shifts in labor demand, e.g., worker attrition, that can alter hiring and separation policies.

Our findings are, however, well in line with the descriptive evidence from the United

[5] In addition, the system allows for a (transitory) residual shock component to soak up any remaining short-run dynamics, including mean-reverting shocks to purely idiosyncratic factor prices.

States. The employment adjustment induced by permanent demand shocks is both rapid and symmetric across hirings and separations. On average, Swedish manufacturing firms adjust employment almost as much through changes in the separation rate as through changes in the hiring rate. This result suggests that both margins should be treated as endogenous when modelling labor adjustments at the firm level. Although we find that the average firm reduces hires in response to negative shocks, it also continues to recruit workers even when forced to reduce employment substantially. Thus, the firms are far from exploiting the full potential of downsizing through reduced hirings. We also show that, in the face of negative shocks, only half of the adjustment through separations is related to short-tenured workers.

Overall, the speed of adjustment, the symmetry between hires and separations as adjustment margins, the rapid separation of long-tenured workers, and the continued recruitment of workers in the face of negative shocks jointly suggest that Swedish firms facing permanent idiosyncratic demand shocks do adjust their labor flexibly. This may appear surprising considering the slightly above-average levels of employment protection in Sweden relative to OECD countries in general (OECD, 2014), which should impose non-trivial employment adjustment costs.

To investigate this further, we extend the model to include the labor effects of transitory demand shocks. The firms' responses to transitory demand shocks is heavily muted, suggesting that they adjust employment only in response to long-lasting shocks. The results are thus consistent with a framework in which firms hoard labor when hit by temporary disturbances, but release workers rapidly when hit by a permanent reduction in their product demand. This implies that both the origin and the time-series properties of the shocks are crucial for the degree of labor adjustments.[6]

The paper is organized as follows: Section 2 outlines a simple model that motivates the long-run restrictions needed to extract our permanent demand and technology shocks. Section 3 introduces the main characteristics of the firm-level data used in the analysis and discusses the empirical implementation of the SVAR and the validation of the shocks. Section 4 reports the main results, distinguishing employment, hiring, and separation margins in response to technology and demand shocks. Section 5 presents extensions of the basic model, including a discussion of (i) the role of worker heterogeneity, and (ii) how temporary demand shocks affect employment adjustment. Finally, section 6 summarizes

[6] See also Guiso, Schivardi, and Pistaferri (2005), which shows that the time-series properties of shocks to firm-level value added have a crucial impact on workers' wages.

our results and findings.

2 Model and empirical strategy

2.1 Shocks to idiosyncratic production functions and demand curves

The purpose of the paper is to document how firms adjust labor in response to "structural" shocks—processes that involve changes in key elements of the firms' profit functions. We focus on two such processes: The first process captures shifts in the firm-specific physical production function; we label these shifts *technology shocks*. The second process captures shifts in the firm-specific demand curve; we label these shifts *demand shocks*.

The key distinction between a technology shock and a demand shock lies in how the shock affects the producing firm, not in the origin of the shock. We therefore refrain from modeling the origins. This approach, which is consistent with the existing (micro) literature (such as Foster et al., 2008, and Syverson, 2011), implies that we do not distinguish between shifts in the firm-specific demand curve that arise from changing preferences among final consumers, those that arise from increased demand among downstream firms, and those that arise from quality changes that increase product demand at a given price.[7]

To identify firm-level structural shocks, we need to make assumptions about the technology and market conditions faced by the firm. Our setup follows Eslava et al. (2004), Foster et al. (2008, 2012) closely, by using a first-order approximation of both production technologies and product market demand and by modeling the key technology and demand shocks as neutral shifters of the production function and the demand curve respectively. Thus, the firm-level production function is approximated by:

$$Y_{jt} = A_{jt} N_{jt}^{\alpha} K_{jt}^{\beta} M_{jt}^{1-\alpha-\beta} \text{ and } \alpha, \beta \in (0,1), \tag{1}$$

where physical gross output Y_{jt} in firm j at time t is produced using technology indexed by A_{jt} and combining labor input N_{jt}, capital input K_{jt}, and intermediate production factors (including energy) M_{jt}. Importantly, our data allow us to account for idiosyncratic

[7] Note that the firm-level price index we use is based on unit prices for very detailed product codes (8/9-digit Harmonized System/Combined Nomenclature codes), which limits the scope for quality changes to be the key component in our demand shock. However, it is straightforward to show that if we added a quality shock to the system developed below (through a wedge between the measured firm-level price, based on unit values, and the quality-adjusted price), it would enter the system symmetrically to the demand shock.

price differences across firms, so that our measure of technology (the Solow residual, A_{jt}) refers to *physical* total factor productivity (TFPQ), rather than to *revenue* total factor productivity (TFPR) in the terminology of Foster et al. (2008). Equation 1 presupposes a constant-returns technology, which is an assumption we maintain in our main specification, but we also present robustness exercises where we relax this assumption.

The firm-level demand curve is approximated by a constant-elastic function according to

$$Y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\sigma} Y_t \Omega_{jt} \text{ and } \sigma > 1, \tag{2}$$

where P_{jt}/P_t is the firm's relative price, Y_t denotes aggregate market demand, and Ω_{jt} is a firm-specific demand shifter. σ denotes the elasticity of substitution between different competing goods and hence captures the demand elasticity for each firm in the economy.

Following Guiso et al. (2008) and Franco and Philippon (2007) we model the key shocks as permanent shifters. More precisely, we specify the evolution of the demand and technology shifters as in Franco and Philippon (2007):

$$A_{jt} = A_{jt-1} e^{\mu_j^a + \Phi_j^a(L)\eta_{jt}^a}, \tag{3}$$

$$\Omega_{jt} = \Omega_{jt-1} e^{\mu_j^\omega + \Phi_j^\omega(L)\eta_{jt}^\omega}, \tag{4}$$

where μ_j^a and μ_j^ω are constant drifts, and $\Phi_j^a(L)$ and $\Phi_j^\omega(L)$ are polynomials in the lag operator, L. The white-noise idiosyncratic technology and demand shocks are denoted by η_{jt}^a and η_{jt}^ω. The assumed functional form implies that the shocks' lag polynomials are linearly related to the log differences of A_{jt} and Ω_{jt}, respectively.[8] As is evident from the formulation, our focus is on permanent shocks, but in a variation of the model we also explicitly analyze the role of transitory disturbances (see section 5.2).

Our model also allows for sectoral shocks to factor prices other than labor. This is useful in the Swedish setting of a small open economy where factor prices are likely to vary across sectors and time (due to exchange-rate volatility, for example). To simplify the notation, we next define a price index (consistent with cost minimization) for input factors other than labor in sector s, $P_{st}^F = \left(P_{st}^K/\beta\right)^\beta \left(P_{st}^M/(1-\alpha-\beta)\right)^{1-\alpha-\beta}$, where P_{st}^K is the capital price and P_{st}^M is the price of intermediate materials in sector s at time t. As for technology and demand, P_{st}^F evolves according to $P_{st}^F = P_{st-1}^F e^{\mu_s^f + \Phi_s^f(L)\eta_{st}^f}$, where μ_s^f

[8] This, in turn, provides a convenient moving average (MA) representation of the VAR specified below (see appendix C for details).

is a constant drift, $\Phi_j^f(L)$ is a polynomial in the lag operator, L, and η_{st}^f is a white-noise, sector-specific factor-price shock.

The specified shocks (together with aggregate conditions) are taken as state variables during firm-level wage determination. In addition, we assume cost minimization and that the firms have the right-to-manage so that factor choices are made taking wages as given.

2.2 Long-run restrictions

We rely on the stylized model presented above to derive a set of long-run restrictions that allow us to filter out the structural shocks of interest (η^a and η^ω). Table 1 summarizes the set of equations that motivate our restrictions, and appendix A presents details and derivations. The second column of the table denotes variables that can all be constructed from our firm-level data as long as we have an estimate of the demand elasticity σ (as detailed in the next section).

The third column summarizes the three key predictions that we rely on for identification:

1. The measured physical Solow residual ($TFPQ$ in the terminology of Foster et al. 2008) is equal to A and hence independent of both demand (Ω) and factor prices (P^F).
2. The "wage-neutral" unit labor costs ($WNULC$), as defined in the second row, is a function of both A and P^F.
3. The "wage-neutral" demand (WND), as defined in the third row, is a function of A, Ω, and P^F.

We use the modifier "wage-neutral" to highlight that the measures are defined to neutralize the impact of potential wage shocks.

These three predictions motivate the recursive sequence of restrictions highlighted in the fourth column: The Solow residual is independent of the innovations η^ω and η^f, and $WNULC$ is independent of η^ω. If invoked in the long run, these restrictions are sufficient to identify a VAR model in these variables using standard structural VAR (SVAR) techniques. In practice, we will also include a fourth residual variable in the system that will soak up all remaining transitory dynamics in the system. We impose that this variable has no long-run impact on the three measures within our core system, but we return to this below.

Table 1: The core structural VAR equations

Variable:	Measured in data as:		Model expression:	Long-run restrictions:
Solow	$Y_{jt}\left(N_{jt}^{\alpha}K_{jt}^{\beta}M_{jt}^{1-\alpha-\beta}\right)^{-1}$	$=$	A_{jt}	Independent of η^{ω} and η^{f}
WNULC	$\left(W_{jt}N_{jt}/Y_{jt}\right)W_{jt}^{-\alpha}$	$=$	$\alpha^{1-\alpha}A_{jt}^{-1}P_{jt}^{F},$	Independent of η^{ω}
WND	$Y_{jt}W_{jt}^{\sigma\alpha}$	$=$	$\psi Y_{t}P_{t}^{\sigma}A_{jt}^{\sigma}\left(P_{jt}^{F}\right)^{-\sigma}\Omega_{jt}$	–

Note: *Solow* is the physical Solow residual (TFPQ), $WNULC$ is wage-neutral unit labor cost and WND is wage-neutral demand. ψ is a constant such that $\psi\equiv\left(\frac{1}{\alpha}\right)^{-\sigma\alpha}\left(\frac{\sigma}{\sigma-1}\right)^{-\sigma}$.

What is important to note at this stage is that we only need to impose the zero-impact restrictions of the last column in the long run. Hence, we do not impose any other elements of the model, and we do not make assumptions regarding the short-run dynamics. Our identification of the technology shocks (η^{a}) are therefore consistent with changes in inventories, factor utilization, markups, or idiosyncratic input prices altering the Solow residual as long as these changes are mean reverting, i.e., as long as they do not affect the Solow residual in the long run.

The key assumption for distinguishing technology shocks from demand shocks is that technology shocks alter the physical Solow residual in the long run, whereas other shocks do not. This assumption implies that changes in the scale of operation are not allowed to permanently alter the efficiency of production. The most straightforward reason why this assumption may prove invalid is that firms might use a production technology with non-constant returns to scale. For this reason, we also provide results from variations of the model where we assume that the production functions exhibit increasing or decreasing returns. Details regarding the modifications are found in appendix A.

3 Data and estimation of the shocks

3.1 Data and measurement

Our primary data source is the Swedish Industry Statistics Survey (IS). It contains annual information on inputs, outputs, and firm-specific producer prices for all Swedish manufacturing plants with 10 employees or more from 1990 through 2002. We perform our analysis at the plant level, but because about 72 percent of the observations in our sample pertain to plants that are also firms, we refer to the plants as firms.

In our model, the technology shock η^a is the only shock that affects the Solow residual in the long run. This assumption is only credible if the Solow residual is calculated from a measure of real output where nominal output has been deflated by firm-specific prices. This is important because gross output deflated by sector-level price deflators (a measure often used in empirical analyses) will be a function of firm-specific idiosyncratic prices, which themselves are likely to depend on shocks other than technology (see Carlsson and Nordström-Skans 2012 for direct evidence). As our data-set contains a firm-specific price index built from plant-specific unit price changes,[9] we can derive a measure of gross output that is robust to changes in relative prices across firms, see Eslava et al. (2004) and Smeets and Warzynski (2013) for a similar strategy.

To take our model to the data, we rely on gross output throughout. We first compute a measure of firm-level changes in the physical Solow residual for firm j at time t. Letting lowercase letters denote logs, we use

$$\Delta a_{jt} = \Delta y_{jt} - \Delta z_{jt}, \tag{5}$$

where Δy_{jt} is the growth rate of real gross output, and Δz_{jt} is a cost-share-weighted input index defined as $C_K \Delta k_{jt} + C_N \Delta n_{jt} + C_M \Delta m_{jt}$ where Δk_{jt} is the growth rate of the capital stock (see details in appendix B), Δn_{jt} is the growth rate of labor input, and Δm_{jt} is the growth rate of intermediate materials and energy. C_J terms are the cost shares of factor J in total costs. We measure the cost shares as time-constant averages by two-digit industry code.

Using data on factor compensations, changes in output, and changes in inputs, we can thus calculate the residual Δa_{jt}, which provides an estimate of changes in the physical Solow residual. As argued above, this might not accurately measure technology shocks (η^a) due to varying factor utilization, inventories, or truly idiosyncratic factor prices, but the SVAR will filter out true technology shocks from eq. (5) as long as η^a is the only factor that permanently shifts A_{jt}. Material inputs are deflated using three-digit sectoral price indices, which implies that our model allows not only for an arbitrary set of transitory

[9] The index uses Paasche-type links. In cases where a plant-specific unit-value price is missing (e.g., when the firm introduces a new good), Statistics Sweden uses a price index for similar goods defined at the minimal level of aggregation (starting at the four-digit goods-code level). The disaggregated sectoral producer-price indices are only used when a plausible goods-price index is not available. Our identification is fully resilient to transitory errors in measured prices.

factor price shocks but also for permanent sectoral input price shocks at the three-digit level. However, our model cannot (for data reasons) account for truly idiosyncratic input price shocks if these are permanent.

We next compute $\Delta wnulc_{jt}$ and Δwnd_{jt}. Relying on cost minimization, we use C_N as the estimate of α and thus let it vary by two-digit industry. The rest of the components of $\Delta wnulc_{jt}$ are directly observed in the firm-level data. However, to compute wage-neutral demand (Δwnd_{jt}) we also need an estimate of the demand elasticity σ. We obtain this by estimating the demand equation (2) while instrumenting the firm idiosyncratic price using the Solow residual, as in Foster et al. (2008). The instrument is consistent with our initial assumptions, because the Solow residual is expected to affect firm-level sales only through firm-level prices. The results of this procedure suggest an elasticity of substitution equal to 3.306 (se 0.075), which we use when computing Δwnd_{jt}. The σ estimate is well in line with standard calibration exercises (see e.g., Erceg, Henderson, and Levin, 2000) as well as recent Swedish micro-evidence provided by Heyman, Svaleryd, and Vlachos (2008). As robustness checks, we also show that the main results are robust to using sector-specific estimates of σ and to using a very wide span of assumed values of σ.

We extract our baseline shocks by estimating a VAR on a sample of $6,137$ firms and $53,379$ firm/year observations (see appendix B for additional details on the data and for details on the construction of the final sample). Since the model uses lags, we can extract structural shocks for $41,105$ firm/years.

To analyze the impact of the shocks on the use of labor and the flows into and out of the firms, we link a longitudinal employer-employee data base (Statistics Sweden's register-based labor market statistics, or RAMS) to the firm-level data. These data are based on tax records and include the identity of all employees within the plants at the end of the year (November). We restrict the analysis to full-time employees within their main jobs. In the end, we are able to match shocks and labor flows for $40,451$ firm/year observations in $6,125$ firms. The final sample covers nearly two-thirds of all manufacturing employees.

