Firm-Level Shocks and Labour Flows[*]

Mikael Carlsson, Julián Messina, and Oskar Nordström Skans

May 27, 2020

Abstract

We analyse how labour flows respond to permanent idiosyncratic shifts in firm-level production functions and demand curves using very detailed Swedish micro data. Shocks to firms' physical productivity have only modest effects on firm-level employment decisions. In contrast, we document rapid and substantial employment adjustments through hires and separations in response to firm-level demand shocks. The choice of adjustment margin depends on the sign of the shock: firms adjust through increased hires if these shocks are positive and through increased separations if the shocks are negative.

Keywords: Technology, Demand, Job Flows, Worker Flows

JEL classifications: D22, J63, J23

1 Introduction

About one in every five jobs is created or destroyed every year (Davis, Faberman, and Haltiwanger 2012). The bulk of this firm-level labour adjustment is truly idiosyncratic, as firms operating in the same sector and area shrink and grow side-by-side. Hence, jobs are rapidly created and destroyed, even in sectors with stable net employment. Following the seminal work of Davis, Haltiwanger, and Schuh (1996), the importance and magnitude of these job flows have been documented for a large number of countries. Similarly, for every job created or destroyed at the firm level, there is typically a larger number of worker hires and separations. Following in the tracks of Abowd, Corbel, and Kramarz (1999), the relationship between employment adjustments of different signs and magnitudes on one side, and worker flows on the other side, has been thoroughly investigated in several studies.[1] However, while the empirical regularities of job and worker flows have been abundantly documented, little is known about how job and worker flows respond to structural firm-level shocks.[2] In this paper, we use detailed Swedish register data to show that permanent shocks to firms' idiosyncratic product demand are a much more important source of job reallocation between firms than idiosyncratic technology shocks. We further show that the job reallocation induced by these permanent demand shocks generates excessive worker flows, because firms adjust to negative demand shocks through additional separations, instead of cutting down on hires.

We follow Foster, Haltiwanger, and Syverson (2008) and let technology and demand shocks summarize the set of idiosyncratic disturbances that may alter firms' demand for labour inputs. Technology shocks are defined as shifts in the firms' ability to produce at a given level of inputs (i.e., shifts in the firms' physical production function), whereas demand shocks are defined as shifts in the firms' ability to sell at a given price (i.e., shifts in the firms' demand curve). To generate empirical measures of permanent technology and demand shocks, we derive a set of conditions arising from a stylized model of monopolistically competitive firms. To avoid imposing un-

---

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

[2]A small macro-oriented literature aims at identifying the employment responses to technology-driven changes in firm-level productivity: see, for example, Carlsson and Smedsaas (2007) and Marchetti and Nucci (2005). The macro literature also contains several related studies, such as Galí (1999) and Michelacci and Lopez-Salido (2007), the latter of which distinguishes between neutral technology shocks and investment-specific technology shocks and derives the consequences for job reallocation.

necessary restrictions on firms' short-run adjustment behaviour, we only assume that our derived conditions are valid in the long run. Therefore, inspired by Franco and Philippon (2007), we make use of structural vector autoregression (SVAR) methods, as originally outlined in Blanchard and Quah (1989), for estimation of the shocks. Using the SVAR allows us to filter out empirical measures of permanent idiosyncratic demand and technology shocks from the observed data without any assumptions about short-run dynamics. The focus on permanent shocks is motivated by the findings in Franco and Philippon (2007); Roys (2016); and Guiso, Schivardi, and Pistaferri (2005), suggesting that workers are insulated from transitory idiosyncratic shocks, a feature that is echoed in our empirical application.[3]

The most important imposed restriction is that demand shocks cannot affect the physical gross Solow residual in the long run. The long-run aspect of this restriction is crucial, since it implies that demand shocks, changes in factor utilisation, or inventories are allowed to have a transitory impact on the physical Solow residual without affecting our measured technology shocks. We further derive sufficient restrictions to identify permanent demand shocks without imposing any restrictions on the nature of short-run shocks or dynamics but explicitly allowing for shocks to factor prices, which are likely to be important in a small open economy.

