Firm-Level Shocks and Labor Adjustments*

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

October 30, 2013

Preliminary Version

Abstract

This paper analyzes the relationship between firm-level productivity and labor adjustments with the purpose of understanding the coexistence of a sustainable cross-firm dispersion in technological efficiency and substantial idiosyncratic labor flows. Using a unique Swedish data set merging information about firms' inputs, outputs and prices to a linked employer-employee data set, we analyse how firms adjust their labor in response to permanent shifts in their idiosyncratic production functions and demand curves. We show that permanent shocks to firm-level demand is the main driving force behind both job and worker reallocation. Furthermore, we show several pieces of evidence suggesting that the adjustment in response to permanent shocks is a relatively unconstrained process. Notably, most labor adjustment takes place within a year and firms adjust through increased separations even when they could have adjusted through reduced hires. Jointly, these results suggest that the technology dispersion is maintained in equilibrium because labor flows are driven by demand differences rather than differences in technology, whereas labor market rigidities appear to be a less important factor.

Keywords:
JEL classifications:

*We would like to thank Richard Freeman, Nils Gottfries, Francis Kramarz, Claudio Michelacci, Frédèric Warzynski, Nikolaj Angelov, and seminar participants at the Bank of Spain, Carlos III, UCLS Advisory Board Workshop 2012, IAE-MOVE Workshop on Industry and Labor Market Dynamics, Barcelona, 2012, CAP data workshop in Arhus 2013, The Nordic Summer Institute in Labour Economics, Bergen, 2012, IFAU, Sveriges Riksbank, Linnaeus University, Gothenburg University and the Greater Stocholm Macro Group. We are also grateful to Karl Harmenberg, Mats Levander and Matias Quiroz for excellent research assistance. The data used in this paper are confidential but the authors' access is not exclusive. 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.

†Uppsala University, Uppsala Center for Labor Studies and Sveriges Riksbank. E-mail: mikael.carlsson@nek.uu.se

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

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

1 Introduction

Recent research has shown that firms with very different levels of technological efficiency coexist simulatneously at the same market. Less is known as to why these differences are sustainable. One reason may be frictions on the labor adjustment side (see e.g. Mortensen, 2003). Poor-performing firms may be prevented from adjusting their labor force due to institutional rigidities and workers may choose to remain in their jobs because the individual costs of mobility are higher than the benefits. Still, this explanation seems at odds with the enormous magnitudes of idiosyncratic, firms specific, labor flows documented by Davis, Haltiwanger, and Schuh (1996)) and followers (see e.g. Davis, Faberrman and Haltiwanger, 2012 and references therein). These studies uniformly show that the bulk of firm-level labor adjustments is truly idiosyncratic: Firms operating in the same sector and area shrink and grow, side-by-side. It thus seems reasonable to expect that labor should flow in the direction of the more efficient firms. In this paper we derive novel evidence on the relationships between firm performance and labor flows in an attempt to reconcile the survival of relatively inefficient firms in a market characterized by immense labor flows and provide an account of the extent to which separation rigidities prevent firms from reducing they labor when needed.

To this end, we measure permanent shifts in firm-level physical productivity and product demand (for details, see below) and relate these shifts to idiosyncratic job and worker flows. The aim is to present direct evidence on how the magnitudes and signs of permanent changes in determinants of firm performance affect firms' labor adjustments along different possible adjustment margins. We do this by combining a unique Swedish data base which links measures of firm level input, output and prices to individual worker-flow data. The data allows us to analyze how these shocks affect firms' adjustments of net employment (jobs) through hires and separations. To corroborate the interpretation of the shocks, and to shed light on the co-movement of different firm-level adjustment margins, we relate the shocks to idiosyncratic output prices and physical output.

There exits a large body of theoretical research (B & B etcetera) on the relationship between firm-level revenue productivity (marginal revenue productivity of labor) and labor adjustments. However, the recent empirical literature, starting with Foster, Haltiwanger, and Syverson (2008), has drawn the attention to two distinct shocks which may shift revenue productivity despite having different structural interpretations and with different relationships to firm-level outcomes. In line with the existing conventions we may define these as *technology shocks* shifting the firm-level physical production function (i.e. reducing the needs for inputs) and *demand shocks* shifting the firm-level demand curve (i.e. increasing the ability to sell at a given price). Importantly, these shocks are defined according to their effects on firm-level optimization, not according to their origin. Findings in this empirical literature include Foster, Haltiwanger, and Syverson (2008) who show that firm closures primarily are driven by changes in idiosyncratic demand and only to a lesser extent by changes in idiosyncratic physical productivity. Similarly, recent evidence by Foster, Haltiwanger, and Syverson (2012) suggests that the growth of new firms is due to a shrinking product demand gap relative to incumbent firms.[1]

Since we, in this paper, are interesting in understanding the sustained co-existence of firms with different levels of physical productivity and demand, we want to focus our attention to the permanent component of these two types of structural shocks (see Guiso, Schivardi, and Pistaferri (2005) for a similar decomposition in another context).[2] Notably, the focus on permanent shocks also direct our attention to the aspects of the labor adjustment process which is likely to be most relevant from a policy perspective. Whereas key labor market institutions are explicitly set up to reduce the labor flows in the face of temporary shocks, a common fear is that these insitutions may hamper the necessary strucural process of adjusting labor flows between firms with different long-run prospects.

Our analysis depart from a model which closely follows the set-up of Foster, Haltiwanger, and Syverson (2008) and Foster, Haltiwanger, and Syverson (2012). The model presumes monopolistic competition which allows us to separate between firm-level technology and demand, and further assumes that the physical gross Solow residual is independent of all shocks except technology. Importantly, our application of the model allows other shocks, or changes in factor utilization or inventories, to affect the physical Solow residual temporarily without affecting the measured technology shocks. Empirically, we use a strategy similar to **?** and deflate the (nominal) firm-level output series with firm-level price indices. Thus, in a strict sense, we are not measuring

[1] There is also a small macro-oriented literature which tries 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), where the latter distinguish between neutral technology shocks and investment specific technology shocks and derive the consequences for job reallocation.

[2] The intuition behind our strategy resembles Guiso, Schivardi, and Pistaferri (2005) who extract the permanent component of firm-level value added. The key difference, apart from technicalities, is that we first separate between technology shocks and demand shocks.

physical output units of a homogeneous good, as in Foster, Haltiwanger, and Syverson (2008), but our empirical strategy handles cross-firm differences in quality by relying on firm fixed effects throughout.

In order to derive empirical measures of permanent idiosyncratic demand and technology shocks, we make use of an excellent longitudinal Swedish firm-level data set which includes information on inputs, outputs and output prices at the firm level. These data are then further linked to a longitudinal employer-employee data set which allows us to analyze worker flows. Using these data, we impose restrictions implied by the model as long-run restrictions in a structural VAR (SVAR) setting in the spirit of Franco and Philippon (2007). The system of equations we derive, which also explicitly allows for a sectoral factor price shock,[3] is recursive in a set of long-run restrictions allowing us to separately identify the structural innovations to technology and product demand although the system is unconstrained towards all short-run adjustments. In contrast to standard time-series implementations, including Franco and Philippon (2007), we rely on a large panel-data set for identification. Using Arellano and Bond (1991) dynamic panel data methods, which are well suited for our broad cross-sectional panel, we estimate both the parameters of the reduced form equations and the covariance matrix of the error terms with considerable precision, thus avoiding standard macro-data concerns regarding the practical implementation of SVARs. The focus on idiosyncratic shocks also allows us to analyze the direct impact of the shocks in a stable market environment, effectively abstracting from feedback effects through changes in market wages or aggregate unemployment. Also, shocks are relative ==> negative technology shocks are not so strange...

Due to the nature of our data, we are able to analyze many dimensions of firm responses to the shocks. In particular we corroborate the interpretation of the processes by showing that prices and output respond as expected to the identified shocks; firm-level prices are reduced in response to firm-level technology shocks, whereas they remain unaffected when demand shifts. Similarly, we find that idiosyncratic output increase in response to shocks to both demand and technology.

Our focus lies on how employment, hires and separations of workers respond to permanent idiosyncratic shocks. We find that, despite being crucial for both firm-level prices and output, firm-level technology shocks a very marginal impact on labor inputs. Our analysis instead clearly points towards product demand as the key driving force behind firm-level labor adjustments. A one

[3] The system also allows for mean reverting shocks to truly idiosyncratic factor pricess)

standard deviation shock to idiosyncratic demand increases employment by six percentage points.

Importantly, the adjustment is both rapid and symmetric across possible adjustment margins. The bulk of the adjustment takes place within a year. On average, firms adjust employment almost as much through changes in the separation rate as through changes in the hiring rate. They also adjust employment as much in response to negative as positive shocks. Although hires are reduced in response to negative shocks, the average firm continues to recruit workers even when hit by substantial negative shocks. We show that firms are very far from exploiting the full potential of downsizing thorugh reduced hirings, suggesting that the cost of inducing separations has a relatively minor impact on the adjustment behavior when shocks are permanent. In contrast, we show that firms adjust very little in response to temporary demand shocks.

Overall, 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 firms that face permanent idiosyncratic demand shocks do adjust their labor flexibly. We interpret this as suggesting that labor market rigiditites and adjustment costs which may hamper adjustments in the face of temporary shocks, are likely to be of a relatively minor importance when trying to understand the coexistence of firms with very different productivities.

The paper is organized as follows: Section (2) outlines our method and discusses the data. Section (4) reports the results. Finally, section (5) concludes. Throughout, we refer to appendices for robustness checks and computation details.

2 Model and empirical strategy

2.1 Modelling shocks to production functions and demand curves

The purpose of the paper is to measure how firms adjust their labor input in response to permanent idiosyncratic shocks. We are primarily interested in two types of processes which we can think of as structural shocks in the sense that they involve changes in key elements of the firms' profit functions: The first 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 and we label these shifts *demand shocks*. We refrain from modelling the origin of these shift in structural firm-level parameters. This implies, e.g., that we do not separate between shifts in the firm-specific demand curves that arise as a response to changes in the preferences of final consumers from shifts in the firm-specific demand curves that occur because of technological advancements among down-streams firms, or from shifts in the firm-specific demand curves due to quality changes that increases product demand for a given price. However, quality differences across firms will not affect our analysis since we exclusively focus on permanent changes within firms across time. [4] The key distinction between technology shocks and demand shocks therefore lies in how the shocks affect the producing firm, and not in the origin of the shock. This approach is fully consistent with the existing (micro) literature, see e.g. Foster, Haltiwanger, and Syverson (2008) or the survey by Syverson (2011).

To identify firm-level structural shocks we need to make assumptions about the technology and the market conditions faced by the firm. Our set-up follows Foster, Haltiwanger, and Syverson (2008) and Foster, Haltiwanger, and Syverson (2012) closely by using a first order approximation of both production technologies and product market demand and by modelling the key shocks as neutral shifters of the production function and the demand curve respectively. Thus, firms' production functions are 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} 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 allows us to account for idiosyncratic price differences across firms, so that our measures of technology refer to *physical* TFP (or TFPQ), rather than revenue productivity (or TFPR) in the terminology of Foster, Haltiwanger, and Syverson (2008). The equation presupposes a constant returns technology which is our main specification, but we also present robustness exercises where we relax this assumption. Furthermore, firm-level demand curves are approximated by constant-elastic functions according to

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

[4] It is straightforward to show that if we added a quality shock to the system derived below that drives a wedge between the measured price, based on unit values, and the quality adjusted price, it would enter the system symmetrically to the demand shock Ω_{jt} since the demand function, derived from preferences, should be in terms of goods of constant quality and the quality adjusted relative price. Note, however, 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), which limits the scope for quality changes to be the key component in our demand shock.

where P_{jt}/P_t is the firm's relative price and where $Y_t\Omega_{jt}$ denotes market demand and Ω_{jt} is a firm-specific demand shifter..