3.2 Estimation

To derive the shocks of interest, we estimate a SVAR on the three variables defined in Table 1: Δa_{jt}, $\Delta wnulc_{jt}, \Delta wnd_{jt}$, and a fourth residual variable, which will be output (Δy_{jt}) unless otherwise noted. In practice, we first estimate four reduced-form equations where Δa_{jt}, $\Delta wnulc_{jt}, \Delta wnd_{jt}$, and the residual variable are explained by two lags of all four variables. We then invoke the long-run restrictions (including the long-run independence of the core system to the fourth residual shock) to derive the impulse responses of the structural shocks. Details regarding identification and estimation are found in appendix C.

The model also allows for firm-specific fixed effects to capture the drift terms of equations (3) and (4) as well as year dummies to capture aggregate shocks shared by different firms within the manufacturing sector, hence allowing us to concentrate on idiosyncratic disturbances. As a robustness check, we also estimate models accounting for sector-specific year dummies.

We use dynamic panel data methods building on Arellano and Bond (1991) for estimation because the asymptotic properties of the estimator rely on the cross-sectional dimension. This is a very useful feature in the current context of a large N ($6,137$ firms), but short T (12 years) panel because the identification of structural shocks with long-run restrictions crucially hinges on the quality of the estimated reduced-form coefficients and covariance matrix.

Table 2 shows descriptive statistics of the structural shocks derived for our baseline sample and model. The standard deviation of the demand shock is about 35 percent larger than the technology shock (16.02 and 11.1, respectively). Appendix C depicts the shock distributions in graphs and also shows impulse responses and variance decompositions related to the main SVAR model. In addition, the appendix discusses specification tests.

Two particular results are relevant for the analysis ahead. First, we find a fairly limited amount of dynamics, in particular in the Solow residual. The main reason for this finding is that the Solow residual is defined in physical gross terms and much of the dynamics in standard measures of Solow residuals appear to be due to the dynamics of idiosyncratic prices (see Carlsson and Nordström-Skans 2012, for direct evidence on relative-price dynamics). Second, shocks to the residual fourth variable explain little of the variance in our key variables at all horizons. Since the model is estimated conditional on time dummies, this finding is in line with the result of Franco and Philippon (2007), which shows that transitory shocks, although highly correlated across firms (and therefore of macroeconomic importance), matter only marginally at the firm level.

Table 2: Demand and technology shocks

	(1) Mean	(2) sd	(3) p(25)	(4) p(75)	(5) Firms	(6) Observations
Technology (η_a)	–	0.101	-0.056	0.058	6,137	41,105
Demand (η_ω)	–	0.162	-0.086	0.085	6,137	41,105

Note: p(N) denotes the Nth percentile of the data.

3.3 Validation

Because the shocks we are analyzing are idiosyncratic, we cannot use correlations with known aggregate shocks such as oil-price or exchange-rate movements to cross-validate their interpretation, at least not without strong priors regarding differences between firms in the sensitivity to these aggregate shocks. Instead, we perform two alternative corroboration exercises.

A first piece of evidence supporting our interpretation of the shocks is presented in appendix C, which shows theory-consistent impulse responses for the three unrestricted responses within the VAR system: The estimated response of $\Delta wnulc_{jt}$ to a technology shock is negative, as predicted from the theoretical model. Similarly, the estimated responses from both technology shocks and factor prices on Δwnd_{jt} are negative.

A second piece of evidence comes from relating the structural shocks to the firm-specific price index and to output. It should be clear from the model presented in section 2 that a positive technology shock can affect sales only if prices go down (since the demand curve is fixed). In contrast, demand shocks, defined as shifts in the firm-specific demand curve, allow the firm to sell more at a given price. This suggests that prices should remain unchanged in the response to a demand shock unless (i) the firm's technology features non-constant returns to scale, (ii) input prices change when the scale of production is altered.

Hence, theory suggests that technology and demand shocks should affect output, whereas prices should primarily respond when technology changes. To assess these predictions, we reestimate the SVAR and compare responses of output and prices to the two shocks (using output and prices, in turn, as the fourth variable in the SVAR system).

Figure 1 shows the impulse responses of output and idiosyncratic prices to technology and demand shocks—indicating that both types of shocks are important for firm-level aggregates. The figure also clearly validates the theoretical predictions discussed above: A 1 standard deviation (sd) technology shock increases output by 6 percent in the long run.

In the case of a 1 sd demand shock, output rises by 10 percent. Moreover, as expected, prices go down in the case of a technology shock, but prices change very little when the demand curve shifts. That the demand shock permanently changes output without altering relative prices strongly supports the interpretation of the demand shock as an idiosyncratic shift in the demand curve. Note that these results are not imposed from the construction of our variables: in particular, prices could well (from a pure measurement standpoint) respond in either direction to structural innovations in technology and demand.

Figure 1: Output and price responses

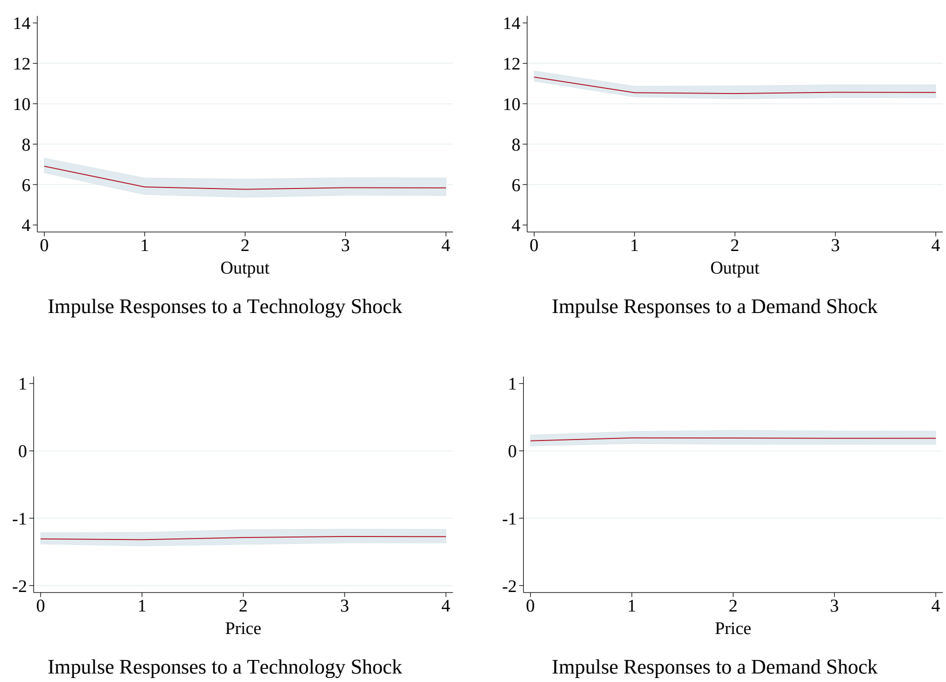

Note: Inpulse responses are expressed in percentage points. Lines depict the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervals calculated from 1000 replications.

4 Results

4.1 Idiosyncratic shocks and employment adjustment

A first objective of our analysis is to illustrate how firm-level employment responds to permanent shifts in idiosyncratic production functions and demand curves. Figure 2 shows impulse responses of log employment with bootstrapped confidence bands. The responses are derived from our SVAR system, using employment as the fourth variable.

The figure shows that idiosyncratic demand shocks have substantially more impact than the corresponding technology shocks on firm-level labor adjustments. A positive demand shock of 1 sd increases employment by slightly more than 6 percentage points, whereas the impact of an equivalent technology shock raises employment by only 0.5 percentage points.

It is also evident from figure 2 that the dynamics of labor adjustments are fairly limited. More than 90 percent of the long-run adjustments in response to the permanent shocks occur within the first year.

Figure 2: Employment responses

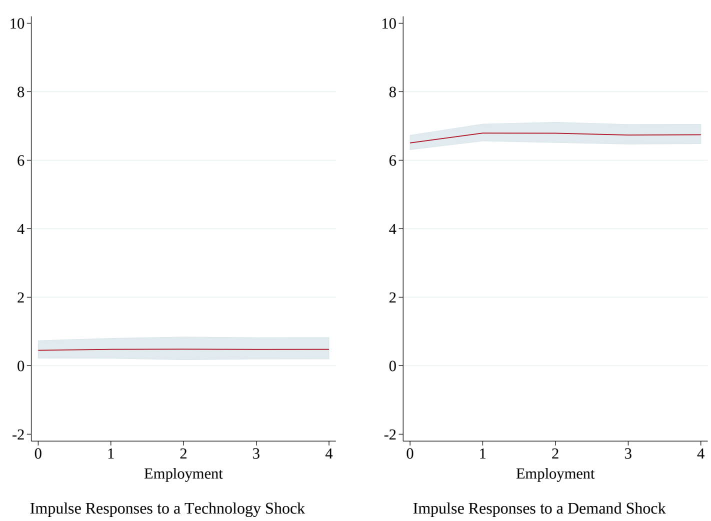

Note: Inpulse responses are expressed in percentage points. Lines depict the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95 percent confidence intervals calculated from 1,000 replications.

Much of the analysis that follows below relies on exploring non-linear responses to the estimated shocks. To this end, we proceed in three steps: We (i) estimate the SVAR, (ii) extract the ensuing measures of structural shocks, and (iii) relate the structural shocks to different outcomes in a standard regression framework.[10] This gives us additional flexibility in the specifications, which will be exploited in the next sections to assess potential asymmetries and non-linearities in the labor adjustment responses and allows us to present the results in a more compact table format. Empirically, we estimate the following equation in the baseline specification:

$$Outcome_{jt} = \eta_{jt}^{a}\delta_1 + \eta_{jt}^{\omega}\delta_2 + \rho_t\beta_\rho + \mu_j + \xi_{jt}, \tag{6}$$

where *Outcome* denotes employment (or hirings and separations in the following sections) for firm j at time t. The coefficients δ_1 and δ_2 capture the impact of the firm-level structural shocks on the outcomes. Moreover, we include time, ρ_t, and firm-fixed effects, μ_j, in line with the SVAR formulation above. This ensures that identification is driven by idiosyncratic, rather than aggregate, shocks.[11] Equation (6) shows the short-run impact of the shocks. We also present the long-run impact, measured as the sum of the contemporary effect and the impact of the first lag in the shock series.

Our baseline specification, following equation (6) is presented in the first column of table 3, and as is evident (and expected), the results closely mimic the impulse responses presented in figure 2: In the short run, employment increases by 6 percentage points in response to a positive demand shock of 1 sd. The coefficient of an equivalent technology shock is 0.15, and is non-statistically different from 0. If we add one lag of the shocks to the regression and calculate the long-run employment responses (column 4), the technology shock becomes somewhat larger and also statistically significant. However, long- and short-run responses are of a fairly similar magnitude, which corroborates our findings of limited dynamics in the labor adjustment. As before, firms' demand shocks continue to be the main driver of employment adjustments: A positive 1 sd shock to the demand curve increases employment in the long run in 6.4 percentage points, while the equivalent technology shock increases employment by 0.5 percentage points.

The constant returns to scale (RTS) assumption used in the empirical model and in the construction of the Solow residual is potentially controversial. In Carlsson, Messina, and

[10] Formally, this model is exposed to a potential generated regressor bias, but we show that all key results hold when either estimating them internally in the VAR or when relying on an IV-strategy (see below), both of which are insensitive to generated regressor biases.

[11] Since the shocks are identified as structural orthogonal innovations, they are uncorrelated with each other conditional on the year and firm-fixed effects of the SVAR.

Table 3: Contemporaneous and long-run effect on log employment under different returns to scale assumptions

	SHORT RUN			LONG RUN		
	(1)	(2)	(3)	(4)	(5)	(6)
	RTS=1	RTS=0.9	RTS = 1.1	RTS=1	RTS=0.9	RTS = 1.1
Technology (η_a)	0.153	0.955**	-0.492**	0.504*	1.378**	-0.244
	(0.159)	(0.161)	(0.149)	(0.214)	(0.211)	(0.232)
Demand (η_ω)	5.986**	6.149**	5.541**	6.357**	6.313**	5.978**
	(0.233)	(0.233)	(0.223)	(0.310)	(0.310)	(0.301)
Observations	40,451	41,132	39,788	34,414	35,031	33,811
Firms	6,125	6,193	6,065	6,116	6,184	6,055
Sd. η_a	10.06	10.04	10.37	10.06	10.04	10.37
Sd. η_ω	16.18	18.74	13.45	16.18	18.74	13.45

Note: Effect of one s.d. shock. Robust standard errors in parenthesis. Regression includes firm fixed effects and time dummies. Long-run estimates are obtained by adding the contemporaneous impact and one lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

Nordström-Skans (2014), we estimate RTS separately for the durables and non-durables sectors among Swedish manufacturing firms, obtaining 0.98 for non-durables and 0.9 for durables. In both cases we cannot reject the null of constant RTS, but there is some uncertainty in particular for sectors producing durable goods.

The model can be altered to accommodate increasing or decreasing RTS. Changing the assumed RTS affects the measures that are fed into the SVAR (for details, see appendix A) and hence also the estimated magnitudes of employment adjustments. However, the main message remains robust throughout. Column 2 in table 3 reestimates the model, imposing 0.9 in the construction of the Solow residual. A positive technology shock of 1 sd raises employment now by 1 percentage point in the short run (1.4 in the long run, see column 5). But this estimate still remains far below the estimated impact of a demand shock: an increase of 6.1 percentage points in the short run and 6.3 in the long run. If instead we impose an RTS coefficient of 1.1, the results change in the other direction (the impact of technology turns negative), but the main message regarding the strong relative importance of demand remains unaltered.

4.2 Robustness

We proceed by carrying out a battery of checks to assess the robustness of our first set of findings—namely, that (i) firm-level demand shocks are more important in the determination of labor adjustments than firm-level technology shocks, and (ii) employment adjustment to the permanent shocks is very rapid, exhibiting limited short-term dynamics. In all cases we use the specification presented in equation (6). We discuss the main findings here, but present the regression tables in appendix D to conserve space.

Demand elasticity. The baseline model uses an estimated demand elasticity of 3.3. As a robustness check we have allowed for industry-specific estimates of the demand elasticity. As appendix's table D1 shows (column 2), this does not alter the results. The main results also remain unaltered if we instead replace the year dummies for industry-by-year dummies, which controls for different employment trends across sectors (column 3). In addition, we have verified that our key results are robust to demand elasticities that vary within what we believe to be the full range of plausible values (from 1.1 to 10); the results in table D1 (columns 4 and 5) show that the estimated coefficients of technology and demand shocks are fairly stable despite this large interval of demand elasticities. The main reason for the low sensitivity of the estimated employment adjustment to alternative calibrations of the demand elasticity is that the latter enters the system with a weight equal to the labor share in gross production, which is fairly low (around 0.25 on average).

Sectoral heterogeneity. The dynamic panel approach used for estimation took advantage of our large-N small-T panel setting to estimate the VAR system with considerable precision. This is a key advantage relative to standard SVAR estimations in the macro literature. A potential cost, however, is that the underlying dynamic processes are assumed to be equal across different firm types. To address this concern, we have reestimated the model to allow for separate dynamics for each two-digit industry, and the employment adjustment results remain unchanged (see column 6 in appendix's table D1).