When taking the analysis to the data, we benefit from detailed Swedish register data covering the universe of workers and manufacturing firms with at least 10 workers during a 12-year period. A key aspect of our data is that they contain a firm-specific price index, which allows us to derive measures of firm-level real output volumes, which are needed to separate technology shocks from demand shocks. The data further allow us to mitigate standard macro-data concerns about the practical implementation of SVARs arising from imprecisely estimated parameter vectors and covariance matrices of the underlying set of reduced-form equations. In our case, we estimate these components using dynamic panel data methods, building on Arellano and Bond (1991), which allow us to use cross-sectional variation for identification of the crucial parameters. This provides substantial gains in power relative to standard time-series applications, including, for example, Franco and Philippon (2007), who

---

[3]Demand shocks have a nontrivial transitory component, which we abstract from in the main analysis and then study separately in section 4.2. In contrast, the bulk of movements in the Solow residual are persistent enough to emerge as permanent shocks in our SVAR. This is consistent with Carlsson, Messina, and Nordström-Skans (2016), who, when estimating an AR(1) process for the level of technology using Swedish data similar to ours, find a persistence estimate as high as 0.88. Eslava et al. (2004) find an even higher persistence of 0.92 for Colombia.

2

estimate similar processes using time-series variation within firms.

The empirical analysis provides two important new insights. The first is that the main driving force behind employment adjustments is changes in labour demand that arise through permanent shocks from the product demand side, and not from cost-saving technology shocks. Firm-level technology shocks have a relatively limited effect on labour inputs, despite large effects on firm-level prices and output. In contrast, we find that permanent idiosyncratic variations in product demand have a major impact on firm-level labour adjustments. Our preferred estimates indicate that a permanent idiosyncratic demand shock of one standard deviation increases employment by 6 percent, whereas the corresponding number for technology shocks is 0.5 percent. These results are robust to a wide range of variations in measures and variations in the empirical approach.[4] We also find that most of the adjustment takes place within a year, which implies that the short- and long-run employment adjustments induced by permanent demand shocks are similar in magnitude. In contrast to the responses to these permanent demand shocks, we show evidence suggesting that the firms' responses to transitory demand shocks are more muted.

Our second main result is that the responses of worker flows (hires and separations) are asymmetric, i.e., they depend on the sign of the underlying shock. Here we build on the seminal work of Abowd, Corbel, and Kramarz (1999) and Davis, Faberman, and Haltiwanger (2012), who provide descriptive decompositions of how worker flows are related to employment changes at the firm level. In contrast to these decomposition exercises, we analyse how hires and separations respond to job creation or destruction induced by permanent shifts in a firm's product demand schedule. We find that, while most of the response to permanent positive shocks is through increased hirings, by far most of the adjustment to permanent negative shocks is through increased separations, and not through reduced hirings. Firms continue to recruit workers at a rate that is not far off the average hiring rate, even while they are destroying jobs in response to a permanent fall in product demand. These results thus concur with the descriptive picture provided by Davis, Faberman, and Haltiwanger (2012) for the United States, but differ from Abowd, Corbel, and Kramarz (1999), who document

---

[4]The fact that employment responds so much more forcefully to permanent demand shocks relative to technology shocks is difficult to reconcile with a constant product-demand elasticity. However, we show that the magnitudes can be reconciled with models (consistent with our identifying assumptions) where the demand elasticity moves with the structural shocks. Furthermore, the implied sensitivity of the elasticity is modest in magnitude.

3

that employment reductions in French firms are primarily associated with reduced hiring rates.

Our analysis adds to a vibrant empirical literature, surveyed by Syverson (2011), who documents the distinct impacts of firm-level technology and demand shocks on productivity and other firm-level outcomes. Most notably, Foster, Haltiwanger, and Syverson (2008) show that firm closures are driven primarily by changes in idiosyncratic demand and only to a lesser extent by changes in idiosyncratic physical productivity. Eslava and Haltiwanger (2018) study the contribution of economic fundamentals, most notably technology and demand differences across firms, versus adjustment frictions to output and sales growth of Colombian establishments. Foster, Haltiwanger, and Syverson (2016) show that the growth of young firms in the United States is due to a shrinking product-demand gap relative to incumbents. Pozzi and Schivardi (2016), who use Italian data to analyse how technology and demand affect firm output, show that firm-level technology shocks have a surprisingly low impact on firm growth, and that demand shocks are at least as important. Studies analysing employment adjustments in response to shocks also include Caballero, Engle, and Haltiwanger (1997) and Eslava et al. (2010). In addition, Carlsson, Messina, and Nordström-Skans (2016) show that firm-level technology shocks affect workers' wages, using Swedish data.