We modelling the time series properties of the shocks, we follow Guiso, Schivardi, and Pistaferri (2005) and Franco and Philippon (2007) and allow shocks to be permanent. More precisely, we specify the evolution of the demand and technlogy shifters as in Franco and Philippon (2007):

$$\begin{aligned} A_{jt} &= A_{jt-1}e^{\mu_j^a+\Phi_j^a(L)\eta_{jt}^a}, \\ \Omega_{jt} &= \Omega_{jt-1}e^{\mu_j^\omega+\Phi_j^\omega(L)\eta_{jt}^\omega} \end{aligned}$$

where μ_j^a and μ_j^ω are constant drifts, $\Phi_j^a(L)$ 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. In a variation of the model we also explicitly analyze the role of transitory shocks (see section 4.5).

2.2 Implied Restrictions and Empirical model

We use restrictions implied by the stylized model presented above in order to filter out shocks that permanently shifts the firms' production functions and demand curves. Table 1 sums up the restrictions (see Appendix A for details on derivations). Importantly, we only invoke the implied restrictions on the long-run behavior of the firms. The restrictions imply that: (i) Only the technology shock has a long-run impact on the measured physical total factor productivity (the Solow residual). As we ony impose this restriction in the long run, using gross output, the shock will however differ from standard measures of Solow residuals. Our measure of technology shocks is purged of relative changes output prices as well as changes in inventories, factor utilization and idiosyncratic input prices as long as these are mean reverting. (ii) Both technology and factor-price shocks, but not the demand shock, will have a long-run impact on wage neutral unit labor costs ($WNULC_{jt}$) as defined in table 1. (iii) All three shocks have a permanent effect on wage neutral demand (WND_{jt}) as defined in the final row of table 1. The long-run values of these variables are, under the assumptions of the model, independent of permanent or transitory wage shocks. As explained below, our empirical strategy also allows for an arbitrary set of alternative transitory shocks.

Table 1: The Core Structural VAR Equations

Variables:	Measured as:		
$Solow:$	$Y_{jt}\left[N_{jt}^{\alpha}K_{jt}^{\beta}M_{jt}^{1-\alpha-\beta}\right]^{-1}$	$=$	A_{jt}
$WNULC:$	$W_{jt}N_{jt}/Y_{jt} * \left(P_{jt}^{N}\right)^{-\alpha}$	$=$	$\alpha^{1-\alpha} * A_{jt}^{-1}P_{jt}^{F},$
$WND:$	$Y_{jt}\left(P_{jt}^{N}\right)^{\sigma\alpha}$	$=$	$\psi Y_t P_t^{\sigma} * (A_{jt})^{\sigma}\left(P_{jt}^{F}\right)^{-\sigma}\Omega_{jt}$

Note: DEFINE THE XI OF WND HERE, I CANT DO MATH IN NOTE HOWEVER

The left hand side variables can all be constructed from our firm-level data (see below for details). These three equations motivate a recursive sequence of long-run restrictions and in order to extract the shocks of interest from these series, we estimate a structural VAR. Since our interest lies in how other variables (such as output, prices and employment) respond to the shocks of interest, we include these other variables as fourth variables in the system. In practice, we rotate across these variables while keeping the core system of the first three equations intact as in **?**.[5] We allowing the fourth variable have a long-run effect on itself, but not on the other variables in the core system which implies that they will soak up all remaining transitory dynamics.

3 Data and Measurement

3.1 Data

Here we briefly describe the data we use, and how we measure the variables we include in the VAR and the final regressions. For details we refer to Appendix B.

Our firm-level data set is primarily drawn from the Swedish Industry Statistics Survey (IS). It contains annual information for the years $1990-2002$ on inputs and output for all Swedish manufacturing plants with 10 employees or more. About 72 percent of the plant/year observations in our sample pertains to plants that are also a firm and we therefore refer to the plants as firms.

One of our key assumptions is that only technology affects the Solow residual in the long run. For this assumption to be valid, sales must be deflated by a firm-level deflator since firm-specific relative prices are likely to respond to other idiosyncratic shocks. A crucial feature of the data is

[5] Parts of our analysis relies 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.

that it includes a firm-specific producer price index constructed by Statistics Sweden. The firm-specific price index is a chained index with Paasche links that combines plant-specific unit price values.[6]

To take the model outlined above to the data, we first compute a measure of firm level total factor productivity growth using

$$\Delta a_{jt} = \Delta y_{jt} - \Delta z_{jt}, \tag{3}$$

where Δy_{jt} is the growth rate (i.e. the log difference) of real gross output (obtained using the firm-specific price index), 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 (taken from the IS survey) and Δm_{jt} is the growth rate of intermediate materials (including energy). Moreover, C_J is the cost share of factor J in total costs. We measure the cost-shares as averages by 2-digit industry.

Given data on factor compensation, changes in output and inputs, the resulting residual Δa_{jt} provides a times series of changes in technology for the firm. Inputs are deflated using three-digit sectoral price indices, which imply that our model allow for three-digit sectoral input price shocks. ITs important to note that although (3) may not provide a good measure of technology due to varying factor utilization, inventories or truly idiosyncratic factor prices, our VAR filters out true technology shocks from (3) as long as technology shocks is the only factor that permanently shifts A_{jt} (i.e. as long as variations in factor utilization, inventories and factor prices within three-digit sectors are mean reverting).

When computing $\Delta wnulc_{jt}$ and Δwnd_{jt}, we use C_N as the estimate of α and thus let it vary by 2-digit industry. In order to compute wage neutral demand (Δwnd_{jt}), we need an estimate of the elasticity of substitution, σ. This estimate can be obtained from the demand equation (equation 2) by instrumenting the firm idiosyncratic price using the Solow residual, as in Foster, Haltiwanger, and Syverson (2008). As discussed earlier, our measure of the Solow residual should approximate firm's technology relatively well, since we measure true output volumes by using firm level prices as deflators. As such, the Solow residual is expected to affect firm level output only through firm level prices. The results suggest an elasticity of substitution equal to 3.306 (s.e. 0.075),which we

[6] 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 4-digits goods code level). The disaggregate sectoral producer-price indices are only used when a plausible goods-price index is not available.

use when computing Δwnd_{jt}. This result 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). In addition, the main results are insensitive to large changes in the calibrated value, and to the inclusion of sector specific values of σ.

Since 1996, Statistics Sweden are imputing the allocation of production across different plants within multi-plant firms. To ensure that our results are not driven by this procedure, we also present a series of robustness checks where we focus on single plants firms throughout, or use multiplant firms before 1996 but only single-plant firms thereafter.

In the end we construct series for Δa_{jt}, $\Delta wnulc_{jt}, \Delta wnd_{jt}$, Δy_{jt} for $7,940$ ongoing firms (observed at least during 5 consecutive years), over the period $1991-2002$. All in all, this amounts to $70,077$ firm/year observations. The sample covers nearly two thirds of all manufacturing employees. Further details regarding the data are given in the Appendix B.

The top panel of table 2 displays descriptive statistics of the structural shocks as well as firm level prices for the sample ($41,105$ observations in $6,137$ firms) for which we can construct the structural shocks (given the restrictions from the estimation approach; see next section).

Table 2: Summary Statistics

	Category	(1) Mean	(2) S.d.	(3) p(25)	(4) p(75)	(5) Observations
			Industry Statistics Data (IS)			
η_a	overall	-	0.101	-0.056	0.058	$41,105$
	within		0.101			
η_ω	overall	-	0.162	-0.086	0.085	$41,105$
	within		0.162			
Output Growth	overall	0.028	0.190	-0.073	0.133	$41,105$
	within		0.176			
Price Growth	overall	0.022	0.069	-0.001	0.044	$41,105$
	within		0.064			
			Registry Based Labor Market Statistics (RAMS)			
Employment Growth	overall	0.012	0.267	-0.062	0.089	$40,451$
	within		0.252			
Net Employment Rate	overall	0.012	0.208	-0.062	0.089	$40,451$
	within		0.195			
Hiring Rate	overall	0.150	0.151	0.063	0.200	$40,451$
	within		0.127			
Separation Rate	overall	0.138	0.152	0.061	0.174	$40,451$
	within		0.131			

Note: The "Within" rows show the dispersion within a firm. p(N) denotes the N:th percentile of the data.

In order to relate the firm-level structural shocks to employment and entry and exit of workers we link the firm-level data to detailed information on each of the employees in each of the firms in the sample using the Register Based Labor Market Statistics data base (RAMS) maintained by Statistics Sweden. This data base contains information about annual labor earnings with links to the employing firm for all private sector employees in Sweden. We are primarily interested in how firms adjustment the number of non-marginal employees as measured at the end of the year. We therefore measure employment in November following the practice of Statistics Sweden and only use employees working close to full time, focusing on each worker's primary job.[7]

Using these data, we compute various measures of job and worker flows using the metrics in the spirit of Davis, Haltiwanger, and Schuh (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 is defined such that the flow equation of employment holds.[8] All in all, we are able to match these flow measures to $6,130$ firms in the firm data (described above). Descriptive statistics of these measures for the sample where we can construct the structural shocks ($40,451$ observations) are found in the lower panel of 2.

3.2 The Shocks

3.2.1 Estimation

We extract the shocks from a VAR model conditional on year dummies and firm fixed effects. When estimating the VAR, we use an Arellano and Bond (1991) estimator which is developed in order to avoid the problems involved with fixed effects and lags of the dependent variable in a panel data setting (see Nickell, 1981). A key point here is that the Arellano and Bond (1991) estimator

[7] The raw data was compiled by the Swedish Tax Authority in order to calculate taxes. Data include information on annual earnings, as well as the first and last remunerated month received by each employee from each firm. Using this information, we can construct a measure of monthly wages for each employee in each of the firms in our sample. In order to restrict attention to workers employed working close to full time, we only keep employees whose (monthly) wage exceeds 75 percent of the mean wage of janitors employed by municipalities.

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

relies on asymptotics in the cross-sectional dimension. Since the identification of structural shocks with long-run restrictions crucially hinges on the quality of the estimated VAR coefficients and the covariance matrix, this is a very useful feature in the current context of a but wide, but short (12 years), panel of about $6,000$ firms. The procedure implies that we account for any firm-specific trends in the levels of the variables in the model. We also include time dummies which capture all aggregate shocks shared by different firms within the manufacturing sector.

Appendix C discusses specification tests and shows impulse responses and variance decompositions of the main VAR model using output as a fourth variable alongside the core system. Two particular results are relevant for the analysis ahead. The first is that we find a fairly limited amount of dynamics, in particular in the Solow residual. The main reason for this somewhat surprising 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 Nordstrom Skans, 2012, for direct evidence on relative price dynamics).

The second finding is that the residual shock explain very little of the variance in our key variables. Since the model is estimated conditional on time dummies, this finding is in line with the result of Franco and Philippon (2007) who show that transitory shocks, although highly correlated across firms (and therefore of macroeconomic importance), matter only marginally at the firm level.

The appendix also displays the shock distributions, normalized to have a unit standard deviation which will be the normalizatio we will rely on in most of what follows. However, 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 units, respectively).

3.2.2 Validation: The Impact on Prices and Output

Since the shocks we are analyzing are idiosyncratic, we cannot cross-validate their interpretation through correlations with aggregate shocks such as oil price or exchange rate movements unless we have strong priors regarding heterogeneous impacts of these shocks across firms. Instead, we perform a number of 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, as predicted from the theoretical model, negative. Similarly, the estimated responses from both technology shocks and factor prices on Δwnd_{jt} are negative, as expected from the model.

A second piece of evidence comes from relating the structural shocks to our firm-specific price index and to output. It should be clear from how the shocks are defined that a positive technology shock only can affect sales if prices goes down (since the demand curve is fixed). In contrast, demand shocks, defined as shifts in the firm-specific demand curve, allows the firm to sell more at a given price suggesting that prices should remain unchanged in the response to a demand shock unless the firm's technology features non-constant returns to scale, or input prices change when the scale of production is altered. Hence, theory suggest that technology and demand shocks should affect output, whereas prices primarily should respond when technology changes. Figure 1 shows the impulse responses of prices and output to our measured shocks of interest. The picture clearly validates the theoretical predictions; output responds to both shocks (although somewhat more to a one standard deviation demand shock), but prices only respond to the technology shock.

Finally, results presented in section 4.5 explore the relationship of our estimates to outcomes from other, alternative, empirical methods proposed in the literature.