Sample selection. The data appendix (appendix B) explains that the output allocation across plants within (the relatively few) multi-plant firms after 1996 is imputed in the IS data set. We have therefore redone the analysis for the single-plant firms in the sample (column 2 of appendix table D2), as well as for a mixed sample including multi-plant firms until 1996, but not thereafter (column 3 of table D2). The results are robust in these alternative samples which is unsurprising since the bulk of the original sample is unaffected.[12] The results are also unchanged when the shock distribution is truncated into the Lester range of -2 to 2 sd (see column 4 of table D2).

Alternative fourth variable in the SVAR. We have varied our model to ensure that the limited dynamics in the employment adjustments we find is not due to the specific way we handle the residual dynamics in the system. In particular, we have used value added per worker, output, and employment from our two data sources (RAMS and IS) as alternative fourth variables. Table D3 shows that these variations only have minor impacts on both the estimated dynamics and the long-run adjustments.[13]

Firm exit. Finally, a possible concern with the analysis is that we disregard the firm exit process. Firms are likely to exit in response to severe negative demand or technology shocks, and this process may impact labor dynamics. To address this concern, we have analyzed the relationship between the shocks and the probability of exit from the sample using a kernel-weighted local polynomial regression. The results are shown in figure 3. As is evident, the main driver of firm exit is large negative demand shocks.[14] That demand shocks are more important for firm exits than technology shocks is well in line with results for the United States in Foster et al. (2008). To determine whether this finding has any bearing on the results regarding employment adjustments, we have analyzed the employment impact of the shocks using a two-periods model instead of the one-period baseline (see appendix's table D4). In practice, this implies that we relate the shock to the net employment growth across two years, defined as the change in employment divided by the average employment in the two years as in Davis et al. (1996). Since the labor flows are defined even if all workers exit the year after the shock, we can calculate the impact of the shocks while excluding or including the firms that exit. Reassuringly, the results are insensitive to whether we include or exclude exiting firms.[15]

Overall, our findings strongly suggest that (i) permanent shifts in firms' idiosyncratic demand curves are a key determinant of firms' idiosyncratic net employment adjustments, and (ii) the pace of labor adjustment is relatively fast. In contrast, permanent shifts in

[12] As explained in Section 3, 72 percent of plants are in single-plant firms and the imputation only affect the later half of the sample.

[13] That the fourth variable plays a negligible role in employment adjustment is also suggested in the variance decomposition shown in appendix C.

[14] Note that large positive shocks are also associated with firms' exit, although to a lesser extent. This indicates that firms facing large positive shocks are more likely to be absorbed by other companies through mergers and acquisitions.

[15] Note that our sample excludes new establishments because of the dynamic identification strategy, which requires a firm to be observed at least five consecutive years to be included in the estimation. Therefore, the mean exit rate is somewhat lower than in the full sample.

Figure 3: Firms' exit

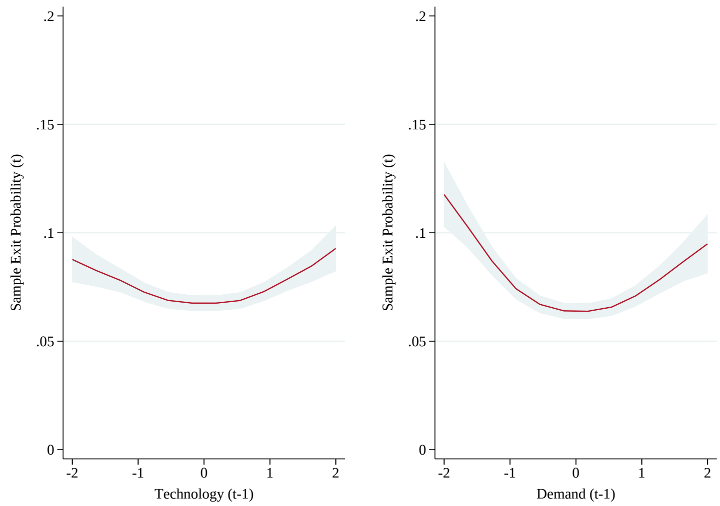

Note: Sample exit probability as a (non-linear) function of an x sd lagged technology or demand shock. Shaded areas depict 95 percent confidence intervals.

firms' physical production functions (i.e., technology shocks) appear to play a much more limited role in firms' labor adjustment, despite being crucial to the evolution of both output and relative prices.

4.3 Idiosyncratic shocks, hires, and separations

To analyze which margins firms use when hit by permanent idiosyncratic demand and technology shocks, we relate our measures of shocks to firm-level measures of worker flows (i.e., hires and separations). Using annual individual-level employment data on end-of-the-year employment that are matched to our firm-level data, we compute measures of job and worker flows using the metrics proposed by Davis et al. (1996).

Net employment growth is defined as the change in employment relative to the preceding year, divided by the average employment during the two years. Similarly, we define the hiring (separation) rate as the number of new (separated) employees between t and $t-1$, divided by the average number of employees during the two years. With these definitions, net employment growth will be the difference between the hiring rate and the separation rate, and the timing of the flows are defined such that the flow equation of employment holds.[16]

We also calculate a measure of the short-tenured (< 3 years) separation rate using the same denominator as for the other rates. All in all, we can match these flow measures to $6,130$ firms in the firm data (described above). Table 4 displays descriptive statistics of these measures for the sample in which structural shocks can be constructed ($40,451$ observations). The average hiring rate during the observation period is 15 percent, and the average separation rate is 14 percent, whereof slightly less than half (6 percent) are separations of short-tenured workers.

We proceed analogously to the employment analysis of the preceding section. Hence, we follow equation (6) for hirings and separations, separately. Table 5 present the effects of a 1 sd technology and demand shocks in both the short and long run, as well as the corresponding elasticities (see appendix C regarding the computation of the elasticities). A normal demand shock is estimated to increase the hiring rate by 2.9 percentage points and reduce the separation rate by 2.7 percentage points in the short run (slightly more in the long run).[17] These numbers should be compared with average hiring and separation rates of about 14 to 15 percent each, as shown in table 4 above. The estimates imply that, on average, 52 percent of the net employment adjustment is obtained using the hiring margin, and 48 percent using the separation margin. Firms thus, on average, rely as much on variations in separations as on variations in hires when responding to the shocks.[18]

The results also imply that the low response of net employment to technology shocks does not mask any substantive counteracting responses in terms of gross flows. Rather, idiosyncratic technology shocks appear to have a limited impact on both hiring and separation rates in both the short run and the long run. As a final result, we see that separations of short-tenured workers make up slightly more than one-third of the total short-run separation response to demand shocks, but a lower fraction of the longer-run responses.[19]

[16] That is, $Employment_t = Employment_{t-1} + Hires_t - Separations_t$.

[17] Note that the difference in the estimated coefficients of the hiring and separation rate result in the impact of the shocks on the net employment change (not reported in the table). However, there is a marginal difference between the implied results on employment changes presented here and those reported in columns 1 and 4 of Table 3. This arises because we here follow the metric proposed by Davis, Haltiwanger, and Schuh (1996) to measure net employment changes, instead of the log differences we relied on in the previous subsection. The differences are, however, far from altering our main conclusions.

[18] This result is also interesting in the light of the literature on labor flows and the business cycles (see Shimer, 2012; and Fujita and Ramey, 2009). It suggests that any quantitatively important asymmetries between hiring and separations over the business cycles should be explained by asymmetries in the market responses, and not as asymmetries in firm-level labor adjustment behavior.

[19] The lower relative contribution of short-tenured separations in the long run is consistent with a reduction in contemporary hirings, which reduces the number of short-tenured workers who can be released in the

Table 4: Summary statistics; worker's data

	Category	(1) Mean	(2) sd	(3) p(25)	(4) p(75)	(5) Firms	(6) Observations
Net Employment Rate	overall	0.012	0.208	-0.062	0.089	6,125	40,451
	within		0.195				
Hiring Rate	overall	0.150	0.151	0.063	0.200	6,125	40,451
	within		0.127				
Separation Rate	overall	0.138	0.152	0.061	0.174	6,125	40,451
	within		0.131				
ST Separation Rate	overall	0.061	0.082	0	0.083	6,125	40,451
	within		0.065				

Note: The"within"rows show the dispersion within establishments. p(N) denotes the Nth percentile of the data.

Table 5: Permanent demand and technology shocks, hirings, and separations

	SHORT RUN			LONG RUN		
	(1) Hiring Rt.	(2) Sep. Rt.	(3) ST Sep Rt.	(4) Hiring Rt.	(5) Sep. Rt.	(6) ST Sep Rt.
A) 1 sd shock:						
Technology (η_a)	-0.050	-0.165*	-0.117**	-0.093	-0.504**	-0.177**
	(0.075)	(0.078)	(0.038)	(0.116)	(0.128)	(0.066)
Demand (η_ω)	2.906**	-2.703**	-1.010**	3.125**	-2.884**	-0.416**
	(0.096)	(0.120)	(0.052)	(0.156)	(0.186)	(0.076)
B) Elasticities:						
Technology (η_a)	-0.005	-0.016*	-0.012**	-0.009	-0.050**	-0.018**
	(0.007)	(0.008)	(0.004)	(0.011)	(0.013)	(0.007)
Demand (η_ω)	0.180**	-0.167**	-0.062**	0.193**	-0.178**	-0.026**
	(0.006)	(0.007)	(0.003)	(0.010)	(0.012)	(0.005)
Observations	40,451	40,451	40,451	34,414	34,414	34,414

Note: Robust standard errors in parenthesis. Hiring Rt: Hiring rate; Sep. Rt: Separation rate; ST Sep. Rt.: Short-tenured separation rate. Hiring and separation rates are measured as the flow between the end points of two years divided by the average employment across these two points in time. ST Sep. Rt. is measured as the number of separations of short-tenured ($<$ 3 years) workers divided by the same denominator as the other rates. Regressions include time dummies and firm fixed effects. The long-run impact is based on the sum of the contemporary effect and the effect of the first lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

4.4 Asymmetry and non-linearity

To examine potential non-linearities in the hiring and separation responses depending on the signs and magnitudes of the shocks, we extend the model presented in equation (6) by allowing for separate second-order polynomials above and below zero. Because the dynamics add few insights, we focus on the short-run impact.

Figure 4 shows how firms adjust hirings in response to positive and negative shocks of different magnitudes. To facilitate the interpretation, the graphs show the sum of the average hiring rate among firms that do not adjust employment (about 10 percent) and the predicted estimates for various deviations from a zero-shock state. For completeness, we show the responses to both technology and demand shocks, but we focus our attention toward the demand-shock responses. (Throughout, we find limited adjustments in response to technology shocks, as expected from the previous subsections.)

Two patterns are particularly noteworthy: First, the hiring response is considerably smaller if the shocks are negative. Second, the impact of a 2 sd positive shock is exactly twice that of a 1 sd positive shock, suggesting that the costs of increasing hirings are a linear function of the magnitude of the adjustment.

Figure 5 shows the corresponding patterns for separations. The shapes and magnitudes (again focusing on the demand shocks) are not far from mirror images of the impact on hirings. Thus, separations primarily respond to negative shocks. Although separations do go down somewhat when shocks are positive, this impact is even smaller than the hiring cuts in response to negative demand shocks. Symmetrically to the hiring response, the estimates imply that a 2 sd negative shock causes a separation response that is exactly twice as large as the response to a 1 sd negative shock, which suggests that the costs of increasing separations are approximately linear on average. Notably, the results of figures 4 and 5 imply that firms primarily use separations when responding to permanent negative demand shocks, an issue to which we return to below.

Finally, figure 6 shows the impact on net employment and, as could be imagined from the combination of figure 4 and figure 5, these effects add up to a fairly linear relationship. The somewhat more curved pattern on the positive side arises because the kink at zero is more pronounced for hires than for separations. This difference in curvature is statistically significant, but the magnitude is fairly small: The net employment changes in response to a 2 sd positive demand shock (9 percentage points) is reasonably close to the response to

next period.

Figure 4: Shocks and the hiring rate

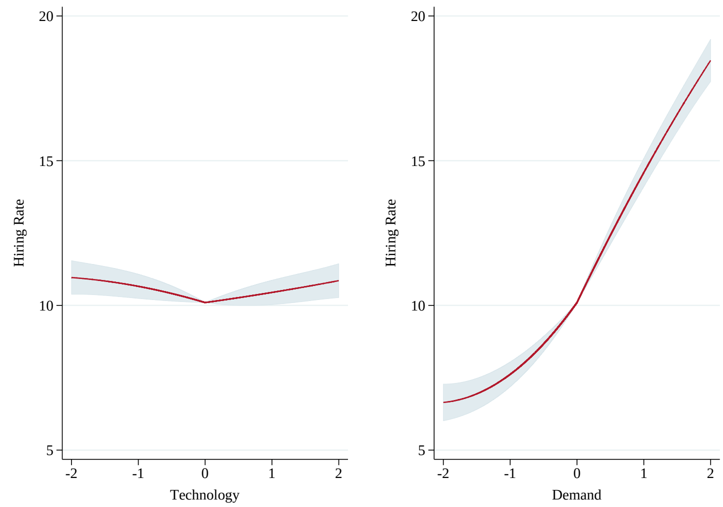

Note: Each line represents the sum of the average hiring rate among firms that do not adjust employment (10 percent) and the response of the hiring rate in percentage units as a (non-linear) function of an x sd of technology and demand shocks. Shaded areas depict 95 percent confidence intervals.

Figure 5: Shocks and the separation rate

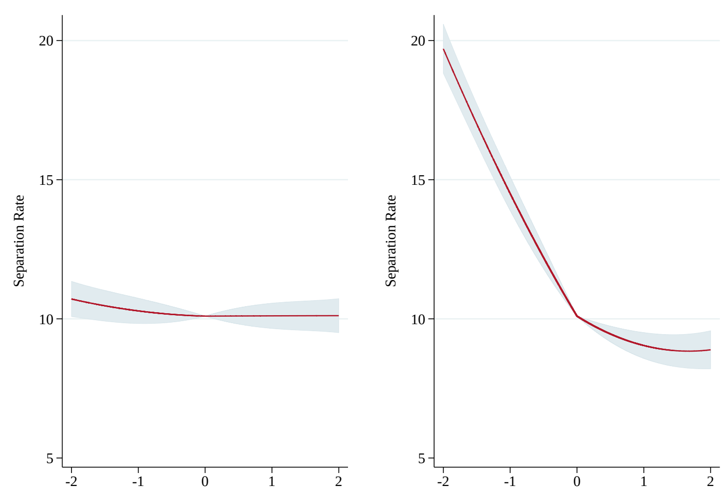

Note: Each line represents the sum of the average separation rate among firms that do not adjust employment (10 percent) and the response of the separation rate in percentage units as a (non-linear) function of an x sd of technology and demand shocks. Shaded areas depict 95 percent confidence intervals.