This paper is the first to show how firm-level technology and demand shocks affect firms' labour adjustments through hires and separations in response to shocks of different nature, signs, and magnitudes. We believe that we are the first to show that firms reduce their labour input though increased separations rather than through reduced hires when hit by a permanent negative idiosyncratic shock. Our finding that transitory demand shocks have a much more limited impact on employment adjustment than permanent demand shocks is, however, fully in line with Guiso, Schivardi, and Pistaferri (2005), who show that firms insure workers' wages relative to transitory (but not permanent) shocks to value added. Perhaps most importantly, we believe that our reduced-form evidence on the importance of idiosyncratic product-demand shocks for worker reallocation should be directly relevant for the theoretical literature on the relationship between firm-level revenue productivity and labour adjustments (e.g., Bentolila and Bertola 1990; Davis and Haltiwanger 1992; Hopenhayn and Rogerson 1993; Mortensen and Pissarides 1994; Cahuc, Postel-Vinay, and Robin 2006; Lise, Meghir, and Robin 2016). This literature tends to emphasize technology shocks as

4

the key driving force behind labour adjustments, but our results instead suggest that a careful modelling of shocks to the product-market environment is a viable way forward in order to provide a better understanding of the process where workers are reallocated across firms.

The paper is organized as follows. Section 2 outlines a simple model that motivates the long-run restrictions needed to extract our permanent demand and technology shocks. Section 3 introduces the main characteristics of the firm-level data used in the analysis and discusses the empirical implementation of the SVAR and validation of the shocks. Section 4 reports the results on the relationships between net employment, hires, and separations with technology and demand shocks. Section 5 documents how hires and separations respond to positive and negative shocks. Finally, section 6 concludes. Appendices, containing background information and additional results, are published online.

2   Model and Empirical Strategy

2.1   Long-Run Model with Permanent Shocks

In this section, we outline a stylized model of monopolistically competitive firms that allows us to identify two exogenous idiosyncratic driving forces of firms' relative performance: *technology shocks,* which affect firms' physical productivity, and *demand shocks,* which affect firms' ability to sell their products at a given price. The purpose of the paper is to analyse how these two disturbances affect firms' hiring and separation policies. Our approach only requires that we define a set of restrictions that are valid in the long run. The model is therefore deliberately stylized and ignores many factors that may be important for short-run dynamics.[5]

---

[5]The key distinction between our technology shocks and demand shocks lies in how the shock affects the producing firm, not in the origin of the shock. This approach, which is consistent with the existing (micro) literature (such as Foster, Haltiwanger, and Syverson (2008), and Syverson [2011]), implies that we do not distinguish between shifts in the firm-specific demand curve that arise from changing preferences among final consumers, those that arise from increased demand among downstream firms, and those that arise from quality changes that increase product demand at a given price. Franco and Philippon (2007) label these as shocks to market shares, and model them formally as preference shocks. The firm-level price index we use is based on unit prices for very detailed product codes (8/9-digit Harmonized System/Combined Nomenclature codes), which limits the scope for quality changes to be the key component in our demand shock. However, it is straightforward to show that if we added a quality shock to the system developed below (through a wedge between the measured firm-level price, based on unit values, and the quality-adjusted price),

To identify firm-level structural shocks, we need to make assumptions about the technology and market conditions faced by the firm. Our setup follows Eslava et al. (2004) and Foster, Haltiwanger, and Syverson (2008, 2016) closely, by using a first-order approximation of production technologies and product market demand and modeling the key technology and demand shocks as neutral shifters of the production function and demand curve, respectively. Following these papers, the firm-level production function is approximated by:

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

where physical gross output Y_{jt} in firm j at time t is produced using technology indexed by A_{jt} and combining labour input N_{jt}, capital input K_{jt}, and intermediate production factors (including energy) M_{jt}. Importantly, our data allow us to account for idiosyncratic firm-specific price changes, so that our measure of technology (the Solow residual, A_{jt}) refers to *physical* total factor productivity (TFPQ), rather than *revenue* total factor productivity (TFPR) in the terminology of Foster, Haltiwanger, and Syverson (2008). Equation (1) presupposes a constant returns technology, which is our baseline assumption. However, in robustness exercises we relax this assumption.