4 Results

4.1 Idiosyncratic Shocks and Employment Adjustment

In order to measure how employment respond to permanent idiosyncratic technology and demand shocks, we first estimate the SVAR-system using employment as the fourth variable. Figure 2 shows the impulse responses with confidence bands. A key result emerging from this picture is that idiosyncratic demand shocks are substantially more important for firms' adjustments of labor inputs than idiosyncratic technology shocks. A one standard deviation shock to demand increases employment by slightly more than 6 percentage points, whereas the impact of a technology is negligible (and statistically insignificant). A second important result is that the time dynamics in labor adjustments is limited. More than 90 percent of the long run adjustments in response to permanent demand shocks takes place within the first year.

In order to analyze the robustness of these results, we have estimated a wide set of variations of

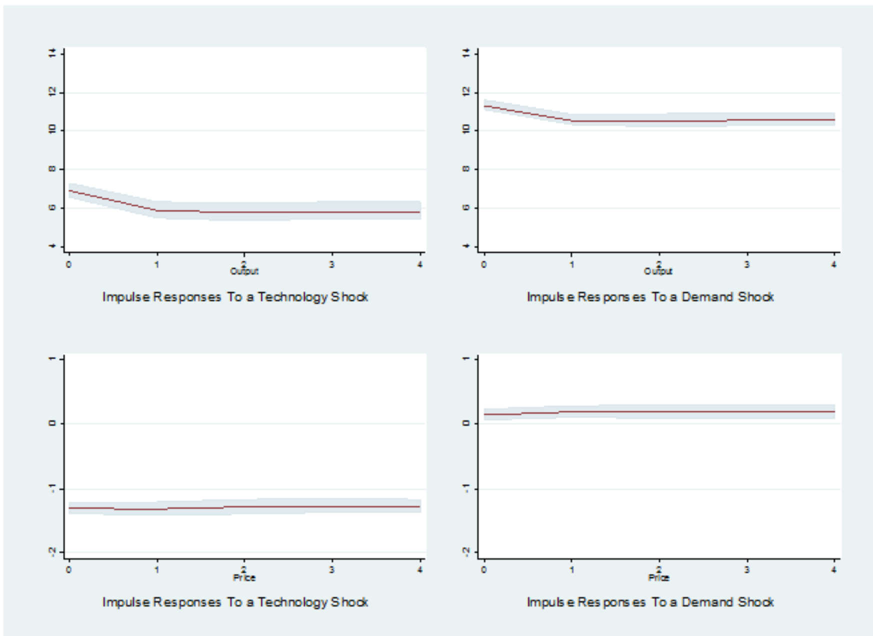

Figure 1: Impulse Responses in percentage points of Output and Price to a Technology- and a Demand Shock. Lines depict the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

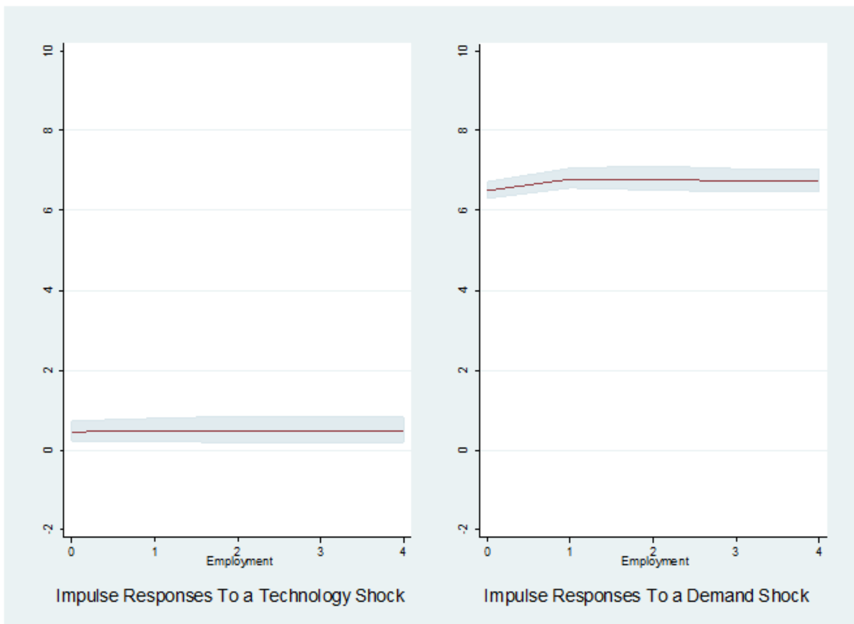

Figure 2: Impulse Responses in percentage points of Employment to a Technology- and a Demand Shock. Lines depict the means of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

the original model. First, we have varied the modelling of the fourth variable, which is included in the system in order to improve the model's performance in terms of short-run dynamics. In the end, however, this variable plays a negligible role for the results. This is already hinted at in the variance decomposition shown in Appendix C. To make this point more explicitly, we have re-estimated the VAR using different variables as the fourth variable (sales per workers, employment, prices and wages) extracting the shocks of interest, and then analyzed their impact on employment, without any major impact on the results, see Appendix D.

We have also re-estimated the VAR-model using alternative values of the demand elasticity. Our main model uses an elasticity of $3,3$ based on estimates using the Solow residual as an instrument for prices (following Foster *et. al.*, 2008). We have experimented with a wide range of assumptions for σ (from 1.1 to 10) with robust results. We have also reestimated the model allowing for industry specific estimates of the demand elasticity instead, but the results are unchanged (see Appendix E).The main reason for the low sensitivity to the estimated demand elasticity is that it enters the system with a weight equal to the labor share, which is around 0.25 since we rely on a gross production function.

The main model assumes constant returns to scale, but it is straightforward to change the model to incorporate increasing or decreasing returns to scale instead. Changing the assumed returns to scale (see Appendix E) affects both the estimated sign and statistical significance of the technology shock in the employment regression. However, the main message still holds: in spite of substantial variations of the assumed returns to scale the demand shock remains more important than the technology shock by at least an order of magnitude.

One general advantage of our approach relative to macro data VAR:s is that we are able to estimate the system with considerable precision by relying on cross-sectional asymptotics. This comes at the (potential) cost of assuming that the dynamic process is equal across different firm types. In order to address this concern, we have reestimated the model allowing for separate dynamics for each two-digit industry, and the results remain unchanged (see Appendix E).

Since the allocation of output across plants within multi-plant firms during parts of our sample period (after 1996) is imputed, we have also redone the analysis for the sample of single plant firms, as well as for a mixed sample including multi-plant firms until 1996, but not thereafter. The results, presented in appendix G, are completely robust to these alterations of the sample.

In appendix G, we also estimate models focusing on "normal" shocks (in the Lester range, from

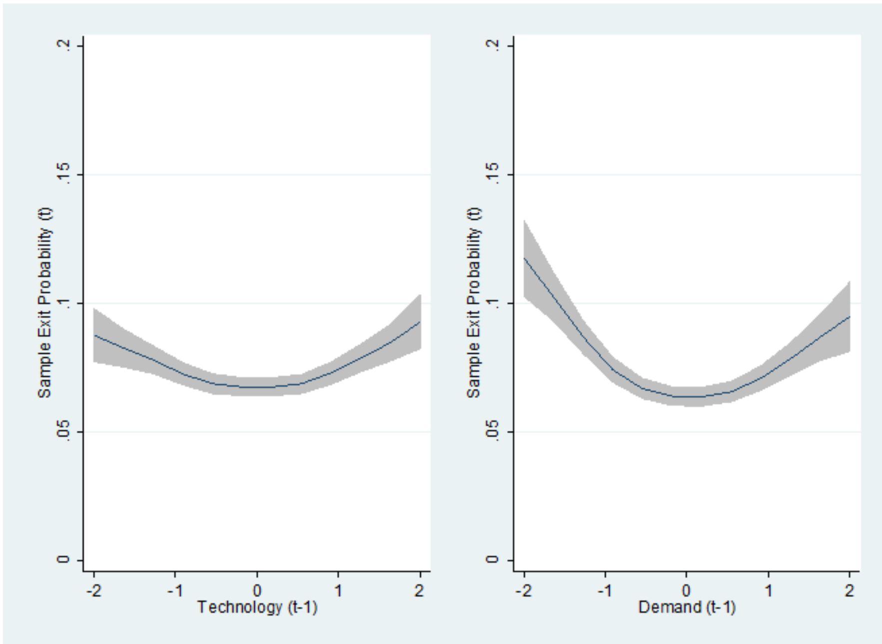

Figure 3: Sample Exit Probability as a (non-linear) function of an x-standard deviation lagged Technology- and Demand shock. Shaded areas depict 95-percent confidence intervalls.

-2 to 2 standard deviations). This restriction does, however, not change any of our results.

Finally, a possible concern with the analysis is that we disregard the process of firm exit. In order 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 very large negative demand shocks. The fact that demand shocks are more important for firm exits than technology shocks is well in line with results in Foster, Haltiwanger, and Syverson (2008). In order to see if this has any bearing on the results of our main analysis, we have analyzed the employment impact of the shocks across two periods instead (which allows us to include exiting firms, see Appendix F for details), but the results are insensitive to weather we include or exclude exiting firms.

IS THIS IN OR OUT? In our view, given the robustness of the result, it seems reasonable to conclude that shifts in firms' idiosyncratic demand curves are a much more important driving force behind firms' idiosyncratic net employment adjustment than shifts in firms physical production functions.

4.2 Idiosyncratic Shocks, Hires and Separations

In order to shed further light on the firms' labor adjustment process, we now turn to a micro-level analysis of firms' different possible adjustment margins. In order to make the analysis more concise, we extract the shocks from the VAR-system using output as the fourth variable and analyze their impact on a set of labor adjustment concepts while focusing entirely on the short run impact of the shocks.

We analyze three measures of labor adjustments in the Davis, Haltiwanger, and Schuh (1996)-tradition: the *Separation rate*, the *Hiring rate*, and *Net employment*. The hiring rate is defined as the number of entrants (i.e. workers that did not work in the firm in $t-1$ but work in the firm in t) divided by the average employment over the current and the lagged year. Similarly, the separation rate is defined as the number of employees who worked in the firm in year $t-1$ but does not work there in t, divided by average employment across the two years. The net employment change is defined as the change in employment between the two years, i.e. the difference between the hiring rate and the separation rate.

We think of the three measures as choice variables for the firm (i.e. firms choose employment growth and turnover based on the shocks), and we estimate an equation of the form

$$Outcome_t = \eta^a_{jt}\delta_1 + \eta^{\omega}_{jt}\delta_2 + \rho_t\beta_\rho + \mu_j + \xi_{jt}, \tag{4}$$

where *Outcome* denotes each of the three variables for firm j at time t. The coefficients δ_1 and δ_2 measures the impact of the firm-level structural shocks on the outcomes.[9] Moreover, we include time, ρ_t, and firm fixed effects, μ_j, in line with the VAR formulation above. This ensures that identification is driven by idiosyncratic, rather than aggregate shocks.

[9] 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 VAR.

Table 3: Contemporaneous Effects

	(1) Net Employment	(2) Hiring Rate	(3) Separation Rate
A) *One s.d. shock*			
η_a	0.115	-0.050	-0.165*
	(0.119)	(0.075)	(0.078)
η_ω	5.609**	2.906**	-2.703**
	(0.173)	(0.096)	(0.120)
B) *Elasticities*			
η_a	0.011	-0.005	-0.016*
	(0.012)	(0.007)	(0.008)
η_ω	0.347**	0.180**	-0.167**
	(0.011)	(0.006)	(0.007)
Observations	40,451	40,451	40,451
Firms	6,125	6,125	6,125

Robust standard errors in parenthesis. Regression includes time dummies and firm fixed effects. Regression sample limited to observations where the absolute value of both the technology and the demand shock is less than or equal to two standard deviations in size.

The results can be found in table 3. We present estimates scaled according to standard deviations of the shocks, and as elasticities (see Appendix A for details). As expected from the results of the previous subsection, the impact of technology on net employment is small (and not statistically significant). [10] The point estimate for the short-run impact suggest that a one standard deviation increase in TFPQ decreases employment by 0.12 percentage points, and the estimate is therefore insignificant despite being estimated with considerable precision. In elasticity form, the short run impact is 0.01. The impact of demand shocks are substantially larger, both in terms of the impact of a normal (one s.d.) shock and in terms of elasticities. A normal shock to demand increases employment by 5.60 percentage points on impact. This effect corresponds to an elasticity of 0.35.