Figure 6: Shocks and the net employment rate

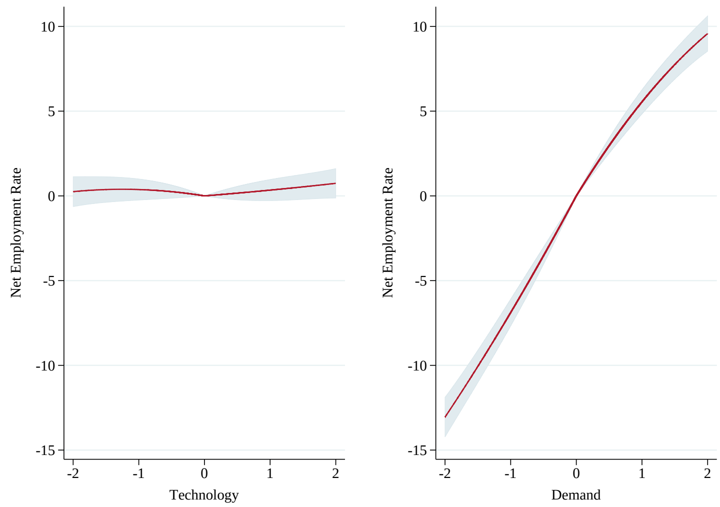

Note: Each line represents the response of the net employment rate in percentage units as a (non-linear) function of an x standard deviation of technology and demand shocks. Shaded areas depict 95 percent confidence intervals.

a 2 sd negative shock (-13 percentage points) in absolute values.

4.5 Decomposing net employment adjustment in response to permanent demand shocks

This subsection provides an explicit decomposition of firm-level net employment adjustments in response to permanent demand shocks.[20] This analysis is similar in spirit to Abowd et al. (1999), and Davis et al. (2012), which provide decomposition exercises of the relative contribution of various worker flows to the observed employment changes in French and U.S. firms, respectively. In contrast to these previous studies, however, we analyze changes in hires and separations induced by employment adjustments due to a demand shock. This allows us to obtain a causal correspondent to the decompositions in the earlier literature. In our case, demand shocks drive the changes in employment, and we can therefore abstract from, for example, possible exogenous separations which may affect firms' employment levels in the short run.

[20] We focus on permanent demand shocks because technology shocks are found to have negligible impacts on net employment.

In practice, we characterize labor adjustments by two second-order polynomials, one for positive values and one for negative values. We then instrument this adjustment by a similarly constructed set of polynomials in the demand shock. We use the hiring rate as our outcome, but since net employment adjustment is identical to the difference between hirings and separations, the impact on separations is easily deduced.[21]

Figure 7: The hiring rate and net employment changes: IV results

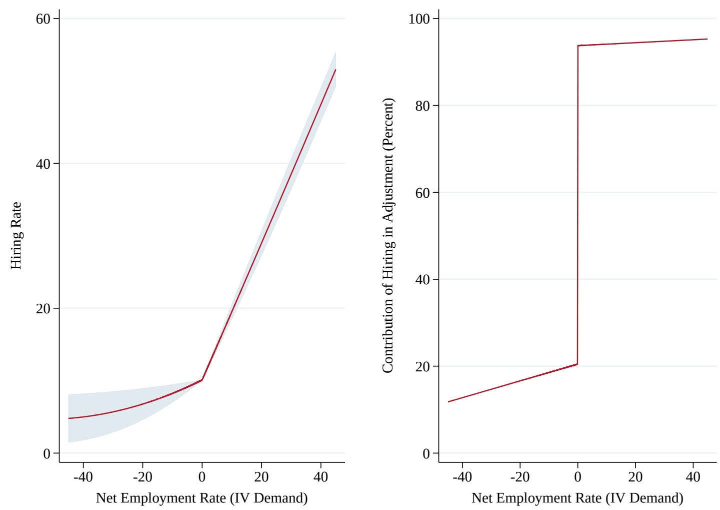

Note: Left-side panel: Contemporaneous hiring rate in percentage units as a (non-linear) function of employment adjustment in percentage units. Employment adjustments are instrumented by demand shocks. The model imposes a separate quadratic polynomial above and below zero for both employment adjustment and the instrument. Shaded areas depicts 95 percent confidence intervals. Right-side panel: Implied fraction of employment adjustment achieved through changes in hirings as a function of the size and magnitude of the employment adjustment.

The results are presented in the left-hand panel of figure 7. They imply a strong and linear relationship between net employment adjustments and hires when the employment adjustments are positive, but a very modest relationship when the employment adjustments are negative. The right-hand panel of figure 7 shows the share of employment adjustment through hires as a function of demand-induced net employment changes. This share jumps from 20 percent to 95 percent when employment adjustments become positive instead of negative.[22]

Figure 7 also suggests that firms are relatively unconstrained in their use of separa-

[21] The instrumental variable (IV) strategy essentially implies that we scale the shock impact on hirings

Figure 8: Actual (IV) and simulated hiring responses to employment changes

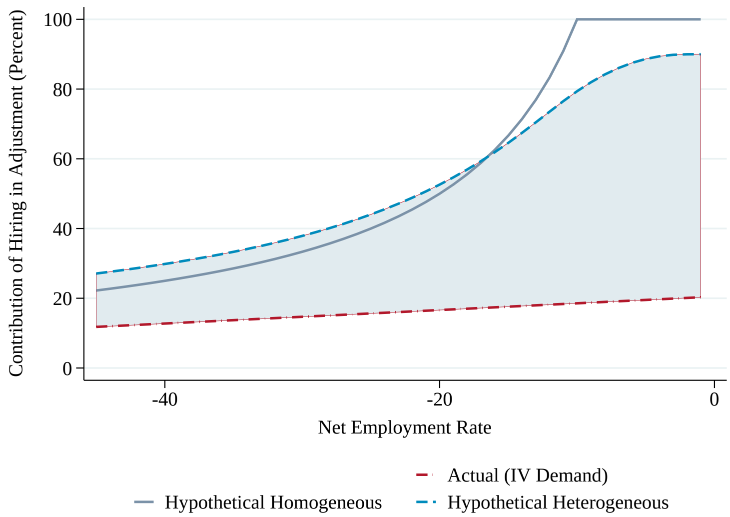

Note: Actual (estimated from data) and hypothetical maximum (simulated) fraction of negtive employment adjustments achieved through changes in hirings. Employment adjustments are instrumented by demand shocks."Hypothetical homogenous" assumes that the same fraction of workers always leaves the firm."Hypothetical heterogenous" imposes a random individual quit rate on the actual firm-size distribution.

tions, since they rely on increased separations even when they could have adjusted through reduced hires. To make this point precise, figure 8 repeats the patterns shown in the right-hand panel of figure 7 but focuses on negative values. As benchmarks illustrating what the firms *could have done*, the figure also depicts two hypothetical adjustment curves. The first, denoted "hypothetical homogeneous," assumes homogenous firms and imposes the empirical steady-state (i.e., without employment changes) separation rate of 10 percent on all the firms. In this case, as long as the need for adjustment is 10 percent or less, reduced hires could fully accommodate the necessary adjustments. If the shock is 20 (30) percent instead, the firm could instead accommodate half (one-third) of the adjustment through reduced hires. Notably, this curve assumes that 10 percent of employees leave each firm every year, which clearly cannot be the case.

presented in figure 4 above with the first stage, which corresponds to figure 6.

[22] Note that, in contrast to figures 4 and 5 (where the zeros refer to the absence of an idiosyncratic shock), zero here refers to the state when net employment adjustment is predicted to be zero based on the full first stage (i.e., based on the combination of the shock polynomials, the year dummies, and the firm-fixed effects).

We therefore also provide a second benchmark, assuming instead that the *individual* probability of leaving a firm is 10 percent. By randomly allocating quits across the workers in our full sample and then aggregating to the firm level, we get the firm-level distribution of quit rates. With this distribution, which naturally widens if firms are small, some firms will not experience any quits at all, which means that they cannot accommodate even the smallest employment adjustment through reduced hires, whereas other firms will experience many random separations, allowing them to accommodate large employment reductions through reduced hires. The curve denoted "hypothetical heterogeneous" displays the simulated frontier of adjustments with random individual quits using our actual distribution of firm sizes.

The logic behind the hypothetical curves is that they provide a baseline indicating how firms would behave in a completely rigid world where firing is prohibitively costly as long as firms are hiring someone. In this case firms would always adjust according to the hypothetical heterogeneous curve in figure 8. As is evident, the observed employment adjustments are far from this rigidity benchmark. The actual share of adjustment through reduced hires is much lower than the hypothetical reliance on separations would allow for. The shaded area between the heterogeneous hypothetical curve and the actual behavior of the firm could be interpreted as a region of flexibility because it depicts the amount of negative labor adjustments through induced separations (i.e., separations above the random rate) which could have been accomplished through reduced hires instead.

One reason for the observed patterns may be that firms adjust by releasing marginal, short-tenured workers who are more likely to be on temporary contracts. Sweden is a country with slightly above-average levels of employment protection (OECD, 2014). The use of temporary contracts is flexible, whereas protection for workers with open-ended contracts remains restrictive. It is thus possible that the labor market responses studied here may hide important heterogeneity across workers, depending on their contract type and tenure with the firm.

We do not observe the contract type in the data, but in order to explore the role played by the (potential) flexibility provided by marginal workers, we have estimated the IV model using the following outcome variable: separation of short-tenured (less than three years) workers divided by average employment across the two years. The results, shown in figure 9, suggest that about half of the response to negative shocks come through reductions of short-tenured workers.

We have also repeated the simulation exercise presented in figure 8 above, but instead

Figure 9: The hiring rate of short-tenured workers and net employment changes: IV results

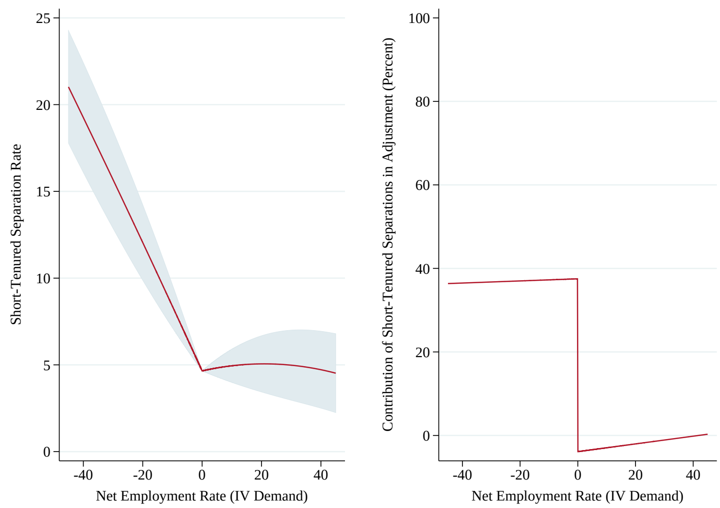

Note: Actual (estimated from data) and hypothetical maximum (simulated) fraction of negtive employment adjustments achieved through changes in hirings. Employment adjustments are instrumented by demand shocks. "Hypothetical homogenous" assumes that the same frction of workers always leaves the firm. "Hypothetical heterogenous" imposes a random individual quit rate on the actual firm-size distribution.

contrasting the actual combined adjustment of reduced hirings and increased separations of short-tenured workers with the maximum possible adjustment levels The results, presented in figure 10, show that firms are far from using the flexibility provided by these two margins. The substantial shaded area in the figure implies that firms rely much more on separations of long-tenured workers than they would have needed to in order to achieve the same level of net employment reduction.

5 Extensions

The results presented above suggest that labor adjustments in response to positive and negative demand shocks are fast and that separations are a flexible margin of adjustment when employment reductions are needed. This section provides two extensions that we believe shed light on the process.

We first explore, in subsection 5.1, the hypothesis that our results are driven by within-

Figure 10: Actual (IV) and simulated hiring plus short-tenured separation responses

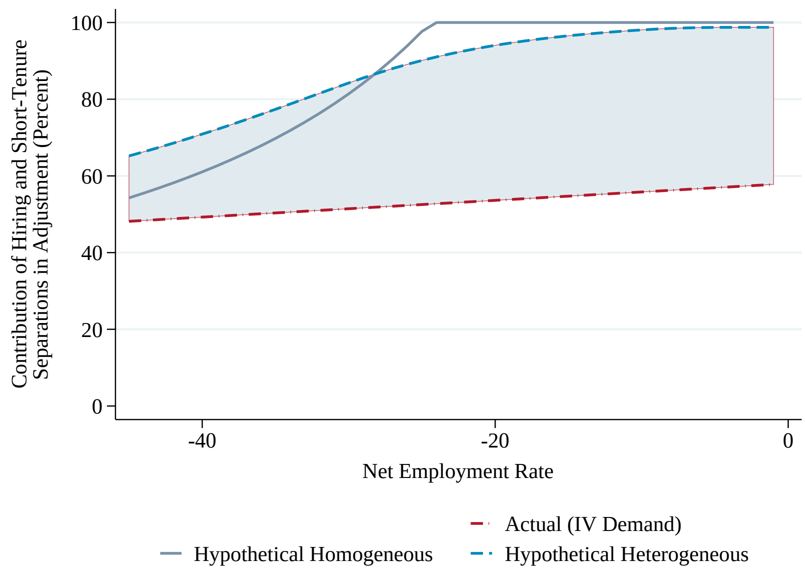

Note: Actual (estimated from data) and hypothetical maximum (simulated) fraction of negtive employment adjustments achieved through changes in hirings and short-tenured separations. Employment adjustments are instrumented by demand shocks. "Hypothetical homogenous" assumes that the same fraction of workers always leaves the firm and are employed on short tenure (less than three years). "Hypothetical heterogenous" imposes a random individual quit rate and short-tenure rate on the actual firm-size distribution.

firm heterogeneity in worker skills, which may create imbalances between the skills of those workers who quit at random and the optimal skill mix within the firm.

Thereafter, we note that if adjustments were costly, firms would be expected to hoard labor and/or refrain from hiring when shocks are perceived as temporary. The focus of the analysis so far, however, has been on how firms adjust employment, hirings, and separations when hit by *permanent* idiosyncratic shocks. As a contrast, subsection 5.2 analyzes the role of transitory shocks to demand.

5.1 Firm-level heterogeneity

Taken at face value, our results imply that firms either bear few costs to separate long-tenured workers, or rely heavily on a well-defined mix of worker types that is hard to change when demand changes. If the latter is true, it is more than likely that the workers who leave, or who are on temporary contracts, differ from the types of workers that the firms would like to separate from. Hence, it is not possible for the firm to fully exploit worker attrition or their pool of short-tenured workers to adjust to the shock.

To explore this further, we have reestimated our models separately for firms with a homogenous workforce in terms of field and education level and for firms with a heterogeneous workforce in the same dimensions. The idea is that firms with a more homogenous set of employees should care less about whom they separate from and thus rely more on attrition and the separation of short-tenured workers when adjusting their net employment.

Figure 11: The hiring rate and net employment changes: firm size and worker heterogeneity: IV results

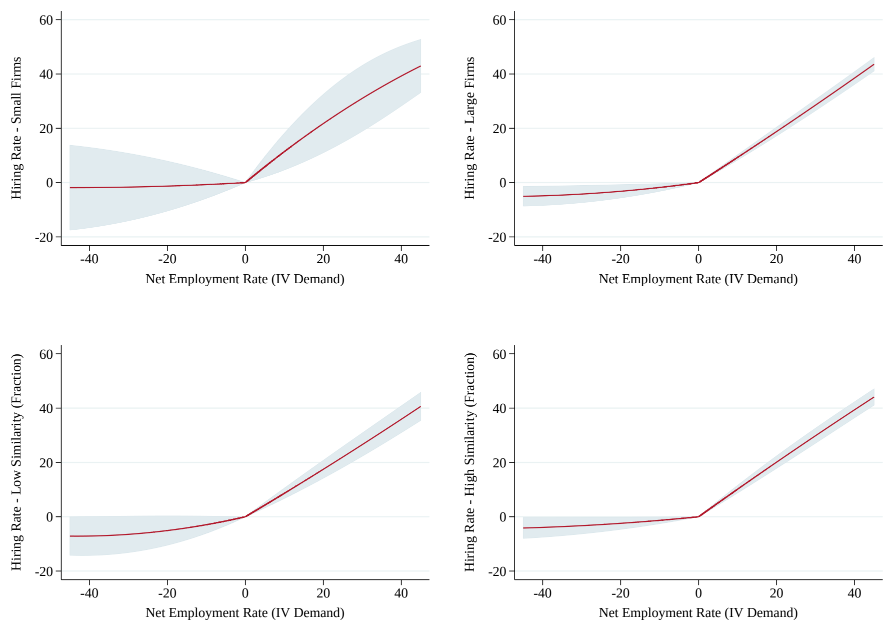

Note: Contemporaneous hiring rates in percentage units as a (non-linear) function of an x sd demand shock in subsamples defined by employee heterogeneity (lower graphs) and firm size (higher graphs). Employment adjustments are instrumented by demand shocks. Low (high) similarity firms are those with a similarity index (described in the text) below (above) the median. Small (large) firms are those with fewer (more) than 20 employees. Shaded areas depicts 95 percent confidence intervals.