The baseline representation of the firm-level demand curve is a constant-elasticity function

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

where P_{jt}/P_t is the firm's relative price, Y_t denotes aggregate market demand, and Ω_{jt} is a firm-specific demand shifter. The parameter σ denotes the elasticity of substitution across products and hence captures the demand elasticity for each firm in the economy. Here, we assume a constant demand elasticity, but below we show that our identification remains valid if we treat σ as a function of the shocks, that is, if $\sigma = \sigma(A_{jt}, \Omega_{jt})$, allowing for Kimball-style (1995) strategic complementarity in price setting.

Following Guiso, Schivardi, and Pistaferri (2005) and Franco and Philippon (2007), we model the key shocks as permanent shifters. Formally:

$$A_{jt} = A_{jt-1} e^{\mu_j^a + \Phi^a(L)\eta_{jt}^a}, \tag{3}$$
$$\Omega_{jt} = \Omega_{jt-1} e^{\mu_j^\omega + \Phi^\omega(L)\eta_{jt}^\omega}, \tag{4}$$

it would enter the system symmetrically to the demand shock.

6

where μ_j^a and μ_j^ω are constant drifts, and $\Phi^a(L)$ and $\Phi^\omega(L)$ are polynomials in the lag operator, L. The white noise idiosyncratic (orthogonal) technology and demand shocks are denoted by η_{jt}^a and η_{jt}^ω. The assumed functional form implies that the shocks' lag polynomials are linearly related to the log differences of A_{jt} and Ω_{jt}, respectively.[6] As is evident from the formulation, our focus is on permanent shocks, but in a variation of the model we also explicitly analyse the role of transitory disturbances (see section 4.2).

Our model also allows for a long-run impact of shocks to factor prices other than labour. This is potentially important in the Swedish setting of a small open economy where factor prices are likely to vary across sectors and time (e.g., due to exchange rate volatility). To simplify the notation, we define a price index (consistent with cost minimization) for input factors other than labour, $P_{jt}^F = \left(P_{jt}^K / \beta \right)^\beta \left(P_{jt}^M / \left(1 - \alpha - \beta \right) \right)^{1-\alpha-\beta}$, where P_{jt}^K is the price of capital and P_{jt}^M is the price of intermediate materials at time t. Since the empirical specification will include time fixed effects, effectively removing the impact of average market prices, P_{jt}^F will only capture idiosyncratic firm-specific disturbances in factor prices, which, in turn, are assumed to follow a stochastic process of the same form as the demand and technology shocks. Thus, P_{jt}^F evolves according to

$$P_{jt}^F = P_{jt-1}^F e^{\mu_j^f + \Phi^f(L)\eta_{jt}^f}, \tag{5}$$

where μ_j^f is a firm-specific drift; $\Phi^f(L)$ is a polynomial in the lag operator, L; and η_{jt}^f is a white noise factor price shock.

2.2 Identifying Permanent Shocks

2.2.1 Recursive Set of Long-Run Restrictions

As our baseline strategy to identify the shocks of interest, we use the motivating model discussed above to derive a set of recursive relationships regarding how our structural shocks can affect observable variables in the long run. The system allows us to incorporate explicit shocks to factor prices and neutralize disturbances through potential wage shocks.[7]

We first note that the assumptions of the model ensure that the only shock that

---

[6]This, in turn, provides a convenient representation of the vector autoregression specified below.
[7]In section 2.4, we discuss alternative strategies to identify the shocks that also are consistent with the same motivating model. Reassuringly, our results do not depend on which strategy we use.