Turning to worker flows instead we see a that 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. Thus, on average, the adjustment of net employment is achieved to 52 percent using the hiring margin and to 48 percent using the separation margin. These numbers should be compared to

[10] Note that results in this and the previous section do not need to be identical since employment changes are measured from two different sources. Employment changes here are obtained from RAMS; which measures employment spells that are observed in November each year. Employment changes in the previous section are obtained from the firm-level data (IS), which measures average yearly employment in the firm.

an average hiring and separation rate of about 14 percent each, see Table 2 above. This result is 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 adjustments costs.

4.3 Idiosyncratic Shocks and Nonlinearities

As shown above, firms appear to adjust their labor in response to the demand shocks almost as much through separations as through hires. But an interesting question is to what extent the adjustments differ between positive and negative shocks.

When analyzing potential nonlinearities in the impact of the shocks on overall employment it is obvious that firms' adjustment to positive and negative shocks may differ due to asymmetries in rigidities on other adjustment margins faced by the firm. In addition, the VAR assumes linearity in the adjustments to the variables that are included in that system. However, Appendix C shows that adjustments of the elements within the VAR, including output, are approximately linear, suggesting that other, non-labor adjustment impediments are of minor importance in this context.

Figure 4 shows how firms adjust hirings in response to positive and negative shocks of different magnitudes. Each figure shows the predicted impact based on regressions allowing for a separate second order polynomial above and below zero.[11] The estimates should be interpreted as deviations from a zero-shock state. For completeness, we show the responses to both technology and demand, but as expected from the results above,we see very little adjustments in response to technology shocks, and therefore focus our attention towards the demand-shock responses. Two patterns are particularly noteworthy. The first is that the response to positive shocks is exactly linear. That is, the impact of a two standard deviations shock is exactly twice that of a one standard deviation shock, suggesting that the costs of increasing hirings are approximately linear. The second is that the response in terms of hiring is considerably smaller if shocks are negative. The results suggest that firms that are hit by a two standard deviations negative demand shock (and thus, on average, reduce their employment by about 10 percent), continue to hire at a rate of about 12 percent, a result which we return to below.

[11] In order to facilitate the interpretation of the graphs, we show the sum of the predicted estimates and the average net employment change in the sample.

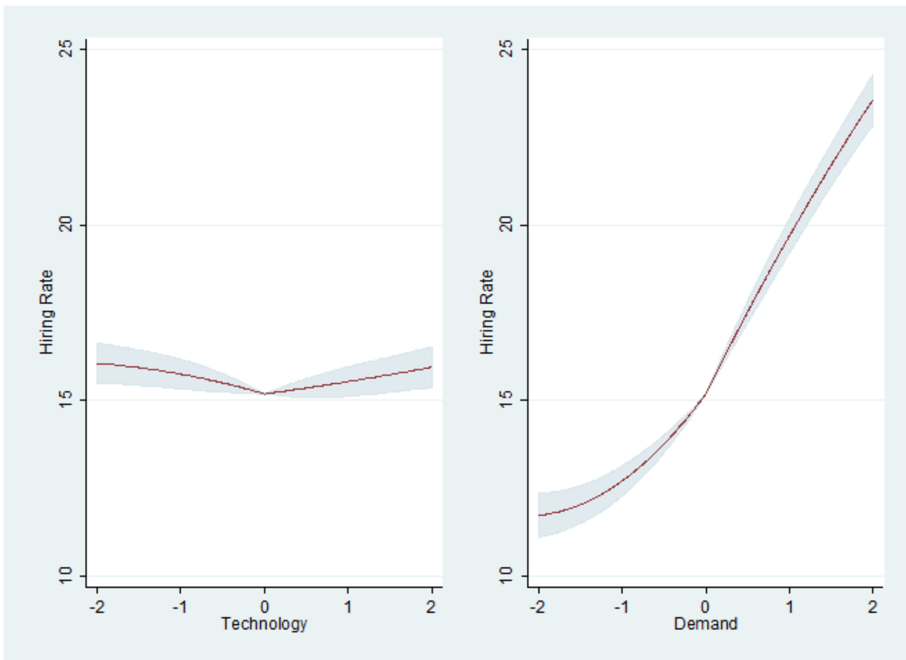

Figure 4: Contemporaneous Hiring Rate in percentage units as a (non-linear) function of an x-standard deviation Technology- and Demand shock. Shaded areas depict 95-percent confidence intervalls.

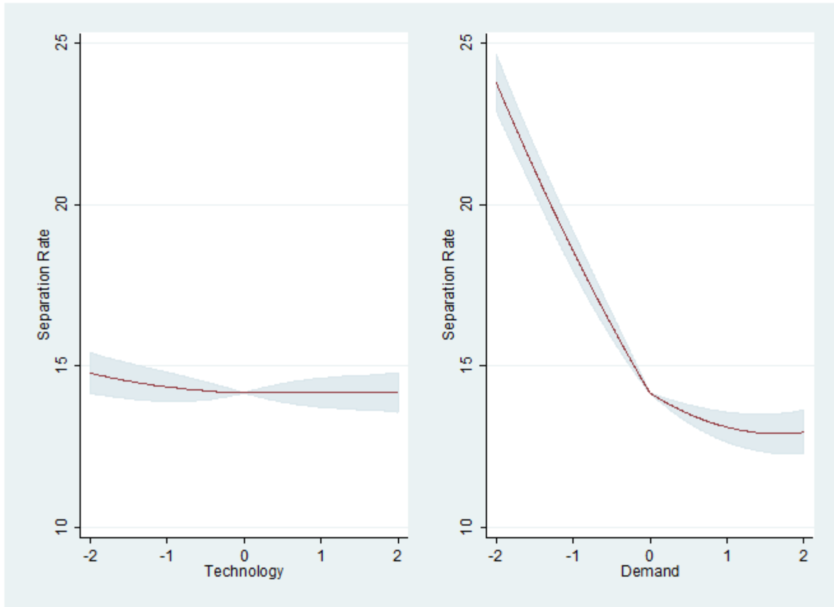

Figure 5: Contemporaneous Separation Rate in percentage units as a (non-linear) function of an x-standard deviation Technology- and Demand shock. Shaded areas depict 95-percent confidence intervalls.

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. When the shock is negative, the response is exactly linear. Thus, the estimates suggest that a two standard deviations negative shock causes a separation response which is exactly twice as large as the response to a one standard deviation shock. This suggests that the costs of increasing separations are approximately linear (or non-existent) on average. The result, together with figure 4, also imply that firms primarily adjust employment through separations in response to permanent negative shocks despite ample opportunities to adjust further through reduced hires. Figure 5 also shows that separations goes down when shocks are positive, but this impact is somewhat smaller than the hiring-response to negative demand shocks.

As could be imagined from the combination of figures 4, and 5, the effects on net employment are fairly linear (see figure 6). The fact that the kink at zero is more pronounced for hires than

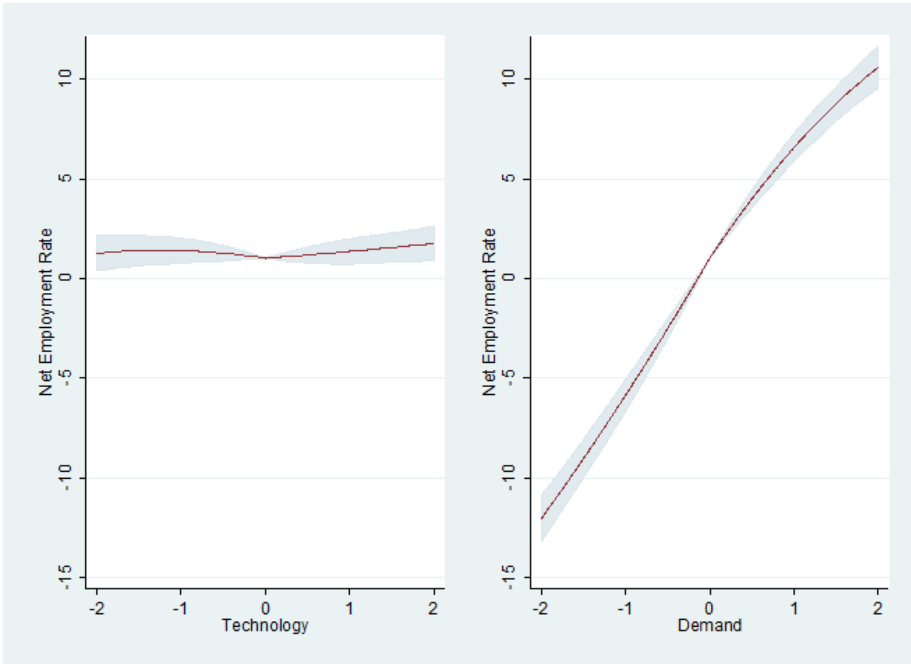

Figure 6: Contemporaneous Net Employment Change in percentage units as a (non-linear) function of an x-standard deviation Technology- and Demand shock. Shaded areas depicts 95-percent confidence intervalls.

for separations, implies that the slope is somewhat lower on the positive side for net employment, but although the difference in slopes is statistically significant, the difference in magnitude is fairly small, impling approximately 11 percent adjustment in response ot a 2 standard deviations positive shock and a 13 percent adjustment in response to a 2 standard deviation negative shock.

.

4.4 Decomposing Net Employment Adjustment in Response to Permanent Demand Shocks (this text is very preliminary)

The results presented above depicts the responses of hires, separations and net employment to firm-level shocks. The results seem to suggest that firms induce separations rather than reduce hiring when hit by negative shocks. To complement this analysis we here present an explicit decomposition of firm-level net employment adjustments in response to permanent demand shocks. We focus on permanent demand shocks since the other shocks are found to have a negligable impact on net employment.

A set of previous papers (Abowd, Corbel and Kramarz, 1999, and Davis, Faberman and Haltiwanger, 2012) have decomposed positive and negative changes in net employments into different components. The results differ between these studies in particular in the dimension that Abowd, Corbel, and Kramarz (1999) find that shrinking French firms reduce employment by reduced hirings, a result which do not appear in Davis, Faberman, and Haltiwanger (2012) data from the U.S. A key difference is that these studies have presented raw decompositions, where short-run fluctuations in employment may be driven by quits rather the other way around. In contrast, we instrument the employment adjustment by the demand shock which allows us to provide a decomposition in response to a well-defined non-labor shock with a long-run impact on the optimal employment level of the firm. As far as we are aware, the flows have never previously been related to measures of structural shocks.

In practice, we proceed by characterizing employment adjustment 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.[12] 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 can be easily be deduced. The results are presented in the left-hand panel of figure 7. Consistent with what we saw when focusing on the direct impact of the shock, we find a strong and linear relationship between net employment adjustments and hirings when shocks are positivebut a very modest relationship between the magnitude of negative employment adjustments and hirings. To make this point clear, the right-hand side of figure 7 shows the share of employment adjustments through hires as function of demand-induced net employment adjustments. This share jumps from 20 to 95 precent when employment adjustments become positive instead of negative. .

One interpretation of the figure is that firms are relatively unconstrained in their use of separations in the sense that they seem to rely on increased separations even when they could have adjusted through reduced hires.To make this point more precise figure 8, repeats the patterns shown

[12] The IV strategy essentially implies that we scale the shock impact on hirings presented in figure 4 above with the first-stage which correspond to figure 6

in the righ-hand side of figure 7, but focusing on negative values. As contrasts, the figure also depicts two hypothetical adjustments. The first, assuming homogenous firms, imposes the observed average steady state (i.e. whithout employmet adjustment) separation rate of 10 percent on all the firms, and showing how much the firms could have reduced employment by only relying on hires. As long as the need for adjustment is 10 percent or less, reduced hires could fully accomodate the necessary adjustments and if the shock is 20 (30) percent instead, the firm could instead accomodate 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 is not the case. If we instead assume that each individual leaves each firm at a rate of 10 percent, we can tease out a distribution of quit rates accross the firms. By randomly allocating quits across the workers within our full sample and then aggregating to the firm level, we get the firm-level distribution of quit rates. Whith this distribution, which naturally gets wider if firms are small, some firms will not experience any quits at all, which means that they are unable to accomodate even the smallest employment adjustments through reduced hires, whereas other firms will experience many random separations, allowing them to accomodate very large employment adjustments 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 idea here is that in rigid world where employees own their jobs as long as firms are hiring someone, then firms would adjust according to the hypothetical curve. As is evident, real adjustments are far from this baseline. The distance between the heterogeneous hypothertical curve and the actual behavior of the firm could be interpreted as a region of flexibility. It shows 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.