In practice, we calculate the fraction of coworkers (to each worker in the data) that has the exact same type of education (three-digit field and two-digit level) and take the average of this share for each firm. This gives an index of the average worker's exposure to similarly trained workers within the firm. This procedure is a straightforward implementation of standard practices used for measuring segregation (see, e.g., Åslund and Nordström Skans 2009). In a second step, we split our firm-level data across the median of this index and analyze the two samples separately.

Figure 11 presents the results for the two samples, i.e., for firms with high versus low degrees of educational similarity among workers. As before, we characterize labor adjustments by two second-order polynomials, one for positive values and one for negative values. We then instrument this adjustment by a similarly constructed set of polynomials in the demand shock.

Quite surprisingly, we find little support for the notion that within-firm heterogeneity is an important explanation for the low reliance on separations when firms are hit by negative demand shocks. We would, however, like to acknowledge that our measures of staff heterogeneity may well be too crude to capture the role of firm-level heterogeneity in the adjustment patterns.

Figure 11 also shows results separately by firm size (more than 20 employees or fewer than 20 employees). The idea is again that if worker heterogeneity is important for the results, smaller firms are more likely to have difficulties using attrition and separation of short-tenured workers to adjust their staffs. The results, however, are very similar across the two size classes, displaying as before little signs of systematic heterogeneity.

5.2 Transitory idiosyncratic shocks

To derive a measure of transitory demand shocks, we follow the strategy used for the estimation of the demand elasticity σ. Hence, we use the Solow residual as an instrument for prices in an estimation of a log-linearized version of the demand equation (2), where time dummies control for aggregate shocks and firm fixed effects eliminate beteween-firm permanent heterogeneity. Since the ensuing residuals of the estimated equation represent changes in sales without price adjustments, they serve as a measure of demand shocks. This strategy is similar in spirit to the approach followed by Foster, Haltiwanger, and Syverson (2008). Thus, we label these shocks "FHS demand".

In contrast to the SVAR filter, however, this procedure does not differentiate between permanent and transitory shocks.[23] The correlation between FHS demand and our SVAR demand shocks is 0.538 (see table 6) and the standard deviation is considerably higher for

[23] In addition, these alternative series are not necessarily uncorrelated with other shocks, and the processes do not account for factor price shocks.

the FHS demand shocks (0.24 versus 0.16). Thus, the two demand-shock series appear to contain a common component without being identical.

Table 6: Baseline estimates vs. Solow residuals and FHS demand shocks

	SHORT RUN		LONG RUN	
	(1)	(2)	(3)	(4)
	Baseline	FHS	Baseline	FHS
A) 1 sd shock				
Technology (η_a)	0.153	0.333*	0.504*	0.993**
	(0.159)	(0.168)	(0.214)	(0.250)
Demand (η_ω)	5.986**	3.406**	6.357**	4.061**
	(0.233)	(0.183)	(0.310)	(0.252)
Observations	40,451	40,451	34,414	34,414
Firms	6,125	6,125	6,116	6,116

Note: In the FHS column the technology shock is the Solow residual, and the demand shock is FHS demand, as defined in the main text. Robust standard errors in parentheses. Regression includes time dummies and firm fixed effects. Long-run impact is based on the sum of the contemporary effect and the effect of the first lag. Regression sample limited to observations where the absolute value of both the technology and the demand shock is less than or equal to two sd. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

We estimate the impact on employment of FHS demand and the Solow residual as measures of shocks (see column 2 of table 6). To facilitate the comparison, estimates with our SVAR shocks are presented in column 1. The impact of the technology shock is somewhat larger using the Solow residual, and the estimated impact of demand shocks is about half as large when using FHS demand than when using the SVAR demand shock. However, the main message still holds. According to the estimates in column 2, the short-run impact of demand shocks is 10 times that of the technology shock in the short run (about 5 times the impact in the long run, as shown in column 4).

As noted above, FHS demand includes both transitory and permanent shocks. To ascertain whether the impact of demand shocks differs depending on the time-series properties of the shock, we have decomposed the FHS demand into permanent and transitory components. Running a regression with FHS demand as the dependent variable and the SVAR demand shocks as a regressor, we use the residual as a measure of transitory demand shocks. Because this residual is measured in the same units, we can directly compare its impact on employment adjustments with the SVAR demand shock, which represent per-

Figure 12: Net employment, permanent and transitory demand shocks

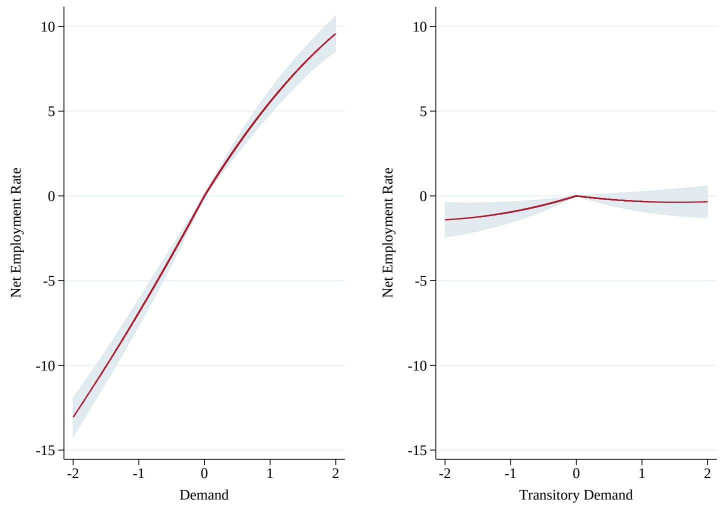

Note: Contemporaneous net employment rate in percentage units as a (non-linear) function of an x sd of the permanent and transitory (residual component of FHS demand) demand shocks. Shaded areas depict 95 percent confidence intervals.

manent shocks.[24]

In Figure 12 we analyze the impact of temporary and permanent shifts in product demand on net employment changes. As before, we allow for second-order polynomials of negative and positive shocks, respectively. The results show that the impact of the transitory shocks is substantially lower than the impact of the permanent shocks, and, for most magnitudes of the shock, is not statistically different from zero.

These results mirror those of Guiso et al. (2005), which shows that wages respond to permanent shocks but not to transitory shocks. In turn, this implies that firms' employment adjustment crucially depends on the time-series properties of the shocks. This is important because the welfare consequences of firms' lack of adjustment are likely to crucially depend on these properties. Labor hoarding in the face of temporary shocks may be welfare-enhancing in the presence of uninsurable labor market risk (Bertola, 2004). However, the

[24] The decomposition resembles Guiso et al. (2005), which extracts the permanent component of firm-level value added using high-order polynomials of lags as instruments. Although the mechanics of the methods differ, the underlying logic is similar. Because the bulk of the technology shock process is persistent enough to show up as permanent in our analysis (the correlation between the two shocks is 0.98), we do not perform a corresponding analysis for the technology shock.

inability of firms to structurally adjust to permanent shocks is much more likely to be welfare-reducing through allocative inefficiencies.

6 Conclusions

This paper has analyzed how firms adjust their labor inputs in response to permanent idiosyncratic firm-level shocks to technology and demand. We identify the shocks by imposing a set of long-run restrictions in an SVAR estimated on firm-level data. The restrictions are derived from a simple model using a first-order approximation of the firm that features constant cost shares and isoelastic demand. The SVAR is estimated using dynamic panel-data methods, allowing us to identify the parameters of the reduced form with considerable precision. To estimate the model, we rely on a unique data-set that merges information about inputs, outputs, and prices of Swedish manufacturing firms with a linked employer-employee data-set.

The shocks derived from the SVAR impact output and prices in a theory-consistent manner, which lends support to their interpretation as demand and technology disturbances. Firm-level output responds vigorously to both technology and demand shocks. In contrast, firm-level prices fall in response to positive technology shocks, but they remain independent of product demand innovations.

Our labor-adjustment results show that both the nature and the time-series properties of the shocks matter. Permanent demand shocks, which affect output but not relative prices, have a pronounced impact on employment. Technology shocks, on the other hand, have much more limited employment effects despite affecting both output and relative prices. Similarly, temporary demand shocks have little effect on employment adjustments.

Further results suggest that employment adjustments in response to permanent shifts in the product demand curve are fast and symmetric. By far the largest part of employment adjustment takes place within a year. Almost as much of the employment adjustments are through changes in the separation rates as through changes in the hiring rates, suggesting that both margins should be considered endogenous at the firm-level. Moreover, there are no signs of non-linear responses in hires or separations. Finally, the sign of the shock determines the primary margin of adjustment: firms primarily adjust through separations if shocks are negative and primarily though hires if shocks are positive.

The speed of adjustment, the symmetry between hires and separations as adjustment margins, and the continued recruitment of workers in the face of negative shocks jointly suggest that labor market rigidities play a very limited role in hampering firm-level labor adjustments in the face of permanent idiosyncratic demand shocks. However, the adjustments with respect to temporary shocks are heavily muted. Considering that OECD ranks Sweden slightly above the average country in terms of institutional labor market rigidities, this suggests that such levels of employment protection may provide enough flexibility for firms to accommodate the impact of permanent shocks, while pushing them toward hoarding labor when hit by temporary innovations.

A possible limitation of our study is the focus on the manufacturing sector, the sector for which technology shocks can be reasonably approximated. However, it seems likely that the overwhelming force of idiosyncratic demand shocks as a source of employment adjustments in manufacturing firms should provide a lower bound for the importance of demand within other sectors. In service sectors where product differentiation is arguably more important than in manufacturing, demand is likely to play an even more important role for reallocation than in manufacturing, where downstream firms are the primary customers.

Overall, our results imply that cross-country comparisons of labor flows need to be careful in accounting for the types of the shocks that hit these economies, because responses depend not only on the nature of the shocks (technology versus demand) but also on the time-series properties of these shocks: Labor market adjustments differ depending on the prevalence of permanent versus transitory components within the shock distribution.

Building on this notion, our empirical approach also suggests a route forward in trying to understand the forces behind the declining rates of labor adjustments observed in the United States in particular. Essentially, our empirical approach provides a tool for assessing whether this development is due to a changing nature of firm-level shocks or due to a reduced impact of these shocks on labor reallocation. Although this question is beyond the scope of this paper, it serves as a good example of the questions that future research can answer by combining data on labor flows and estimated firm-level shocks.

References

Abowd, J. M., P. Corbel, and F. Kramarz (1999): "The Entry and Exit of Workers and the Growth of Employment: An Analysis of French Establishments," *Review of Economics and Statistics*, 81, 170–187.

Andersson, F. (2003): *Causes and Labor Market Consequences of Producer Heterogeniety,* chap. Job Flows in Swedish Manufacturing, pp. 13–52, Economic Studies 73. Uppsala University.

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

Aslund, O., and O. Nordström-Skans (2009): "How to Measure Segregation Conditional on the Distribution of Covariates," *Journal of Population Economics*, 22, 971–981.

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.

Bentolila, S., and G. Bertola (1990): "Firing Costs and Labour Demand: How Bad Is Eurosclerosis?," *Review of Economic Studies*, 57, 381–402.

Bertola, G. (2004): "A Pure Theory of Job Security and Labour Income Risk," *The Review of Economic Studies*, 71(1), 43–61.

Blanchard, O., and D. Quah (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79, 655–673.

Cahuc, P., F. Postel-Vinay, and J.-M. Robin (2006): "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, 74(2), 323–64.

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

Carlsson, M., J. Messina, and O. Nordström-Skans (2014): "Wage Adjustment and Productivity Shocks," *forthcoming Economic Journal.*

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.

D'Agostino, R. B., A. J. Belanger, and R. B. D'Agostino Jr. (1990): "A Suggestion for Using Powerful and Informative Tests of Normality," *American Statistician*, 44, 316–321.

Davis, S., and J. Haltiwanger (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107, 819–863.

Davis, S. J., R. J. Faberman, and J. Haltiwanger (2012): "Labor Market Flows in the Cross Section and Over Time," *Journal of Monetary Economics*, 59, 1–18.

Davis, S. J., J. C. Haltiwanger, and S. Schuh (1996): *Job Creation and Job Destruction.* MIT Press, MA.

Erceg, C., D. Henderson, and A. Levin (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46, 281–313.

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.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 98(1), 394–425.

Foster, L., J. Haltiwanger, and C. Syverson (2012): "The Slow Growth of New Plants: Learning and Demand," Mimeo, University of Maryland.

Franco, F., and T. Philippon (2007): "Firms and Aggregate Dynamics," *Review of Economics and Statistics*, 89, 587–600.

Fujita, S., and G. Ramey (2009): "The Cyclicality Of Separation And Job Finding Rates," *International Economic Review*, 50, 415–430.

Galí, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, 89, 249–271.

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

Heyman, F., H. Svaleryd, and J. Vlachos (2008): "Competition, Takeovers and Gender Discrimination," *forthcoming Industrial and Labor Relations Review.*

Hopenhayn, H., and R. Rogerson (1993): "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101, 915–938.

Lise, J., C. Meghir, and J.-M. Robin (2013): "Mismatch, Sorting and Wage Dynamics," *NBER Working Paper Series*, 18719.

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

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.

Michelacci, C., and D. Lopez-Salido (2007): "Technology Shocks and Job Flows," *Review of Economic Studies*, 74, 1195–1227.

Mortensen, D., and C. Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61(3), 397–415.

OECD (2014): "Indicators of Employment Protection," *Database available at: http://www.oecd.org/employment/emp/oecdindicatorsofemploymentprotection.htm.*

Pozzi, A., and F. Schivardi (2012): "Demand or Productivity: What Determines Firm Growth?," EIEF Working Papers Series No. 1211.

Ramey, V. (2011): "Identifying Government Spending Shocks: It's all in the Timing," *Quarterly Journal of Economics*, 126, 1–50.

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. (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15, 127–148.

Smeets, V., and F. Warzynski (2013): "Estimating Productivity with Multi-Product Firms, Pricing Heterogeneity and the Role of International Trade," *Journal of International Economics*, 90(2), 237–244.

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

Appendices

A Derivation of long-run restrictions

A.1 Baseline model

We use the stylized model presented in the paper to filter out shocks that permanently shift the firms' production functions and demand curves. To filter out the shocks of interest, we first note that the assumptions of the model ensure that the only shock that can affect the physical gross output Solow residual (A) is the technology shock. Since we only impose this restriction in the long run, we can allow for temporary variations in factor utilization and inventories.