7

can affect the physical gross output Solow residual (A_{jt}) is the technology shock. We then 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}. Denoting wages by W_{jt}, marginal cost in optimum is

$$MC_{jt} = A_{jt}^{-1} \left(\frac{W_{jt}}{\alpha} \right)^\alpha P_{jt}^F. \tag{6}$$

In equilibrium, $MC_{jt} = (W_{jt}N_{jt})/(\alpha Y_{jt})$. This expression, together with equation (6), gives

$$(W_{jt}N_{jt}/Y_{jt})W_{jt}^{-\alpha} = \alpha^{1-\alpha}A_{jt}^{-1}P_{jt}^F. \tag{7}$$

Expression (7) will be affected by technology and factor price shocks but not demand shocks. Any direct shocks to the firm-level wage-setting relationship (such as changes in the degree of competition over similar types of labour) will not drive this expression. As the left-hand side of equation (7) shows, this is essentially a measure of unit labour cost ($W_{jt}N_{jt}/Y_{jt}$) net of wage-setting disturbances. We therefore refer to this variable as the wage-neutral unit labour cost ($WNULC_{jt}$).[8]

Using the demand equation (2) and expression (6), we arrive at

$$Y_{jt}W_{jt}^{\sigma\alpha} = \psi Y_t P_t^\sigma A_{jt}^\sigma \left(P_{jt}^F \right)^{-\sigma} \Omega_{jt}, \tag{8}$$

where $\psi = (1/\alpha)^{-\sigma\alpha} (\sigma/(\sigma - 1))^{-\sigma}$. Thus, expression (8) will be driven by shocks to technology, factor prices other than labour, and demand (apart from aggregate factors, which will be captured by time dummies in the empirical implementation of the model). In effect, expression (8) is output changes adjusted for wage-setting disturbances. We refer to this expression as the wage-neutral demand (WND_{jt}), to highlight that this expression helps us recover the permanent demand shocks.

Table 1, column 1, summarizes the derived recursive system of equations that can all be constructed from our firm-level data as long as we have an estimate of demand elasticity σ (as detailed in the next section). Column 3 summarizes the three key restrictions on which we rely for identification:

1. The measured physical Solow residual ($TFPQ$ in the terminology of Foster et al. 2008) equals A and hence is independent of demand, Ω, and factor prices,

---

[8]We also experimented with replacing unit labour cost with unit materials cost in equation (7), and refer to section 4.3 for a discussion of the results. We thank one of the referees for this suggestion.

P^F).[9]

2. The "wage-neutral unit labour cost" ($WNULC$), as defined in the second row, is a function of A and P^F.

3. The "wage-neutral demand" (WND), as defined in the third row, is a function of A, Ω, and P^F.

Table 1: *Core Structural VAR Equations*

	(1)		(2)	(3)
Variable:	Measured in data as:		Model expression:	Long-run restrictions:
Solow	$Y_{jt}\left(N_{jt}^{\alpha}K_{jt}^{\beta}M_{jt}^{1-\alpha-\beta}\right)^{-1}$	$=$	A_{jt}	Independent of η^{ω} and η^f
WNULC	$\left(W_{jt}N_{jt}/Y_{jt}\right)W_{jt}^{-\alpha}$	$=$	$\alpha^{1-\alpha}A_{jt}^{-1}P_{jt}^F,$	Independent of η^{ω}
WND	$Y_{jt}W_{jt}^{\sigma\alpha}$	$=$	$\psi Y_t P_t^{\sigma} A_{jt}^{\sigma}\left(P_{jt}^F\right)^{-\sigma}\Omega_{jt}$	–

Note: *Solow* = physical Solow residual (TFPQ), *WNULC* = wage-neutral unit labour cost and *WND* = wage-neutral demand, $\psi = \left(\frac{1}{\alpha}\right)^{-\sigma\alpha \cdot}\left(\frac{\sigma}{\sigma-1}\right)^{-\sigma}$.

We will also include a fourth residual variable in the system to soak up all remaining transitory dynamics, as in Franco and Philippon (2007). This prevents any other (residual) transitory dynamics from being mechanically loaded into the demand shock in the third row without changing the long-run properties of the system, since the shock is only allowed to have long-run effects on itself. As a baseline we will use output as the fourth residual variable, but a set of robustness exercises (discussed in appendix E) shows that the choice of this variable is irrelevant for the main results in the paper.