To further investigate the role played by size and heterogeneity in restricting adjustments, appendix H shows separate estimates by size brackets and by the amount of skill dispersion within firms. We find very small differences accross these different firm types.

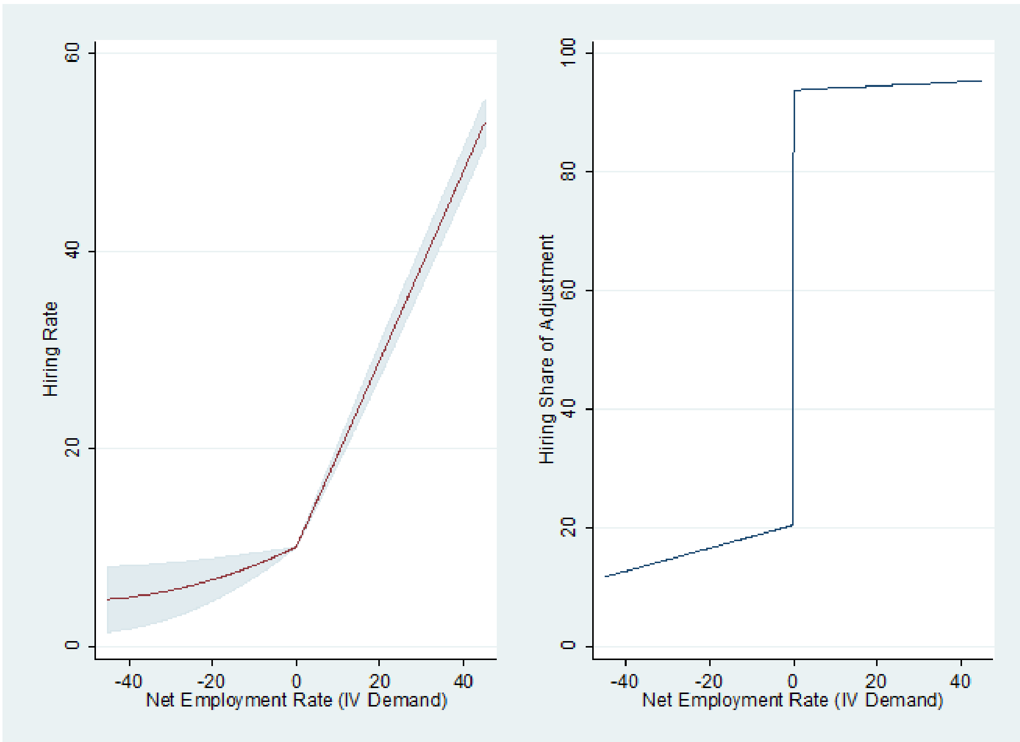

Figure 7: Left-side panel: Contemporaneous Hiring Rate in percentage units as a (non-linear) function of employment adjustment in percentage units. Employment adjustments are instrumented by demand shocks. The model imposes a separate quadratic polynomial above and below zero for both employment adjustment and the instrument. Shaded areas depicts 95-percent confidence intervalls. 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.

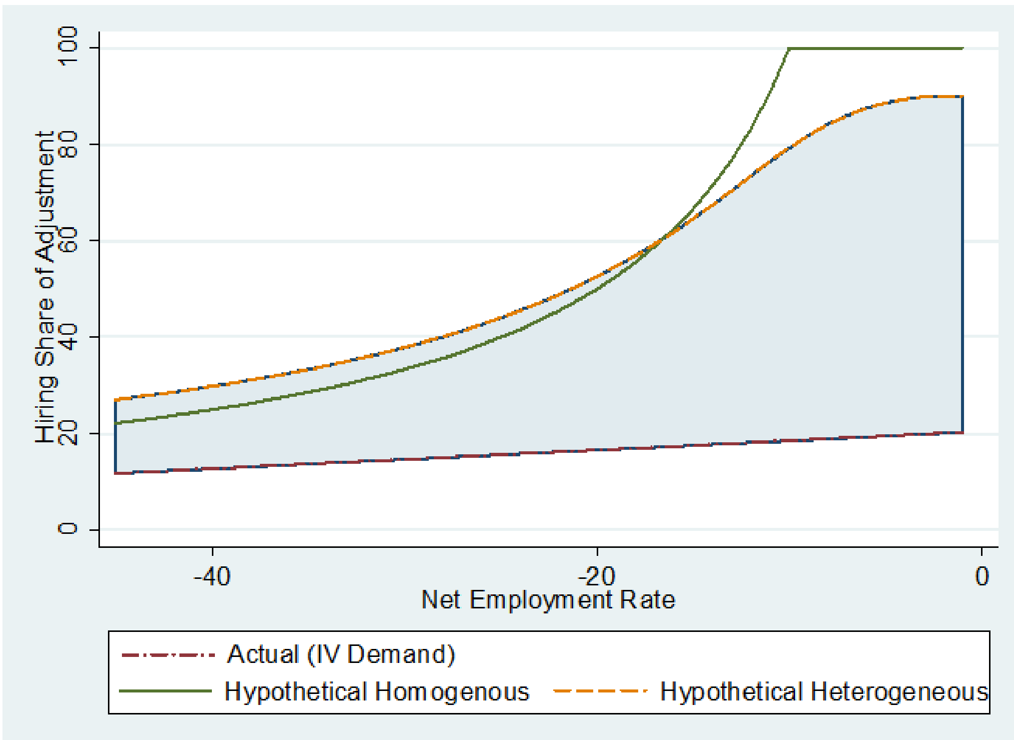

Figure 8: Actual (estimated from data) and hypothetical maximum (simulated) fraction of negtive employment adjustments achieved through changes in hirings. "Hypothetical homogenous" assumes that the same frction of workers always leaves the firm. "Hypothetical heterogeneous" imposes a random individual quit rate on the actual firm-size distribution.

4.5 Transitory Idiosyncratic Shocks (Before IV?)

The focus of the analysis so far has been to disentangle how firms adjust their labor input when hit by permanent idiosyncratic shocks. It is however noteworthy that the variance decomposition shows that the technology shock is very close to being the sole determinant of the physical gross Solow residual at all horizons. As a result, the correlation between the Solow residual and the permanent structural innovation is close to unity (0.980 see table). This has two implications, the first is that the physical gross output Solow residual is a valid measure of technology shocks. The second is that technology shocks appear to have only have a marginal transitory component.

The situation for demand shocks is, however, quite different. To derive a measure of transitory demand shocks, we first follow Foster, Haltiwanger, and Syverson (2008) and use the Solow residual as an instrument for prices in a demand equation. Since, the ensuing residuals represent changes in sales without price adjustments they serve as a measure of demand shocks (henceforth, *FHS-demand*). In contrast to the SVAR filter, this procedure does however not differentiate between permanent and transitory shocks.[13] The correlation between FHS-demand and our regular demand shocks is 0.538 (see table 4) and the standard deviation is considerably higher for the FHS-demand (0.24, vs. 0.16). This suggests that the two demand shock series contain a common component without being identical.

Using FHS-demand and the Solow residual as measures of shocks (see, table 4) and estimating the impact on employment, produces demand estimates that are less than half the size as in our baseline specification (2.8 vs. 6.3), at the same time as the point estimate for technology shocks becomes somewhat larger. However, the magnitude of the technology shock is still tiny and the overall impression that demand is the key driving force behind job reallocation is preserved; demand shocks are estimated to be ten times more important than technology shocks also in this case.

As noted above, FHS-demand includes both transitory and permanent shocks. In order to see if the impact of demand shocks differ depending on their time series properties, we have decomposed the FHS-demand into two components. Running a regression with FHS-demand as the dependent variable and our permanent shocks as regressors, we use the residual as a measure of transitory demand shocks and contrast the impact of these shocks with the impact from the permanent shocks. [14]

[13] In addition, these alternative series are not necessarily uncorrelated with other shocks.

[14] The decomposition resembles Guiso, Schivardi, and Pistaferri (2005) who extract the permanent component of

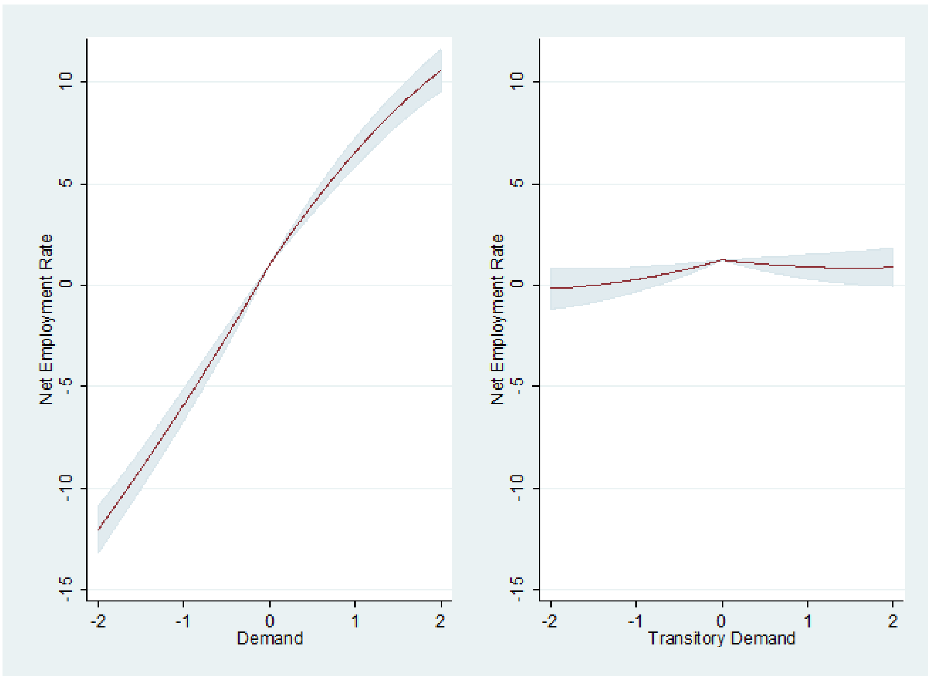

Figure 9: Contemporaneous Net Employment Rate in percentage units as a (non-linear) function of an x-standard deviation (permanent) Demand shock and as a function of a Transitory Demand shock (residual component of FHS-demand). Shaded areas depict 95-percent confidence intervalls.

Table 4: Baseline estimates vs. Solow residuals and FHS-demand

	(1) Baseline	(2) FHS
η_a	0.153	0.33283*
	(0.159)	(0.168)
η_ω	5.986**	3.40621**
	(0.233)	(0.183)
Observations	40,451	40,451
Firms	6,125	6,125
S.d. η_a	10.06	10.58
Correlation with baseline η_a	1	0.980
S.d. η_ω	16.18	23.90
Correlation with baseline η_ω	1	0.538

Robust standard errors in parenthesis. Regression includes time dummies and firm fixed effects. Regression sample limited to observations where the absolute value of both the technology and the demand shock is less than or equal to two standard deviations in size.