Further, we use the standard result that a firm's optimal pricing rule under these conditions is to set the price, P_{jt}, as a constant markup $\sigma/(\sigma-1)$ over marginal cost, MC_{jt}. Marginal cost is, in optimum, equal to

$$MC_{jt} = A_{jt}^{-1}\left(\frac{W_{jt}}{\alpha}\right)^{\alpha} P_{st}^{F}. \tag{A1}$$

Using (A1) and that $MC_{jt} = (W_{jt}N_{jt})/(\alpha Y_{jt})$ in optimum to get

$$(W_{jt}N_{jt}/Y_{jt})W_{jt}^{-\alpha} = \alpha^{1-\alpha}A_{jt}^{-1}P_{st}^{F}. \tag{A2}$$

Thus, expression (A2) will be affected by technology and factor-price shocks but not demand shocks. It is also worth noting that any direct shocks to the firm-level wage-setting relationship (such as changes in the degree of competition over similar types of labor) will not drive this expression. Essentially, expression (A2) is a measure of unit labor cost ($W_{jt}N_{jt}/Y_{jt}$) net of wage-setting disturbances.[25] We therefore refer to the variable as wage-neutral labor cost ($WNULC_{jt}$).

Using the demand equation (2) and expression (A1), we arrive at

$$Y_{jt}W_{jt}^{\sigma\alpha} = \psi Y_t P_t^{\sigma} A_{jt}^{\sigma}\left(P_{jt}^{F}\right)^{-\sigma}\Omega_{jt}, \tag{A3}$$

where $\psi = \left(\frac{1}{\alpha}\right)^{-\sigma\alpha}\left(\frac{\sigma}{\sigma-1}\right)^{-\sigma}$. Thus, expression (A3) will be driven by shocks to technology, factor prices other than labor, and demand (apart from aggregate factors that will

[25] Note also that unit labor cost is proportional to marginal cost.

be captured by time dummies in the empirical implementation of the model). In effect, expression (A3) is demand adjusted for wage-setting disturbances. Thus, we refer to it as wage-neutral demand (WND_{jt}) in the text.

A.2 Non-constant returns to scale

Define the overall returns to scale as $\lambda = \alpha+\beta+\gamma$. Notice that under non-constant returns to scale, it is straightforward to show that the measurement of the variables in the system of equations needs to be changed to those of table A1 to retain the recursive form of the long-run impact of the structural shocks. Also note that the cost share of a factor will equal the output elasticity divided by the overall returns to scale in optimum, which we use in the empirical implementation provided in columns 2 and 3 of table 3 in the main text.

Table A1: Summary of structural system (non-constant returns)

Variables:	Measured as:
$Solow:$	$Y_{jt}\left(N_{jt}^{\alpha}K_{jt}^{\beta}M_{jt}^{\gamma}\right)^{-1}$
$WNULC:$	$\left(W_{jt}N_{jt}/Y_{jt}\right)W_{jt}^{-\frac{\alpha}{\lambda}}Y_{jt}^{\left(1-\frac{1}{\lambda}\right)}$
$WND:$	$Y_{jt}^{\left(1+\sigma\left(\frac{1}{\lambda}-1\right)\right)}W_{jt}^{\sigma\frac{\alpha}{\lambda}}$

B Data

The firm data-set we use is primarily drawn from Sweden's Industry Statistics Survey (IS) and contains annual information for the years 1990–2002 on inputs and output for all Swedish manufacturing plants with 10 employees or more and a sample of smaller plants. Here we focus on firms that have at least 10 employees and that we observe in a spell with at least five observations (the minimum panel dimension required for the SVAR to pass diagnostic tests).

Our measure of real output, Y_{jt}, is the value of total sales taken from the IS deflated by a firm-specific producer-price index. The firm-specific price index is a chained index with Paasche-type links that combines plant-specific unit values and detailed disaggregated producer-price indices (either at the goods level, when available, or at the most disaggregated sectoral level available). Note that when a plant-specific unit-value price is missing (e.g., when the firm introduces a new good), Statistics Sweden tries to find a price index for similar goods defined at the minimal level of aggregation (starting at four-digit goods-code level). The disaggregated sectoral producer-price indices are only used when a plausible goods-price index is unavailable.

To compute the input index (Δz_{jt}), which is necessary for the computation of the Solow residual (Δa_{jt}), real intermediate inputs (M_{jt}) are measured as the sum of costs for intermediate goods and services (including energy) collected from the IS deflated by a three-digit (SNI92/NACE) producer-price index collected by Statistics Sweden. The real capital stock (K_{jt}) is computed using a variation of the perpetual inventory method. In the first step, we calculate the forward recursion

$$K_{jt} = \max((1-\delta)K_{jt-1} + I_{jt}, BookValue_{jt}), \tag{B1}$$

where δ is sector-specific depreciation rate (two-digit SNI92/NACE) and is computed as an asset-share-weighted average between the machinery and buildings depreciation rates (collected from Melander (2009), table 2); I_{jt} is real net investments in fixed tangible assets (computed using a two-digit SNI92/NACE sector-specific investment deflator collected from Statistics Sweden); and $BookValue_{jt}$ is the book value of fixed tangible assets taken from the Firm Statistics data base maintained by Statistics Sweden, deflated using the same deflator as for investment. Moreover, K_{j0} is set to zero if the initial book value is missing in the data. Since, for tax reasons. the firms want to keep the book values low, we use the book values as a lower bound of the capital stock. In a second step, we then calculate the backward recursion

$$K_{jt-1} = \frac{K_{jt} - I_{jt}}{(1-\delta)}, \tag{B2}$$

where the ending point of the first recursion, K_{jT}, is used as the starting point for the second backward recursion. This is done to maximize the quality of the capital-stock series given that we lack a perfectly reliable starting point and the time dimension is small. The labor input (i.e., number of employees) is taken from the IS. To compute the cost shares, we also need a measure of the firms' labor cost, which is defined as total labor cost (including payroll taxes) in the IS.

When computing Δa_{jt}, we take an approach akin to the strategy outlined by Basu, Fernald, and Shapiro (2001). Thus, the C_J (i.e., the output elasticities) are treated as constants. Second, the cost shares are estimated as the time average of the cost shares for the two-digit industry to which the firm belongs (SNI92/NACE).[26] Third, to calculate the cost shares, we take total costs as approximately equal to total revenues.[27] The cost share of capital is then given by one minus the sum of the cost shares for all other factors.

Since 1996, Statistics Sweden has imputed the allocation of production across different plants within multi-plant firms. For this reason, we have explored various cuts of the data either focusing on single-plant firms throughout or use multi-plant firms before 1996 but only single-plant firms thereafter. The results are shown in Table D2 in appendix D and discussed in the robustness section of the paper.

When computing $\Delta wnulc_{jt}$ and Δwnd_{jt}, we use C_N as the estimate of α and the measure of the firms' labor costs together with the measure of real output and labor input (all discussed above). Also, when computing Δwnd_{jt}, we set σ equal to our estimate of 3.306. Finally, we remove 2 percent of the observations in each tail for each of the distributions of Δa_{jt}, $\Delta wnulc_{jt}$, Δwnd_{jt}, and Δy_{jt}. This has little effect on estimated coefficients, but it ensures that the SVAR passes diagnostic tests. We finally require the firm to be observed in spells of at least five years (because we are interested in the within-firm dynamics when estimating the SVAR).

In the end, we construct series for Δa_{jt}, $\Delta wnulc_{jt}$, Δwnd_{jt}, and Δy_{jt} for $7,940$ ongoing firms (observed at least during five consecutive years), over the $1991-2002$ period. All in all, this amounts to $70,077$ firm/year observations. Removing extreme tail events reduces the sample to $6,137$ firms and $53,379$ firm/year observations (in the specification with output growth as the fourth variable). For these firms we can compute the structural shocks for $41,105$ firm/years (due to lags in the model). Finally, we can match on labor flows from RAMS for $6,125$ firms and $40,451$ firm/year observations. Note that the procedure outlined above implies that changing the fourth variable in the VAR introduces small changes in the sample size.

[26] In the calculation we drop firm/year observations in which the (residual) capital share is below -25 percent of sales. This procedure generates reasonable aggregate cost shares, and ensures that the cost shares in all industries are positive.

[27] Using the data underlying Carlsson (2003), and relying on a no-arbitrage condition from neoclassical investment theory (also taking the tax system into account) to calculate the user cost of capital, 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 the approximation of cost shares by revenue shares.

C The SVAR

C.1 Identification

The model outlined in the paper and presented in detail in appendix A provides a set of three equations that depend on the three structural shocks (i.e., demand, technology, and intermediate inputs). The left-hand-side variables in these equations can all be constructed from our firm-level data, and the model motivates a recursive sequence of long-run restrictions regarding the impact of the structural shocks on these variables. To extract the shocks of interest from the system, we estimate a VAR and proceed along the lines of Blanchard and Quah (1989).

Since we are interested in how other variables (such as output, prices, and employment) respond to structural shocks, we start by including these other variables as fourth variables in the system, allowing each to have a long-run effect on itself but not on the other variables in the system. These variables will thus also soak up all remaining transitory dynamics. In practice, we rotate across these variables while keeping the core system of the first three equations intact as in Ramey (2011). Parts of our analysis rely on extracting the technology and demand shocks from the system. In these exercises we use output as the fourth variable, but we also present several robustness checks showing that the results are insensitive to this choice. The VAR system, a fully interacted dynamic system of the variables, can, under standard regularity conditions, be written in a vector moving average (MA) form. Using lowercase letters for logarithms and denoting the fourth variable by θ, the MA representation of the system follows:[28]

$$\begin{bmatrix} \Delta a_{jt} \\ \Delta wnulc_{jt} \\ \Delta wnd_{jt} \\ \Delta \theta_{jt} \end{bmatrix} = \begin{bmatrix} C_{11}(L) & C_{12}(L) & C_{13}(L) & C_{14}(L) \\ C_{21}(L) & C_{22}(L) & C_{23}(L) & C_{24}(L) \\ C_{31}(L) & C_{32}(L) & C_{33}(L) & C_{34}(L) \\ C_{41}(L) & C_{42}(L) & C_{43}(L) & C_{44}(L) \end{bmatrix} \begin{bmatrix} \eta^a_{jt} \\ \eta^f_{jt} \\ \eta^\omega_{jt} \\ \eta^\theta_{jt} \end{bmatrix}. \tag{C1}$$

We assume that the shocks $([\eta^a_{jt}, \eta^f_{jt}, \eta^\omega_{jt}, \eta^\theta_{jt}])$ are structural innovations and hence mutually orthogonal and serially uncorrelated. Because the shock associated with the fourth variable lacks a theoretical interpretation, we refer to it as the "residual" shock in what follows. The terms $C_{rc}(L)$ are polynomials in the lag operator, L, with coefficients $c_{rc}(k)L^k$ at each lag k. The shocks are orthogonal, and using a standard normalization we get $E\boldsymbol{\eta}_t\acute{}\boldsymbol{\eta}_t = \mathbf{I}_t$,

[28] Note that the assumed functional form of the processes for demand and technology shifters specified in equations (3) and (4) directly leads to equation (C1).

where $\boldsymbol{\eta}_t = [\eta_{jt}^a, \eta_{jt}^f, \eta_{jt}^\omega, \eta_{jt}^\theta]'$.

Following standard practice, we denote the elements of the matrix of *long-run* multipliers corresponding to (C1) as $C_{rc}(1)$. Relying on the model outlined above, we know that the technology shock, η_{jt}^a, is the only shock with a long-run impact on a_{jt}, so $C_{12}(1) = C_{13}(1) = C_{14}(1) = 0$ in the matrix of long-run multipliers.[29] Similarly, only the technology and the factor-price shocks have a long-run effect on $wnulc_{jt}$, so $C_{23}(1) = C_{24}(1) = 0$. Finally, since the residual shock has no long-run effects on wage-neutral demand, it follows that $C_{34}(1) = 0$.

Given these assumptions, we can recover the time series of the firm's structural shocks $\boldsymbol{\eta}_{jt}$ from an estimate of the VAR(p) formulation of the system (C1), i.e., from

$$\Delta \mathbf{x}_t = \textstyle\sum_1^P \mathbf{A}_p \Delta \mathbf{x}_{t-p} + \mathbf{e}_t, \tag{C2}$$

where $\mathbf{A}_p$ denotes the matrices with coefficients, $\Delta \mathbf{x}_t = [\Delta a_{jt}, \Delta wnulc_{jt}, \Delta wnd_{jt}, \Delta \theta_{jt}]'$, $\mathbf{e}_t$ is a vector of reduced-form disturbances, and we have suppressed constants to save on notation.

Under standard regularity conditions, there exists a VAR representation of the MA representation (C1) of the form

$$\mathbf{x}_t = \mathbf{A}(L) L \mathbf{x}_t + \mathbf{e}_t, \tag{C3}$$

where $\mathbf{x}_t = [\Delta a_{it},\ \Delta wnulc_{jt},\ \Delta wnd_{jt}, \Delta \theta_{jt}]$, $A_{rc}(L) = \sum_{k=0}^{\infty} a_{rc}(k) L^k$ and $\mathbf{e}_t$ is a vector of reduced-form errors. Since the errors in the VAR, $\mathbf{e}_t$, are one-step-ahead forecast errors, we will have that

$$\mathbf{e}_t = \mathbf{c}(0) \boldsymbol{\eta}_t, \tag{C4}$$

where $\mathbf{c}(0)$ is the matrix of $c_{rc}(0)$ coefficients from the MA representation and $\boldsymbol{\eta}_t = [\eta_{jt}^a, \eta_{jt}^f, \eta_{jt}^\omega, \eta_{jt}^\theta]'$. Thus, if the 16 coefficients in $\mathbf{c}(0)$ were known, we could recover $\boldsymbol{\eta}_t$.

In practice, we first use that $E\boldsymbol{\eta}_t \boldsymbol{\eta}_t' = \mathbf{I}_t$ together with an estimate of $\Omega = E\mathbf{e}_t'\mathbf{e}_t$ from our estimates of equation (C3) to obtain 10 restrictions. In addition, we impose the 6 long-run restrictions. Finally, rewriting equation (C3), we can obtain the MA form by

[29] That is, the coefficients $c_{12}(k)$ are such that $\sum_{k=0}^{\infty} c_{12}(k) = 0$, and similarly for the coefficients $c_{13}(k)$ and $c_{14}(k)$.

using equation (C4) in terms of coefficients from equation (C3) and the $\mathbf{c}(0)$ coefficients as

$$\mathbf{x}_t = [I - \mathbf{A}(L)L]^{-1}\mathbf{c}(0)\boldsymbol{\eta}_t. \tag{C5}$$

Then, our 6 long-run restrictions imply an equal number of restrictions on the matrix $[I - \mathbf{A}(L)L]^{-1}\mathbf{c}(0)$, that together with an estimate of (C3) yields 6 additional restrictions on $\mathbf{c}(0)$. Jointly, these 16 restrictions provide an estimate of the $\mathbf{c}(0)$ matrix, $\hat{\mathbf{c}}(0)$, and using these we can solve for the structural shocks using equation (C4):

$$\hat{\mathbf{c}}(0)^{-1}\hat{\mathbf{e}}_t = \hat{\boldsymbol{\eta}}_t. \tag{C6}$$

When deriving results in term of elasticities, and to obtain an estimate of the standard deviation of the structural shocks, we use a re-normalized $\hat{\mathbf{c}}(0)$ where each element is divided by its column diagonal element.