2.2.2 Implementation of the Long-Run Restrictions

We follow Blanchard and Quah (1989) to implement the long-run restrictions and filter out the permanent shocks of interest. Here, we briefly explain how these restrictions

---

[9]This assumption only remains credible if the Solow residual is calculated from a measure of real output where nominal output has been deflated by firm-specific prices. Using instead *sector-level* price deflators (a measure often used in empirical analyses) will make output a function of firm-specific idiosyncratic prices, which themselves are likely to depend on shocks other than technology (see Carlsson and Nordström-Skans (2012) for direct evidence).

9

are implemented in practice and refer to appendix B for details. Consistent with our postulated stochastic shock processes, we can write the (log difference of the) system outlined in table 1, augmented with the additional transitory variable (output), as a function of all relevant current and past structural shocks. Thus, we have the following vector moving average (VMA) representation (using lowercase letters for logarithms):

$$\Delta \mathbf{x}_t = \mathbf{C}(L)\boldsymbol{\eta}_t, \tag{9}$$

where $\Delta \mathbf{x}_t = [\Delta a_{jt}, \Delta wnulc_{jt}, \Delta wnd_{jt}, \Delta y_{jt}]'$ denotes the first differences of our observed variables and $\boldsymbol{\eta}_t = [\eta_{jt}^a, \eta_{jt}^f, \eta_{jt}^\omega, \eta_{jt}^y]'$ denotes the set of structural shocks (with η_{jt}^y being the fourth shock associated with Δy_{jt}). The elements of the four-by-four matrix $\mathbf{C}(L)$ contain polynomials in the lag operator, L, that is, $C_{rc}(L) = \sum_{k=0}^{\infty} c_{rc}(k) L^k$. Cumulated across all lags, k, the $c_{rc}(k)$ parameters provide the long-run impact of the shocks.

Since the technology shock, η_{jt}^a, is the only shock with a long-run impact on a_{jt}, we know that $\sum_{k=0}^{\infty} c_{12}(k) = \sum_{k=0}^{\infty} c_{13}(k) = \sum_{k=0}^{\infty} c_{14}(k) = 0$. Equivalently, we know that only the technology and factor price shocks can have a long-run effect on $wnulc_{jt}$, so $\sum_{k=0}^{\infty} c_{23}(k) = \sum_{k=0}^{\infty} c_{24}(k) = 0$. Finally, since the residual shock has no long-run effects on wage-neutral demand, it follows that $\sum_{k=0}^{\infty} c_{34}(k) = 0$.

Empirically, we estimate a set of reduced-form equations where the first differences of our observed variables are regressed on the lagged vector of the same variables, that is, the vector autoregression(VAR)-representation of (9):

$$\Delta \mathbf{x}_t = \mathbf{A}(L)L\Delta \mathbf{x}_t + \mathbf{e}_t, \tag{10}$$

where the elements of the four-by-four matrix $\mathbf{A}(L)$ contain lag-polynomials, $\mathbf{C}(L)$, and $\mathbf{e}_t$ is a set of reduced-form errors.[10] To recover the structural shocks from the reduced-form errors, we can compare (9) with (10) to see that the reduced-form errors are related to contemporaneous structural shocks through the expression

$$\mathbf{e}_t = \mathbf{c}(0)\boldsymbol{\eta}_t, \tag{11}$$

where $\mathbf{c}(0)$ is the matrix of the contemporaneous parameters in the lag-polynomials $\mathbf{C}(L)$ from the VMA representation. To recover the structural shocks, we then need

---

[10]Guided by diagnostic tests, the lag length of the VAR is truncated at two in the empirical application.

to estimate $\mathbf{c}(0)$. To this end, we follow standard SVAR protocols and rely on the long-run and orthogonality restrictions together with our estimates of $\Omega = E\mathbf{e}_t'\mathbf{e}$ and $\mathbf{A}(L)$ from the VAR (see appendix B for details).[11]

What matters for precision when extracting the structural shocks from the reduced-form errors is statistical power in the estimation of the VAR parameters $\mathbf{A}(L)$ and covariance matrix Ω. In our setting, we estimate these components using the dynamic panel data methods of Arellano and Bond (1991). Since these methods rely on cross-sectional variation for identification and this dimension is large in our data, we can estimate the components with considerable precision in comparison with standard time-series applications.