We analyze the impact on our measures of worker flows in figure 9. The results show that the impact of the transitory shocks is substantially lower than the impact of the permanent shocks. These results which resembles those of Guiso, Schivardi, and Pistaferri (2005) who show that wages respond to permanent shocks, but not to transitory shocks, implies that firms' adjustment of employment crucially depend on the time series properties of the shocks that hit them. This is important because the welfare consequences of firms' lack of adjustment are likely to crucially depend on these properties; (labor hoarding vs inability to structurally adjust). It also suggest that cross-country comparisons of labor flows would become more informative if they would be able to account for the time series properties of the underlying structural shocks

5 Conclusions

This paper has analyzed how ongoing firms adjust their labor inputs in response to idiosyncratic firm-level innovations in technology, demand, and the prices of other inputs. We estimate the shocks using a structural VAR relying on long-run restrictions derived from a simple firm-level model assuming that constant returns to scale and monopolistic competition with isoelastic demand provide

firm-level value added using high order polynomials of lags as instruments. Although the mechanics of the methods differ, the underlying logic is very similar.

reasonable approximations of firms' long-run behavior. We estimate the model on a unique data set which merges information about inputs, outputs and prices of Swedish ongoing manufacturing firms to a linked employer-employee data set. We estimate the VAR using dynamic panel data methods which allows us to identify the system with considerable precision. The interpretation of the shocks is supported by theory-consistent responses in firm-level prices which are shown to fall in response to technology shocks and increase in response to input prices, but remain independent of product demand innovations.

The results of this paper show that both the nature and the time series properties of the shocks matter. Permanent demand shocks, which affect output, but not relative prices, has a very pronounced impact on employment. Technology shocks on the other hand have very little effects on employment, despite affecting both output and relative prices. Similarly, temporary demand shocks have very little effect on employment adjustments.

Further results suggest that employment adjustments 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 rate as through the adjustments of hiring rates. There are also no signs of non-linear costs of 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. The adjustments towards temporary shocks are however heavily muted. Considering that OECD ranks Sweden as an average country in terms of institutional labor market rigidities, this may suggest that labor markets within the OECD in general are flexible enough for firms to accommodate the impact of permanent shocks, while pushing firms towards hoarding labor when hit by temporary 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.

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.

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.

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

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

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

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. 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.

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*.

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.

Nickell, S. J. (1981): “Biases in Dynamic Models with Fixed Effects,” *Econometrica*, 49, 1417–1426.

Roodman, D. (2006): “How to Do xtabond2: An introduction to "Difference" and "System" GMM inStata,” 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.

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

Appendices

A Identification of Structural Shocks

A.1 Derivation of Long-run Restrictions

We use the stylized model presented in the paper to filter out shocks that permanently shifts the firms' production functions and demand curves. To filter out the shocks of interest, we first note that the assumptions of the model ensures 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 are able to 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

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

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

$$ULC_{jt}\left(W_{jt}\right)^{-\alpha} = \kappa A_{jt}^{-1} P_{jt}^{F}, \tag{A2}$$

where unit labor cost, ULC_{jt}, is defined as $W_{jt}N_{jt}/Y_{jt}$ and $\kappa = \alpha^{1-\alpha}$. Thus, expression (A2) will be affected by technology as well as 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 marginal cost net of wage setting disturbances. Henceforth, we will label the variable wage-neutral labor cost ($WNULC_{jt}$).

Using expressions (2) and (A1), as well as the fact that the price is set as a constant markup over marginal cost we arrive at

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

where $\psi = \left(\frac{1}{\alpha}\right)^{-\sigma\alpha}\left(\frac{\sigma}{\sigma-1}\right)^{-\sigma}$. Thus, expression (A3) will, apart from aggregate factors (which will be captured by time dummies in the empirical implementation below), be driven by shocks to technology, factor prices other than labor, and demand. In effect, expression (A3) is demand adjusted for wage setting disturbances. Thus, we label this variable wage-neutral demand (WND_{jt}).

A.2 The Structural VAR

The model outlined in the paper provides a set of three equations that depend on the three structural shocks. The left hand side variables in these equations, which all can be constructed from our firm level data, motivate a recursive sequence of long-run restrictions. In order to extract the shocks of interest from the four series, we estimate a VAR. Since our interest lies in how other variables (such as output, prices and employment) respond to the shocks of interest, we include 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. They will thus 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 **?**.[15]

The VAR system, which is a fully interacted dynamic system of the variables, can, under standard regularity conditions, be written (on vector moving average form, using lower case letters for logarithms and denoting the fourth variable by θ) as:

$$\begin{bmatrix} \Delta a_{jt} \\ \Delta wnulc_{jt} \\ \Delta wnd_{jt} \\ \Delta\theta_{jt} \end{bmatrix} = \begin{bmatrix} C_{11}(L) & C_{12}(L) & C_{13}(L) & C_{14}(L) \\ C_{21}(L) & C_{22}(L) & C_{23}(L) & C_{24}(L) \\ C_{31}(L) & C_{32}(L) & C_{33}(L) & C_{34}(L) \\ C_{41}(L) & C_{42}(L) & C_{43}(L) & C_{44}(L) \end{bmatrix} \begin{bmatrix} \eta^a_{jt} \\ \eta^f_{jt} \\ \eta^\omega_{jt} \\ \eta^\theta_{jt} \end{bmatrix}. \tag{A4}$$

We assume that the shocks themselves (i.e. $[\eta^a_{jt}, \eta^f_{jt}, \eta^\omega_{jt}, \eta^\theta_{jt}]$) are structural innovations and hence mutually orthogonal and serially uncorrelated. Since 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. Since the shocks are orthogonal (and using a standard normalization) we get $E\boldsymbol{\eta}_t\acute{}\boldsymbol{\eta}_t = \mathbf{I}_t$, where $\boldsymbol{\eta}_t = [\eta^a_{jt}, \eta^f_{jt}, \eta^\omega_{jt}, \eta^\theta_{jt}]'$.

Following standard practice, we denote the elements of the matrix of *long-run* multipliers corresponding to (A4) as $C_{rc}(1)$. Relying on the model outlined above, we know that the technology shock, η^a_{jt}, 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.[16] Similarly, only the technology and the factor price shock have a long-run effect on $wnulc_{jt}$, so $C_{23}(1) = C_{24}(1) = 0$. Finally, since the residual shock has no

[15] Parts of our analysis relies 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.

[16] 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 $c_{14}(k)$.

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 (A4), i.e. from

$$\Delta \mathbf{x}_t = \textstyle\sum_0^P \mathbf{A}_p \Delta \mathbf{x}_{t-p} + \mathbf{e}_t, \tag{A5}$$

where $\mathbf{A}_p$ is matrices with coefficients, $\Delta \mathbf{x}_t = [\Delta a_{jt}, \Delta wnulc_{jt}, \Delta wnd_{jt}, \Delta \theta_{jt}]'$, $\mathbf{e}_t$ is vector of reduced form disturbances and we have suppressed constants to save on notation (see Appendix A for details regarding the relationship between reduced form and structural representations). Under standard regularity conditions, there exist a VAR representation of the MA representation (A4) of the form

$$\mathbf{x}_t = \mathbf{A}(L)L\mathbf{x}_t + \mathbf{e}_t, \tag{A6}$$

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

$$\mathbf{e}_t = \mathbf{c}(0)\boldsymbol{\eta}_t, \tag{A7}$$

where $\mathbf{c}(0)$ is the matrix of $c_{rc}(0)$ coefficients from the MA representation and $\boldsymbol{\eta}_t = [\eta^a_{jt}, \eta^f_{jt}, \eta^\omega_{jt}, \eta^\theta_{jt}]'$. Thus, if the $\mathbf{c}(0)$ coefficients where known we could recover $\boldsymbol{\eta}_t$. To derive the first 10 restrictions we need to solve for the 16 elements in $\mathbf{c}(0)$. We first use (A7) and that $E\boldsymbol{\eta}_t\boldsymbol{\eta}_t' = \mathbf{I}_t$ together with an estimate of $\Omega = E\mathbf{e}_t\acute{\mathbf{e}}_t$ obtained from an estimate of (A6). To get the additional six restrictions we impose three long-run restrictions. Rewriting (A6) we can obtain the MA form, by using (A7), in terms of coefficients from (A6) and the $\mathbf{c}(0)$ coefficients as

$$\mathbf{x}_t = [I - \mathbf{A}(L)L]^{-1}\mathbf{c}(0)\boldsymbol{\eta}_t. \tag{A8}$$

Then, the six long-run restrictions implies equally many restrictions on the matrix $[I - \mathbf{A}(L)L]^{-1}\mathbf{c}(0)$, that together with an estimate of (A6) yields six additional restrictions on $c(0)$. Finally, given an estimate of the $c(0)$ matrix, $\hat{c}(0)$, we can solve for the structural shocks using (A7)

$$\hat{\mathbf{c}}(0)^{-1}\hat{\mathbf{e}}_t = \hat{\boldsymbol{\eta}}_t. \tag{A9}$$

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 $\hat{c}(0)$ where each element is divided by the column diagonal element.

B Data

The firm data set we use is primarily drawn from the 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 with at least 10 employees and that we observe in a spell with at least 5 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 links which combines plant-specific unit values and detailed disaggregate producer-price indices (either at the goods level when available, or at the most disaggregate sectoral level available). Note that in the case in which a 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 4-digits goods-code level). The disaggregate sectoral producer-price indices are only used when a plausible goods-price index is unavailable.

To compute the input index, Δz_{jt}, used to compute Δa_{jt}, real intermediate inputs, M, 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

$$K_{jt} = \max((1-\delta)K_{jt-1} + I_{jt}, BookValue_{jt}) \tag{B10}$$

where δ is sector-specific depreciation rate (two-digit SNI92/NACE) and is computed as an asset-share weighted average between 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 firm 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.

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

where the ending point of the first recursion, K_{jT}, is used as the starting point for the second backward recursion. This is done in order to maximize the quality of the capital-stock series given that we lack a perfectly reliable starting point and the time dimension is small. Labor input, i.e. number of employees, are 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 pay-roll 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). Third, to calculate the cost shares we take total costs as approximately equal to total revenues.[17] The cost share of capital is then given by one minus the sum of the cost shares for all other factors.

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 3.016. Finally, we remove 2 percent of the observations in each tail for each of the distributions of $\Delta a_{jt}, \Delta wnulc_{jt}, \Delta wnd_{jt}, \Delta y_{jt}$ and again require the firm to be observed in spells of at least five years (since we are interested in the within firm dynamics when estimating the VAR). This has little effect on estimated coefficients, but ensure that the VAR passes diagnostic tests.

[17] 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 assumption made here.

C The SVAR: Tests, Impulse Responses and Variance Decompositions

Relying on the Arellano and Bond (1991) autocorrelation test of the differenced residual, two lags in the VAR is 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 with the second to the fourth lag of the 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. As a second precaution, we collapse the instrument set in order to avoid overfitting. That is, we impose the restriction that the relationships in the "first stage" are the same across all time periods (see e.g. Roodman, 2006, for a discussion). For all equations, the Hansen test of the overidentifying restrictions cannot reject the null that the model is correctly specified and the instruments are valid.

In figure 10 to 13 we plot the impulse responses to 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 very surprising, given that we estimate the impulse responses on $34,968$ firm/year observations across 6,137 firms (after considering all restrictions in terms of lags, instruments and differencing). Unfortunately, we have not been able to find any statistical tests of stationarity that are suitable for a setting with a short but wide panel. However, it should be clear from the figure that this issue is of little importance in the current setting. Importantly, the figure is expressed in log-levels and the clear flat, non-zero, end-segments in the responses implies that shocks do have permanent effects on the levels of the series (i.e. levels are I(1)) and that the differenced series are stationary (I(0)).