C.2 Impulse responses, variance decompositions, and tests

Relying on the Arellano and Bond (1991) autocorrelation test of the differenced residual, two lags in the VAR are enough to remove any autocorrelation in the residuals in all four equations. Here we rely on the two-step Arellano and Bond (1991) difference estimator, using the second to the fourth lag levels as instruments. It is worth noting, though, that the parameter estimates are not sensitive to the actual choice of where to cut the instrument set. The results are also insensitive to the inclusion of more lags as instruments. As an additional precaution, we collapse the instrument set to avoid overfitting. That is, we impose the restriction that the relationships in the "first stage" are the same across all time periods (see Roodman, 2006, for a discussion). For all specifications, the Hansen test of the overidentifying restrictions cannot reject the null that the model is correctly specified and the instruments are valid.

Figure C1 shows the impulse responses of each of the variables in the baseline VAR in levels to each of the structural shocks. Since the estimated system converges fairly rapidly, we only plot the initial five periods. All impulse responses are precisely estimated as indicated by the tight (95 percent) confidence bands based on $1,000$ bootstrap replications. The high level of precision is not surprising, given that we estimate the impulse responses on a much larger sample than is common in macroeconomic applications.

Unfortunately, we have not been able to find any statistical tests of stationarity that

Figure C1: Impulse responses

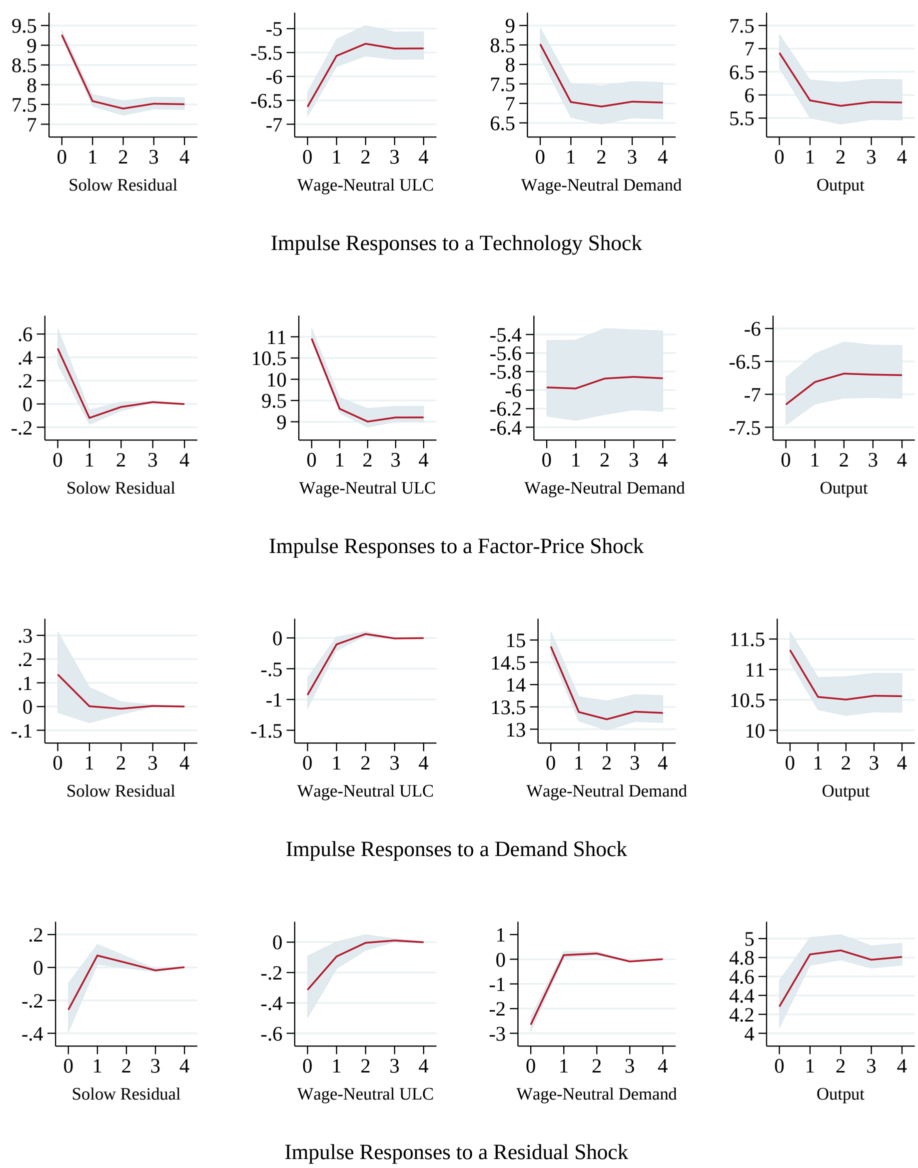

Note: Impulse responses of the Solow residual, wage neutral unit labor costs (wage neutral ULC), wage-neutral demand and output in the baseline VAR in percentage points. Each line depicts the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95 percent confidence intervals calculated from 1,000 replications.

Figure C2: Variance decompositions

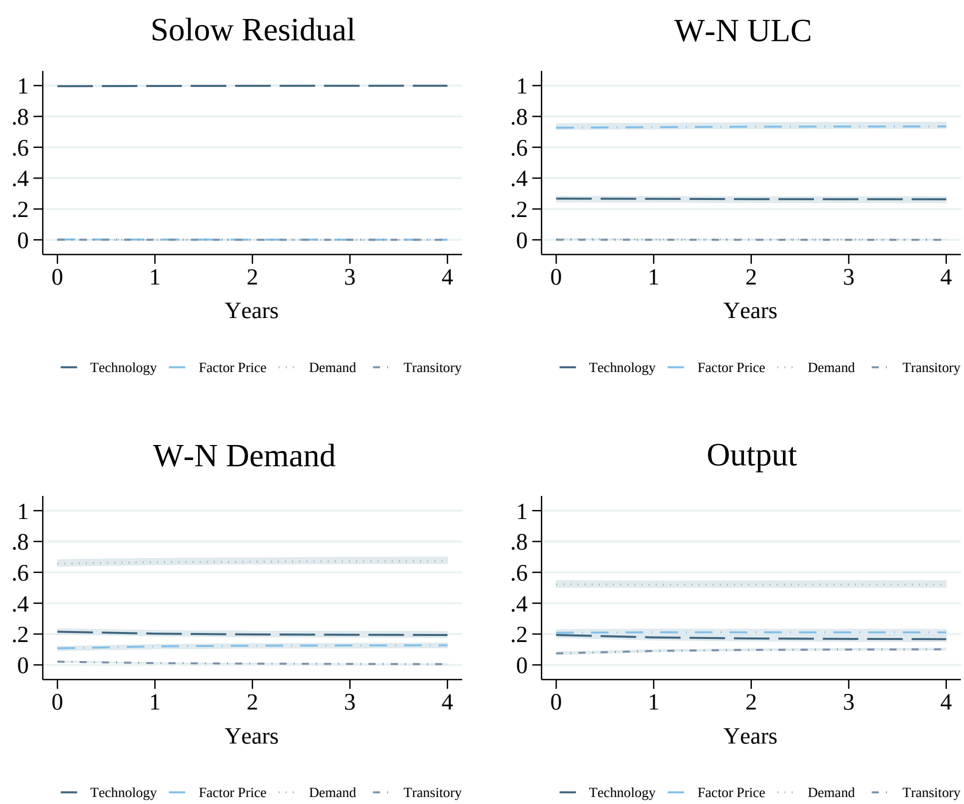

Note: Forecast-error variance decompositions of the VAR in levels. W-N Demand denotes wage-neutral demand. W-N ULC denotes wage-neutral unit labor costs.The left-most panel shows the percentage of the forecast-error variance in the Solow residual that can be explained by each structural shock at different horizons. Each line depicts the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95 percent confidence intervals calculated from 1,000 replications.

Figure C3: The distribution of demand and technology Shocks

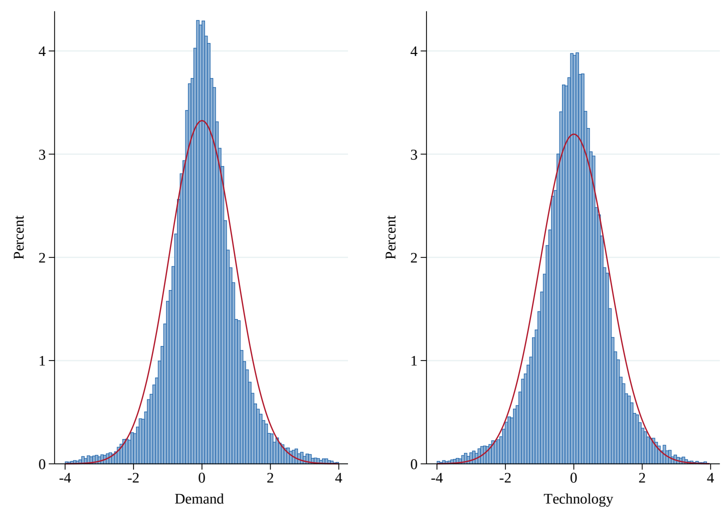

Note: Histograms of demand and technology shocks. Dashed lines depict a normal distribution.

are suitable for a setting with a short but wide panel. However, it should be clear from figure C1 that this issue is of little importance in the current setting. Importantly, the figure is expressed in log-levels, and the flat, non-zero-end-segments in the responses imply that shocks do have permanent effects on the levels of the series (i.e., the levels are I(1)) and that the differenced series are stationary (I(0)).

The first row of figure C1 traces out the impulse responses of the Solow residual, the *wnulc*, the *wnd*, and output to a 1 sd technology shock, η^a_{jt}. Technology shocks have a positive permanent effect on the Solow residual: a "normal" (i.e., 1 sd) shock increases the Solow residual slightly less than 10 percent in the long run. The estimated model does not impose any restrictions on how technology shocks affect *wnulc* and *wnd*. However, the results do concur with predictions from expression (A2) in the sense that *wnulc* falls permanently in response to the (permanent) technology shock. Similarly, we find that a permanent technology shock raises *wnd*, as predicted from expression (A3).

The second row in figure C1 reports the impulse responses to a 1 sd permanent factor-price shock. A "normal" factor-price shock increases *wnulc* and lowers *wnd* permanently (theoretically working through marginal cost, price setting, and demand). The latter result is, again, an unconstrained result in line with predictions from expression (A2). By the

Figure C4: Non-linear responses to a technology shock

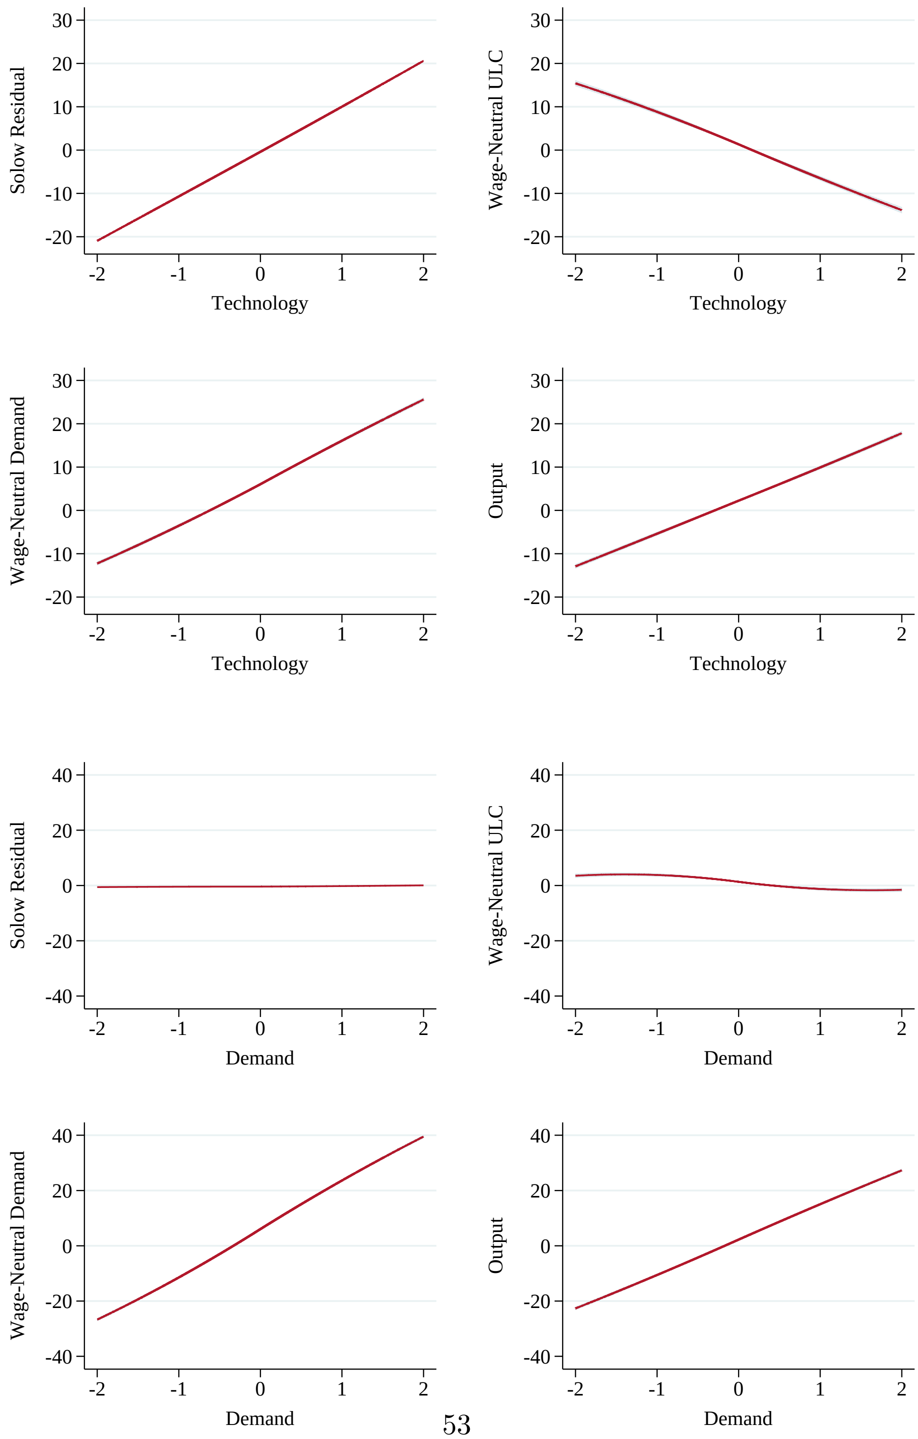

Note: Contemporaneous response of variables included in the baseline VAR in percentage units as a (non-linear) function of an x sd of the technology and demand shocks. Shaded areas depict the 95 percent confidence intervals.

same logic, output also falls permanently in response to a factor-price shock. The Solow residual is affected in the very short run by factor-price shocks but converges to the long-run restriction fairly rapidly.

The impulse responses to a permanent demand shock are shown in the third row of figure C1. In this case, *wnd* is permanently increased in response to a permanent demand shock. In the short run, demand shocks increase the Solow residual and reduce *wnulc*. As expected, a demand shock also has permanently positive effects on output. A "normal" demand shock increases it by about 10 percent in the long run. For completeness, figure C1 also reports the responses to the residual shock in the last row. A "normal" residual shock raises output permanently by slightly more than 5 percent.