2.3 Benefits of the Empirical Approach

The proposed approach offers several advantages. *First*, the zero-impact restrictions in the last column in table 1 are imposed as *long-run* restrictions, thus only restricting the accumulated effect of the shock over time, that is, the restrictions take the form $\sum_{k=0}^{\infty} c_{rc}(k) = 0$. Hence, we do not make any assumptions about short-run dynamics or transitory measurement errors. Notably, our identification of the technology shocks (η^a) is therefore consistent with changes in inventories, factor utilisation, markups, or idiosyncratic input prices altering the Solow residual, as long as these changes are mean reverting in levels, that is, as long as they do not affect the level of the Solow residual in the long run.

Second, we do not require that all aspects of the motivating model are true, even in the long run. We only require that the impact of the shocks on the three variables ($Solow$, $WNULC$, and WND) measured in table 1, column 1, does not violate the restrictions listed in column 3 of the same table. These restrictions are consistent with a wider class of models than the one proposed here.[12] A particular possible extension, which for reasons discussed in section 4.3 will turn out to be useful, is to let the elasticity of demand (and thus the markup) change in response to the shocks, which we can allow for without affecting the long-run restrictions of table 1.

Third, it is straightforward to incorporate non-constant returns to scale into the

---

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

[12]The key assumptions are the relevance of the first-order approximation of the production function, monopolistic competition, and that firms behave optimally given the former restrictions.

11

model. This is important, since the key assumption for distinguishing technology shocks from demand shocks is that technology shocks alter the physical Solow residual in the long run, whereas other shocks do not. This assumption implies that changes in the scale of operation are not allowed to alter permanently the efficiency of production as measured by TFPQ. The most straightforward reason why this assumption may prove invalid is that firms might use a production technology with non-constant returns to scale. Appendix B.4 modifies the model for the case of non-constant returns to scale. Section 4 shows that the main results in the paper are robust to variations in the returns to scale.

2.4 Comparison with Alternative Identification Schemes in the Literature

Our system approach can be compared with the static single-equation method employed by Eslava et al. (2004, 2010), Foster, Haltiwanger, and Syverson (2008, 2016), and Pozzi and Schivardi (2016). Essentially, these studies derive the Solow residual from the production function (1) and the demand shock from the demand function (2) using an estimate of σ. This approach does not separate transitory and permanent shocks and may be sensitive to transitory deviations from the assumed functional forms. However, we could embed this approach into the one taken here by running a two-variable SVAR system with the Solow residual as the first equation, $Y_{jt} / \left(N_{jt}^{\alpha} K_{jt}^{\beta} M_{jt}^{1-\alpha-\beta} \right) = A_{jt}$, and the demand shock backed out from the demand function (2), $(Y_{jt}/Y_t)(P_{jt}/P_t)^{\sigma} = \Omega_{jt}$, as the second equation. In principle, both strategies are equally valid if the motivating model is true and the data are error free, but the approaches differ in how sensitive they are to possible misspecifications in different dimensions. Compared with our baseline formulation, the two-variables alternative relaxes the assumptions about optimal firm behaviour. This comes at the expense of having to rely on price data to capture all the shocks to factor inputs and wages when calculating the demand shocks. In contrast, our strategy uses direct measures of wages and unit labour costs to purge the analysis of these input-price disturbances. We explore this two-equation system in the robustness section (4.3) and, reassuringly, the results are very similar.[13]

The standard approach in the literature is to orthogonalise the structural shocks by

---
[13]We thank one of the referees for suggesting this exercise.

12

using the Solow residual as an instrument of the demand equation when estimating σ. The orthogonalisation is also imposed in our identification scheme. However, Forlani et al. (2016) allow for a correlation between demand and technology processes and find it to be strongly negative.[14] We allow for firm fixed effects in our empirical setup, which would remove the impact of persistent decisions related to the business model of the firm. For example, the choice of market segment may involve a trade-off between demand and technological efficiency generating a negative relationship between technology and demand, This would be factored out by the introduction of firm fixed effects in the econometric model.