In figure 10, we trace out the impulse responses of the Solow residual, the *wnulc*, the *wnd* and output to a one standard deviation technology shock, η^a_{jt}. Technology shocks have a positive permanent effect on the Solow residual with a "normal" (i.e. one standard deviation) shock increasing the Solow residual with slightly less than ten 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

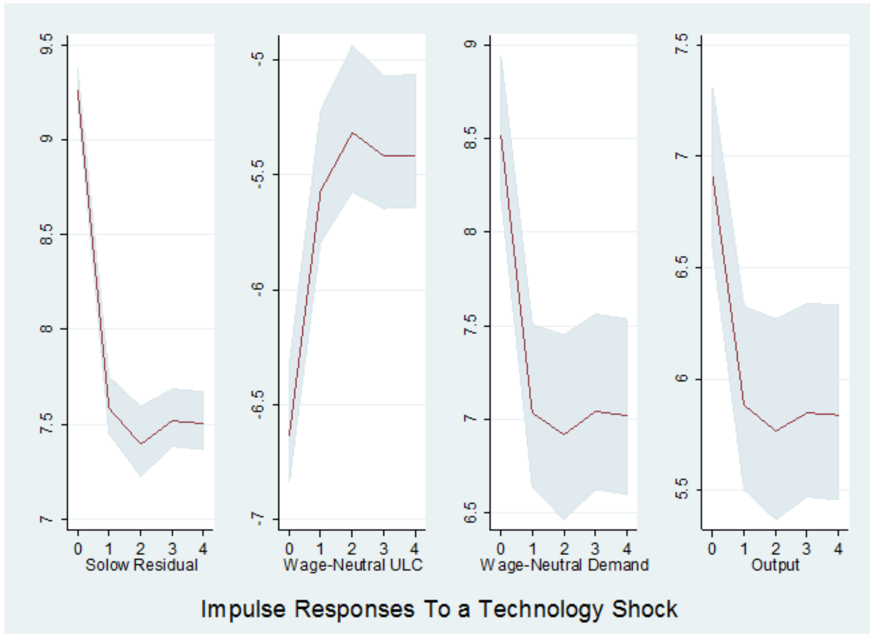

Figure 10: Impulse Responses of the baseline VAR in percentage points to a Technology Shock. Lines depict the means of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

to the (permanent) technology shock. Similarly, we find that a permanent technology shock raises wnd, as predicted from expression (A3). Moreover, output increases with about five percent in the long run.

In figure 11, we report the impulse responses to a one standard deviation (permanent) factor-price shock. A "normal" factor-price shock increases the $wnulc$ permanently 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 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.

In figure 12 we plot the impulse responses to a permanent demand shock. As shown in the figure, wnd is permanently increased in response to a permanent demand shock. In the short run, demand

Figure 11: Impulse Responses of the baseline VAR in percentage points to a Factor Price Shock. Lines depict the means of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

Figure 12: Impulse Responses of the baseline VAR in percentage points to a Demand Shock. Lines depict the means of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

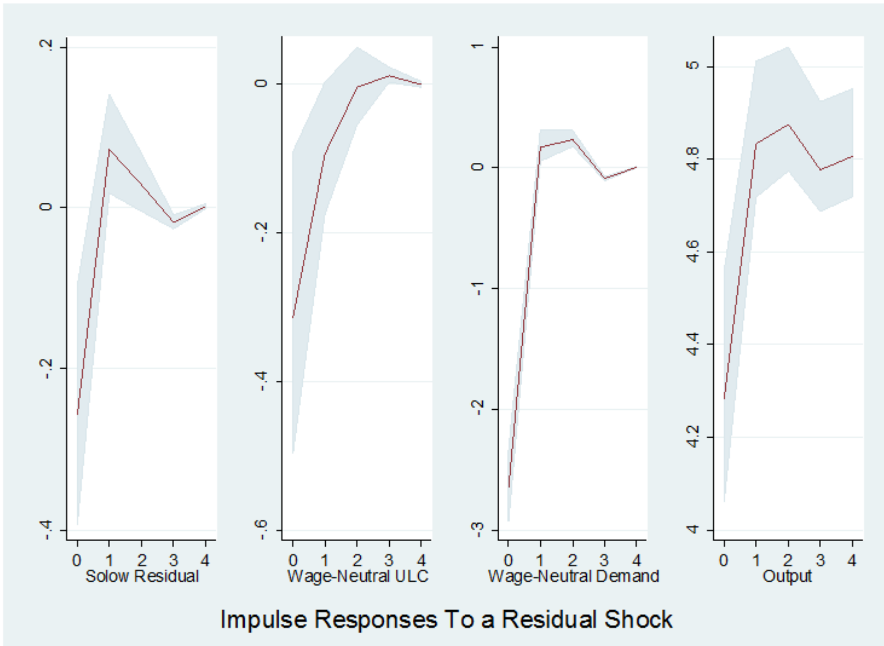

Figure 13: Impulse Responses of the baseline VAR in percentage points to a Residual Shock. Lines depict the means of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

shocks increases the Solow residual and reduces *wnulc*. A demand shock also has permanently positive effect on output as expected, increasing by about 10 percent.

For completeness, figure 13 reports the responses to the residual shock, which raises output permanently by slightly more than five percent.

Figure 14 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 (80% of the variation) and partly by technology shocks (20%). Demand shocks explain about 70% of the movements in *wnd*, whereas factor-price shocks explains about 15%. We also see in the figure that there is a role for technology shocks in explaining wage-neutral demand movements, it account for about 15%. For output, we see that about 55% of the variation is driven by demand shocks, and the rest being explained by

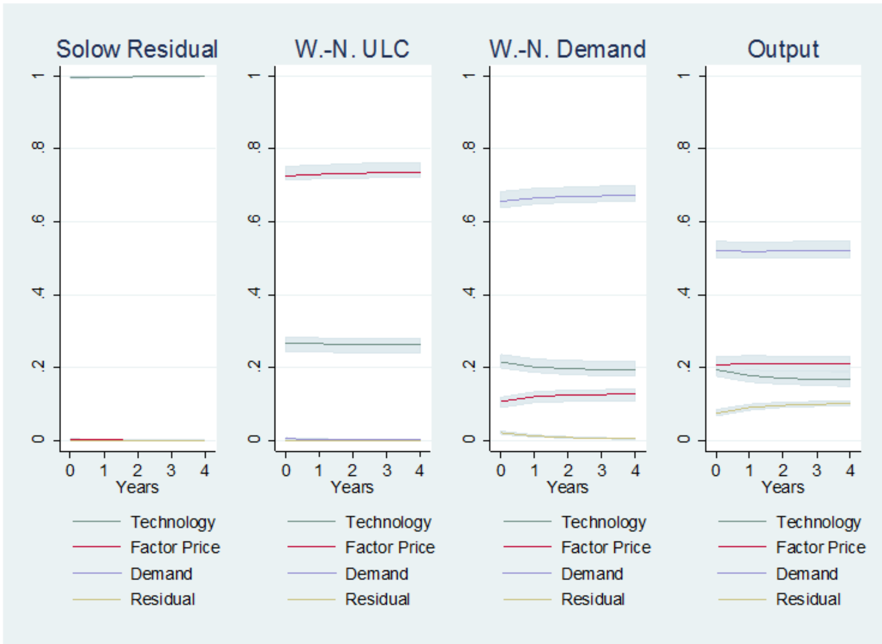

Figure 14: Forecast-error variance decompositions of the VAR in levels. The left-most panel show the percentage of the forecast-error variance in the Solow residual that can be explained by each structural shock at different horizons. Lines depict the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95-percent confidence intervalls calculated from 1000 replications.

factor price shocks (about 20%), technology (about 15%) and the transitory shock (about 10%). Overall though, we find the transitory shock to be of little importance. Given that we include time dummies in the VAR, this finding is in line with the result of Franco and Philippon (2007), who find that transitory shocks is not very important on the firm level, but account for most of the volatility of aggregates since they are correlated across firms.

Figure 15 shows the distributions for extracted innovations to technology and demand. As can be seen in the two panels of the figure, neither the demand nor the technology shock distributions are particularly skewed (skewness coefficients of -0.02 and 0.15, respectively), whereas they are both leptokurtic (kurtosis coefficients of 5.76 and 4.67). This is also clearly visible in the graphs where the dashed line depicts a normal distribution, and a standard skewness/kurtosis test (D'Agostino,

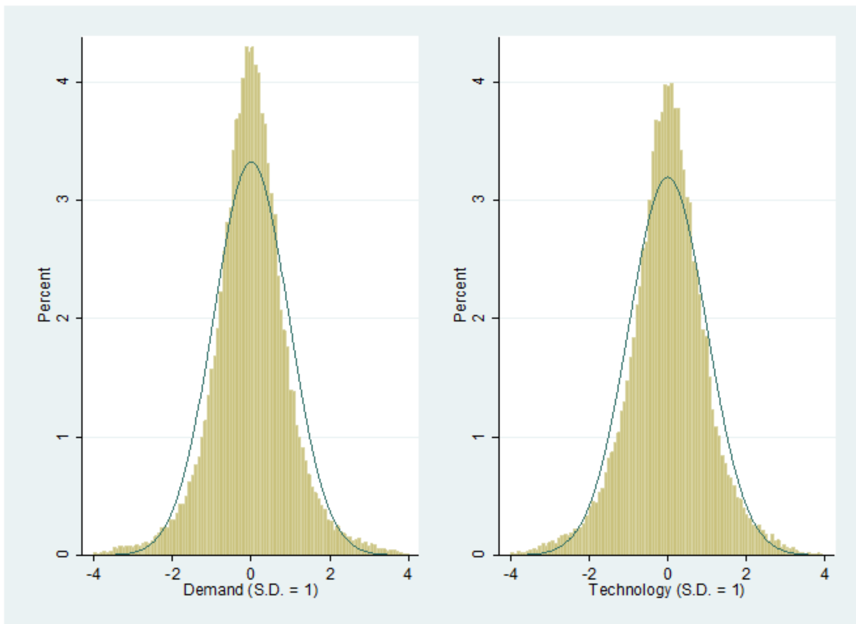

Figure 15: Histograms of the demand- and the technology shock distribution. Dashed line depicts a normal distribution.

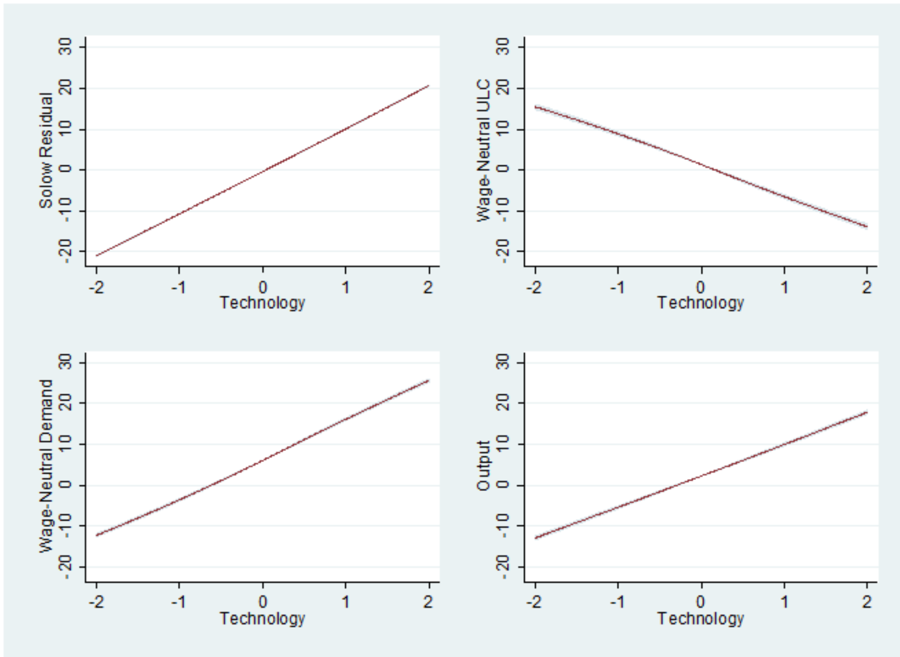

Figure 16: Contemporaneous response of variables included in the baseline VAR in percentage units as a (non-linear) function of an x-standard deviation Technology shock. Shaded areas depict the 95-percent confidence intervalls.

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 15 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 units, respectively).

A maintained assumption in the analysis is that the baseline VAR is linear in the structural shocks. In figures 16 and 17 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 can be seen in the figures, the results do support the maintained linearity assumption.