Figure C2 presents forecast error variance decompositions for each of the variables in the VAR in levels, decomposing the movements of the three variables. Again, bootstrapped confidence bands are extremely tight. Quantitatively, the Solow residual is solely driven by technology shocks on all horizons. The *wnulc* is mostly driven by factor-price shocks (75 percent of the variation) and partly by technology shocks (25 percent). Demand shocks explain about 65 percent of the movements in *wnd*, whereas factor-price shocks explain about 20 percent. We also see in figure C2 that there is a role for technology shocks in explaining wage-neutral demand movements, accounting for about 15 percent. For output, we see that about 55 percent of the variation is driven by demand shocks, the rest being explained by factor-price shocks (about 20 percent), technology (about 15 percent), and the transitory shock (about 10 percent).

Overall though, we find the residual shock to be of little importance. Given that we include time dummies in the VAR, this finding is in line with the results of Franco and Philippon (2007), which finds that transitory shocks are not very important on the firm level but account for most of the volatility of aggregates because they are correlated across firms.

Figure C3 shows the distributions for extracted innovations to technology and demand. As the two panels of the figure show, neither the demand nor the technology shock distributions are particularly skewed (skewness coefficients of -0.02 and -0.14, respectively), whereas both are leptokurtic (kurtosis coefficients of 5.85 and 4.25). This is also clearly visible in the graphs where the dashed line depicts a normal distribution, and a standard skewness/kurtosis test (D'Agostino, Belanger, and D'Agostino, 1990) rejects the null of normality for both distributions (p-value of 0.00 in both cases). The shock distributions depicted in figure C3 are normalized to have a unit standard deviation. When re-normalizing the system (see appendix A), we find that the standard deviation of the demand shock is about 35 percent larger than the technology shock (standard deviations of 16.02 and 11.86 percentage points, respectively).

A maintained assumption in the analysis is that the baseline VAR is linear in the structural shocks. In figure C4 we plot the predicted contemporaneous responses of the variables included in the VAR as (possibly non-linear) functions of structural shocks (allowing for a separate second-order polynomial above and below zero). As the graphs show, the results do support the maintained linearity assumption.

D Appendix tables

Table D1: Contemporaneous and long-run effect on log employment - different values of σ and sectoral dynamics

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline ($\sigma = 3.3$)	σ by sector	σ by sector	$\sigma = 1.1$	$\sigma = 10$	Sectoral Dynamics
			SHORT RUN			
η_a	0.153	0.192	0.147	0.207	0.259	0.126
	(0.159)	(0.161)	(0.162)	(0.158)	(0.152)	(0.162)
η_ω	5.986**	5.693**	5.520**	6.591**	4.060**	5.506**
	(0.233)	(0.221)	(0.222)	(0.240)	(0.198)	(0.225)
Observations	40,451	40,214	39,580	41,046	39,207	39,580
Firms	6,125	6,102	5,997	6,189	5,998	5,997
			LONG RUN			
η_a	0.504*	0.599**	0.510*	0.512*	0.643**	0.490*
	(0.214)	(0.214)	(0.217)	(0.208)	(0.220)	(0.221)
η_ω	6.357**	5.996**	5.811**	7.009**	4.267**	5.737**
	(0.310)	(0.302)	(0.306)	(0.312)	0.291	(0.302)
Observations	34,414	34,198	33,667	34,612	33,291	33,667
Firms	6,116	6,094	5,989	6,181	5,991	5,989
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Sectoral Sigma	No	Yes	Yes	No	No	Yes
Pooled Dynamics	Yes	Yes	Yes	Yes	Yes	No
Sector by Time FE	No	No	Yes	No	No	Yes
sd. η_a	10.06	10.03	9.94	10.16	9.98	9.88
sd. η_ω	16.18	17.09	16.98	13.87	27.23	16.80

Note: Columns (2) and (3) allow for a sectoral σ. In column (4) we reestimate the entire SVAR for each two-digit industry. This explains the small drop in the number of observations. The number of firms in some industries is too small to get sectoral estimates. Columns (5) and (6) impose large variation in values of σ. All estimates are the effect of a 1 sd shock. Robust standard errors in parentheses. Long-run estimates are the sum of the contemporaneous impact and one lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

Table D2: Contemporaneous and long-run effect on log employment - sample variations

	(1)	(2)	(3)	(4)
Sample	Baseline	Single Plant Always	Single Plant After 1996	$\leq \pm 2$ Sd. Shocks
		SHORT RUN		
η_a	0.153	0.421**	0.312*	0.040
	(0.159)	(0.158)	(0.151)	(0.164)
η_ω	5.986**	5.500**	6.244**	6.317**
	(0.233)	(0.236)	(0.238)	(0.205)
Observations	40,451	20,877	30,234	36,072
Firms	6,125	3,246	5,259	6,111
		LONG RUN		
η_a	0.504*	0.534*	0.669**	0.336
	(0.214)	(0.233)	(0.215)	(0.234)
η_ω	6.357**	5.715**	6.657**	6.397**
	(0.310)	(0.309)	(0.326)	(0.294)
Observations	34,414	17,638	25,040	30,693
Firms	6,116	3,246	5,250	6,066
sd. η_a	10.06	9.13	9.41	10.06
sd. η_ω	16.18	15.07	14.79	16.18

Note: Column (2) restricts the sample to single-plant firms; column (3) includes a mixed sample with multi-plant firms until 1996, but not thereafter; column (4) shows results for a trimmed sample where we focus on shocks in the Lester range of ± 2 standard deviations. Estimates are the effects of a 1 sd shock. Robust standard errors in parentheses. Regression includes firm fixed effects and time dummies. Long-run estimates are the sum of the contemporaneous impact and one lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

Table D3: Contemporaneous and long-run effect on log employment - varying the fourth variable in the VAR

	(1)	(2)	(3)	(4)
Fourth Variable of VAR:	Output	Sales per Worker	Employment (IS)	Employment (RAMS)
			SHORT RUN	
η_a	0.153	0.524**	0.263	0.499**
	(0.159)	(0.154)	(0.143)	(0.061)
η_ω	5.986**	5.840**	5.261**	6.986**
	(0.233)	(0.234)	(0.212)	(0.086)
Observations	$41,105$	$40,284$	$38,213$	$37,234$
Firms	$6,125$	$6,113$	$5,879$	$5,703$
			LONG RUN	
η_a	0.504*	0.812**	0.644**	0.643**
	(0.214)	(0.218)	(0.209)	(0.097)
η_ω	6.357**	6.134**	5.477**	7.514**
	(0.310)	(0.317)	(0.266)	(0.121)
Observations	$34,414$	$34,260$	$32,407$	$31,531$
Firms	$6,116$	$6,102$	$5,871$	$5,703$
sd. η_a	10.06	9.980	9.971	9.964
sd. η_ω	16.18	16.39	15.35	15.13

Note: Column (2) derives shocks from a VAR in which the fourth variable is sales per worker; in column (3) the fourth variable is annual employment measured in the IS data set; in column (4) the fourth variable is end-of-the-year employment measured from the RAMS data-set.All estimates are the effect of a 1 sd shock. Robust standard errors in parentheses. All regressions include time dummies and firm fixed effects. Long-run estimates are the sum of the contemporaneous impact and one lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

Table D4: Contemporaneous and long-run effect on net employment growth - two periods models.

	(1) One Period – Baseline	(2) Two Period – Without Exits	(3) Two Period – With Exits
		SHORT RUN	
η_a	0.115	0.328*	0.278
	(0.119)	(0.153)	(0.164)
η_ω	5.609**	5.431**	5.749**
	(0.173)	(0.375)	(0.376)
Observations	40,451	39,822	40,238
Firms	6,125	6,114	6,121
		LONG RUN	
η_a	0.412*	0.420	0.380
	(0.163)	(0.350)	(0.368)
η_ω	6.009**	4.112**	4.696**
	(0.228)	(0.391)	(0.422)
Observations	34,414	33,830	34,243
Firms	6,116	6,099	6,110

Note: The dependent variable in column (1) is the employment change between t and $t-1$ divided by the average employment in the two years. In columns (2) and (3) the dependent variable is defined as the employment change between $t+1$ and $t-1$ divided by the average employment in the two years. Columns (1) and (2) exclude frims that exit the sample in the calculation of the flows, and column (3) includes them. The reported coefficients are the effect of 1 sd shock. Robust standard errors in parentheses. Regression includes firm fixed effects and time dummies. Long-run estimates are the sum of the contemporaneous impact and one lag. ** and * denote statistical significance at the 1 and 5 levels, respectively.

Publication series published by IFAU – latest issues

Rapporter/Reports

2014:1 Assadi Anahita ”En profilfråga: Hur använder arbetsförmedlare bedömningsstödet?”

2014:2 Eliason Marcus ”Uppsägningar och alkoholrelaterad sjuklighet och dödlighet”

2014:3 Adman Per ”Försummas gymnasieskolans demokratiuppdrag? En kvalitativ textanalys av 2009 års svenska gymnasiereform”

2014:4 Stenberg Anders and Olle Westerlund “Utbildning vid arbetslöshet: en jämförande studie av yrkesinriktad och teoretisk utbildning på lång sikt”

2014:5 van den Berg Gerard J., Lene Back Kjærsgaard and Michael Rosholm “Betydelsen av möten mellan arbetslösa och förmedlare”

2014:6 Mörk Eva, Anna Sjögren and Helena Svaleryd ”Blir barn sjuka när föräldrarna blir arbetslösa?”

2014:7 Johansson Per, Arizo Karimi and J. Peter Nilsson ”Könsskillnader i hur sjukfrånvaro påverkas av omgivningen”

2014:8 Forslund Anders, Lena Hensvik, Oskar Nordström Skans, Alexander Westerberg and Tove Eliasson ”Avtalslöner, löner och sysselsättning”

2014:9 Engdahl Mattias ”Medborgarskap, arbetsmarknaden och familjebildning”

2014:10 Hallberg Daniel, Per Johansson and Malin Josephson ”Hälsoeffekter av tidigarelagd pensionering”

2014:11 Karbownik Krzysztof and Sara Martinson ”Svenska högstadie- och gymnasielärares rörlighet på arbetsmarknaden”

2014:12 Hägglund Pathric, Per Johansson and Lisa Laun ”Insatserna inom rehabiliteringsgarantin och deras effekter på hälsa och sjukfrånvaro”

2014:13 Regnér Johan ”Effekter av yrkesinriktad arbetsmarknadsutbildning för deltagare med funktionsnedsättning, 1999–2006”

2014:14 Assadi Anahita and Martin Lundin ”Enhetlighet och träffsäkerhet i arbetsmarknadspolitiken: Hur använder arbetsförmedlare statistisk profilering i mötet med den arbetssökande?”

2014:15 Edmark Karin, Markus Frölich and Verena Wondratschek ”Hur har 1990-talets skolvals-reformer påverkat elever med olika familjebakgrund?”

2014:16 Karimi Arizo “Sen familjebildning, täta födelseintervall och kvinnors inkomster”

2014:17 Eliasson Tove “Bankanställdas ursprungsland och egenföretagande bland utrikesfödda”

2014:18 Ingmanson Staffan ”Fri rörlighet inom den högre utbildningen och tillgång till svenska studiemedel”

2014:19 Andersson Elvira, Petter Lundborg and Johan Vikström ”Arbete, löneutbetalningar och mortalitet”

2014:20 Sibbmark Kristina ”Arbetsmarknadspolitisk översikt 2013”

2014:21 Nordlund Madelene and Mattias Strandh ”Selektivitet och jobbchanser bland arbetslösa”

2014:22 Angelov Nikolay and Marcus Eliason ”Vilka arbetssökande kodas som funktionshindrade av Arbetsförmedlingen?”

2014:23 Angelov Nikolay and Marcus Eliason ”Friställd och funktionsnedsatt”

2014:24 Angelov Nikolay and Marcus Eliason ”Lönebidrag och skyddat arbete: en utvärdering av särskilda insatser för sökande med funktionshinder”

Working papers

2014:1 Vikström Johan "IPW estimation and related estimators for evaluation of active labor market policies in a dynamic setting"

2014:2 Adman Per "Who cares about the democratic mandate of education? A text analysis of the Swedish secondary education reform of 2009"

2014:3 Stenberg Anders and Olle Westerlund "The long-term earnings consequences of general vs. specific training of the unemployed"

2014:4 Boye Katarina "Can you stay at home today? The relationship between economic dependence, parents' occupation and care leave for sick children"

2014:5 Bergemann Annette and Gerard J. van den Berg "From giving birth to paid labor: the effects of adult education for prime-aged mothers"

2014:6 van den Berg Gerard J., Lene Kjærsgaard and Michael Rosholm "To meet or not to meet, that is the question – short-run effects of high-frequency meetings with case workers"

2014:7 Avdic Daniel, Petter Lundborg and Johan Vikström "Learning-by-doing in a highly skilled profession when stakes are high: evidence from advanced cancer surgery"

2014:8 Mörk Eva, Anna Sjögren and Helena Svaleryd "Parental unemployment and child health"

2014:9 Johansson Per, Arizo Karimi and J. Peter Nilsson "Gender differences in shirking: monitoring or social preferences? Evidence from a field experiment"

2014:10 Eliasson Tove and Oskar Nordström Skans "Negotiated wage increases and the labor market outcomes of low-wage workers: evidence from the Swedish public sector"

2014:11 Engdahl Mattias "Naturalizations and the economic and social integration of immigrants"

2014:12 Hallberg Daniel, Per Johansson and Malin Josephson "Early retirement and post-retirement health"

2014:13 Karbownik Krzysztof "The determinants of teacher mobility in Sweden"

2014:14 Karbownik Krzysztof "Job mobility among high-skilled and low-skilled teachers"

2014:15 Karbownik Krzysztof "Do changes in student quality affect teacher mobility? Evidence from an admission reform"

2014:16 Edmark Karin, Markus Frölich and Verena Wondratschek "Sweden's school choice reform and equality of opportunity"

2014:17 Karimi Arizo "Effects of the timing of births on women's earnings – evidence from a natural experiment"

2014:18 Karimi Arizo "The spacing of births and women's subsequent earnings – evidence from a natural experiment"

2014:19 Eliasson Tove "Immigrant entrepreneurship and the origin of bankers"

2014:20 Johansson Per, Lisa Laun and Mårten Palme "Pathways to retirement and the role of financial incentives in Sweden"

2014:21 Andersson Elvira, Petter Lundborg and Johan Vikström "Income receipt and mortality – evidence from Swedish public sector employees"

2014:22 Felfe Christina and Rafael Lalive "Does early child care help or hurt children's development?"

2014:23 Nordlund Madelene and Mattias Strandh "The relation between economic and non-economic incentives to work and employment chances among the unemployed"

2014:24 Mellander Erik "Transparency of human resource policy"

2014:25 Angelov Nikolay and Marcus Eliason "Factors associated with occupational disability classification"

2014:26 Angelov Nikolay and Marcus Eliason "The differential earnings and income effects of involuntary job loss on workers with disabilities"

2014:27 Angelov Nikolay and Marcus Eliason “The effects of targeted labour market programs for job seekers with occupational disabilities”

2014:28 Carlsson Mikael, Julián Messina and Oskar Nordström Skans “Firm-level shocks and labor adjustments”

Dissertation series

2014:1 Avdic Daniel “Microeconometric analyses of individual behaviour in public welfare systems”

2014:2 Karimi Arizo “Impacts of policies, peers and parenthood on labor market outcomes”

2014:3 Eliasson Tove “Empirical essays on wage setting and immigrant labor market opportunities”

2014:4 Nilsson Martin “Essays on health shocks and social insurance”

2014:5 Pingel Ronnie “Some aspects of propensity score-based estimators for causal inference”