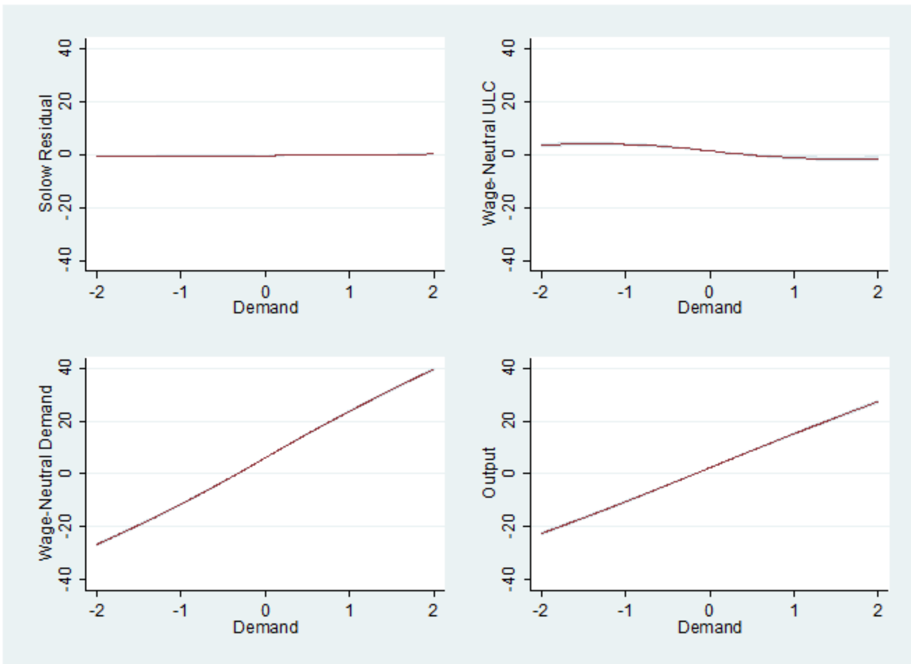

Figure 17: Contemporaneous response of variables included in the baseline VAR in percentage units as a (non-linear) function of an x-standard deviation Demand shock. Shaded areas depict the 95-percent confidence intervalls.

D Robustness: Alternative Residual Shocks

In table 5 we show results from regressions using the log employment as the dependent variable in regressions where the extracted shocks serve as independent variables. The different columns show what happens if we include different variables as the fourth variable in the systems before extracting the shocks. Column (1) show our baseline specification where we use output as the fourth variable. Columns (2) to (3) show estimates when relying on sales per worker, annual employment (from IS) and end of the year employment (from RAMS) as the fourth variable. Overall, the picture is very close to that of the main text. The lower panel shows the long-run impact which displays a somewhat larger impact of technology, but the effect remains muchs smaller than for demand.

Table 5: 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)
		Contemporaneous Effect		
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	40,451	40,284	38,213	37,234
Firms	6,125	6,113	5,879	5,703
		Long-Run Effect		
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
S.d. η_a	10.06	9.980	9.971	9.964
S.d. η_ω	16.18	16.39	15.35	15.13

Effect of one s.d. shock. Robust standard errors in parenthesis. Regression includes time dummies and firm fixed effects. Long-run estimates is the sum of the contemporaneous impact and one lag.

E Robustness: Returns to Scale and Sectoral Processes

In table 6 we first allow for specific estimates of the demand elasticity (σ) for each two digit industry. We then add variations in the assumed returns to scale, allowing for both increasing and decreasing returns. We then control for sector time dummies, and finally reestimate the entire SVAR for each two digit industry. The number of observations is slightly smaller for the sectoral estimates since some sectors are too small for us to get any sensible precision within these, none of the differences in estimates depend on the sample, however. Although magnitudes, signs and significance levels of the estimates vary across columns, the main quantitative results that demand is the primary driver of employment adjustment remains robust. Even in column (4), where we allow for substantially decreasing returns which increases the estimated importance of technology, is the estimated impact of a normal demand shock on employment more than 6 times as large as the impact of a normal technology shock.The lower panel, shows the long-run impact. As before the impact of technology is larger here throughout, but significantly smaller than for demand.

In the second table below we instead show robustness to cases where we allow the full process to be sector specific by estimating the VAR by sector. For this to be meaningful, we ned to exclude the smallest sectors, and we therefore also show estimates from the pooled model for this sample. Again, we see very small differences in the results.

Table 6: Contemporaneous and Long-Run Effect on Log Employment - Sectoral Heterogeniety

	(1) Baseline	(2) Sectoral σ	(3) RTS = 1.1	(4) RTS = 0.9
		Contemporaneous Effect		
η_a	0.153	0.192	-0.492**	0.95505**
	(0.159)	(0.161)	(0.149)	(0.161)
η_ω	5.986**	5.693**	5.541**	6.14867**
	(0.233)	(0.221)	(0.223)	(0.233)
Observations	40,451	40,214	39,788	41,132
Firms	6,125	6,102	6,065	6,193
		Long-Run Effect		
η_a	0.504	0.599	-0.244	1.378
	(0.214)*	(0.214)**	(0.232)	(0.211)**
η_ω	6.357	5.996	5.978	6.313
	(0.310)**	(0.302)**	(0.301)**	(0.310)**
Observations	34,414	34,198	33,811	35,031
Firms	6,116	6,094	6,055	6,184
S.d. η_a	10.06	10.03	10.37	10.04
S.d. η_ω	16.18	17.09	13.45	18.74

Effect of one s.d. shock. Robust standard errors in parenthesis. Regression includes firm fixed effects and time dummies. Long-run estimates is the sum of the contemporaneous impact and one lag.

F Robustness: A Two-Period Model, Allowing for Firm Exit

Our baseline model ignores the role of firm exits. As shown in the paper, demand shocks appears to be more important than technology also in this dimension, in particular when shocks become sufficiently large. Our baseline strategy is to relate yearly shocks to the end of the year employment, but if the firm exit during the year, we are unable to measure the shocks. In order to see if this affects the results, we have reestimated the model, taking a two period approach. In practice this implies that we relate the shock to the labor flows across two years. Since the labor flows are defined even if all workers exit, we are able to calculate the impact of the shocks while excluding ad including the firms that exit during the year after the shock. Since the relationship between the shocks and firm exit appear to be non-linear, we perform this robustness analysis in the non-linear framework of section 4.3. The results, indicate that the results are completely insensitive to how we treat the exiting firms, see the table below .

Table 7: Contemporaneous and Long-Run Effect on Log Employment - Sectoral dynamics and sector by time dummies

	(1)	(2)	(3)	(4)
	Contemporaneous Effect			
η_a	0.153	0.192	0.147	0.126
	(0.159)	(0.161)	(0.162)	(0.162)
η_ω	5.986**	5.693**	5.520**	5.506**
	(0.233)	(0.221)	(0.222)	(0.225)
Observations	40,451	40,214	39,580	39,580
Firms	6,125	6,102	5,997	5,997
	Long-Run Effect			
	0.504*	0.599**	0.510*	0.490*
	(0.214)	(0.214)	(0.217)	(0.221)
	6.357**	5.996**	5.811**	5.737**
	(0.310)	(0.302)	(0.306)	(0.302)
Observations	34,414	34,198	33,667	33,667
Firms	6,116	6,094	5,989	5,989
Sectoral Sigma	No	Yes	Yes	Yes
Pooled Dynamics	Yes	Yes	Yes	No
Large Sectors Only	No	No	Yes	Yes
Sector by Time Dummies	No	No	Yes	Yes
S.d. η_a	10.06	10.03	9.94	9.88
S.d. η_ω	16.18	17.09	16.98	16.80

Effect of one s.d. shock. Robust standard errors in parenthesis. Regression includes firm fixed effects and (columns 1-2) time dummies or ((columns 3-4)) two-digit NACE sector by time dummies (columns 3-4). Long-run estimates is the sum of the contemporaneous impact and one lag.

G Robustness: Data window

The allocation of output across plants within multi-plant firms during parts of our sample period (after 1996) is imputed. We have therefore redone the analysis for the sample of single plant firms, as well as for a mixed sample including multi-plant firms until 1996, but not thereafter. The results, presented below, are completely robust to these alterations of the sample.

Table 8: Contemporaneous and Long-Run Effect on Log Employment - using outcomes over one and two periods, with and without firm exits in the second period.

	(1) One Step	(2) Two Step - No Zeros	(3) Two Step Zeros
		Contemporaneous Effect	
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 Effect	
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

Effect of one s.d. shock. Robust standard errors in parenthesis. Regression includes firm fixed effects and time dummies. Long-run estimates is the sum of the contemporaneous impact and one lag.

H Size and Employee Heterogeneity

A possible reasons for why firms rely so heavily on separations rather than refraining from hiring when faced by a negative shock, is that they may be constrained in terms of the composition of the workers. Even though some workers always leave the firm, so that the firm potentially could shrink by not recruiting, it may be the wrong workers who leave. To shed some further light on this issue we have analyzed small (less than 20 employees) and large firms separately, as well as separated the firms according to the dispersion of educational fields and level within the firm.

The idea behind the second exercise is that firms with a more homogenous workforce should care less about who leaves, whereas firms with a more heterogeneous workforce should be more prone to hire and separate at the same time when hit by a negative shock. In practice, we calculate the fraction of coworkers (to each worker in the data) that has the exact same type of education (3-digit field and 2-digit level) and take the average of this share for each firm. This gives an

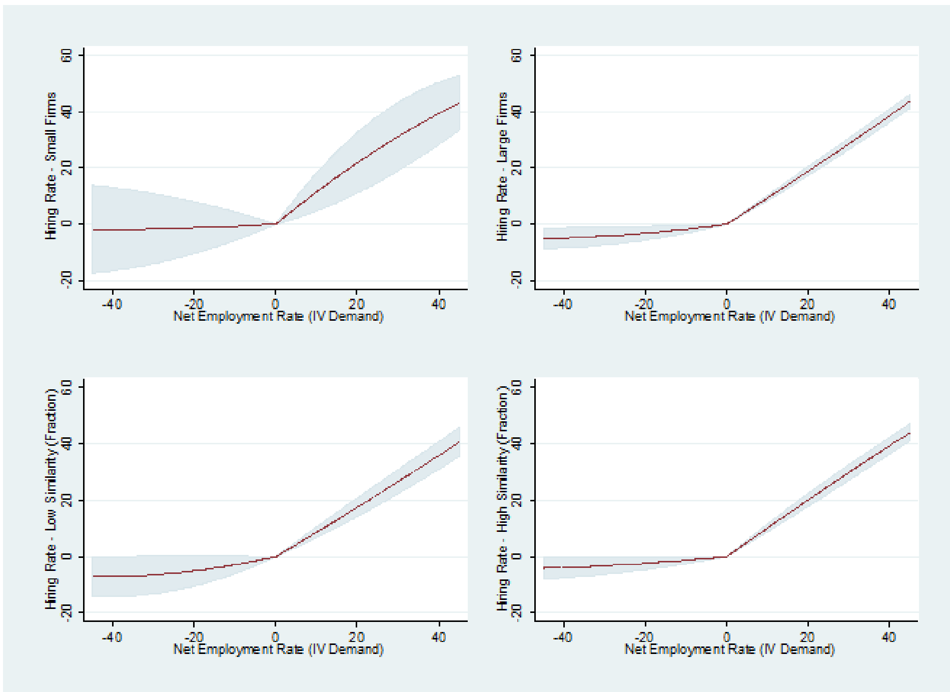

Figure 18: Contemporaneous Separation Rates in percentage units as a (non-linear) function of an x-standard deviation Demand shock in subsamples defined by employee heterogeneity (upper graphs) and firm size (lower graphs). Shaded areas depicts 95-percent confidence intervalls.

Table 9: Contemporaneous and Long-Run Effect on Log Employment - Sample Variations

	(1)	(2)	(3)	(4)
Sample	Baseline	Single Plant Always	Single Plant After 1996	$\leq \pm 2$ S.d. Shocks
		Contemporaneous Effect		
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 Effect		
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
S.d. η_a	10.06	9.13	9.41	10.06
S.d. η_ω	16.18	15.07	14.79	16.18

Effect of one s.d. shock. Robust standard errors in parenthesis. Regression includes firm fixed effects and time dummies. Long-run estimates is the sum of the contemporaneous impact and one lag.

index of how exposed the average worker within each firm is to similarly trained workers. This procedure is a straightforward implementation of standard practices when measuring segregation, see e.g. **?**. In a second step, we split our firm-level data across the median of this index, and analyze the two samples separately. Figure 18 shows the results for the impact of instrumented employment adjustments on hires.Quite surprisingly, we find very little to support the notion that size or within-firm heterogeneity is an important explanation for the heavily reliance on separations (rather than a reduction in hires) when firms are hit by negative demand shocks.