Hottman, Redding, and Weinstein (2016) develop a structural framework to decompose the firm size distribution in terms of the contributions of the heterogeneity of demand, product scope, marginal cost, and markup. Apart from the difference in focus, a key difference is that their structural framework only relies on price and quantity data, whereas our approach (and all the other studies referenced above), uses factor input data as well to identify technology (or marginal-cost) shocks.

Guiso, Schivardi, and Pistaferri (2005) decompose shocks to firm-level value added into permanent and transitory components under assumptions of the stochastic processes underlying the shocks. We share with them the accent on the permanent versus transitory nature of the shocks. Aside from the differences in the questions posed, Guiso, Schivardi, and Pistaferri (2005) do not have data on firm-level prices and hence cannot provide a structural interpretation of the shocks, as is possible here.

3 Data and Estimation of the Shocks

3.1 Data on Firms

Our primary data sources are the Swedish Industry Statistics Survey (IS) and the Industrins Varuproduktion (IVP). These data sets contain annual information on inputs, outputs, and firm-specific producer prices for Swedish manufacturing plants from 1990 through 2002. The data have complete coverage of all plants, except for plants with fewer than 10 employees, which are subject to random sampling. We therefore exclude all plants with fewer than 10 employees and provide robustness checks to verify that our analysis is robust to endogenous transitions below this threshold. We perform our

---

[14]See also Hottman, Redding, and Weinstein (2016) and Eslava and Haltiwanger (2018).

13

analysis at the plant level, but because about 72 percent of the observations in our sample pertain to plants that are also firms, we refer to the plants as firms.

An important part of our data is the firm-specific price index built from plant-specific unit price changes.[15] These data allow us to derive a measure of gross output that is robust to changes in relative prices across firms, similar to Eslava et al. (2004) and Smeets and Warzynski (2013). Details about data construction are presented in appendix A. The data entering the VAR model cover 6,137 firms and 53,379 firm-year observations, but since the VAR model uses lags, we can extract structural shocks for 41,105 firm-years.

Table 2: *Summary Statistics: Firm Data*

	Category	(1) Mean	(2) S.d.	(3) p(25)	(4) p(75)	(5) Firms	(6) Observations
Output growth	Overall	0.029	0.190	-0.073	0.133	6,125	40,451
	Within		0.175				
Price growth	Overall	0.022	0.069	-0.001	0.044	6,125	40,451
	Within		0.064				
ΔSolow	Overall	0.001	0.112	-0.062	0.062	6,125	40,451
	Within		0.106				
$\Delta wnulc_{jt}$	Overall	0.010	0.155	-0.077	0.097	6,125	40,451
	Within		0.147				
Δwnd_{jt}	Overall	0.064	0.224	-0.060	0.189	6,125	40,451
	Within		0.210				

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

The top two rows in table 2 describe gross output and price growth in the sample used for the final analysis. The average price change of 2.2 percent per year coincides with the average growth rate of the official producer price index for industry provided by Statistics Sweden. However, in the empirical analysis, all the aggregate trends are accounted for by time dummies (sector-by-time dummies in some robustness exercises). The table also shows the distribution statistics of the main variables used in the analysis, overall and within firms.

---

[15]The index uses Paasche-type links constructed from survey information on reported sales and volume at the goods level.

3.1.1 Variables in the VAR

Letting lowercase letters denote logs, firm-level changes in the physical Solow residual for firm j at time t are computed as follows:

$$\Delta a_{jt} = \Delta y_{jt} - \Delta z_{jt}, \tag{12}$$

where Δy_{jt} is the growth rate of real gross output, and Δz_{jt} is a cost-share-weighted input index defined as $C_k \Delta k_{jt} + C_n \Delta n_{jt} + C_m \Delta m_{jt}$. The terms Δk_{jt}, Δn_{jt}, and Δm_{jt} are the growth rates of capital, labour, and intermediate materials and energy, respectively, while C_K, C_N, and C_M are the corresponding cost shares in total costs (see details in appendix A). We approximate the cost shares by revenue shares, which should be innocuous since pure profits are small in Swedish manufacturing firms.[16] Moreover, we use industry-level averages over time and take total costs as approximately equal to total revenues.[17] Since cost/revenue shares sum to one, the share for capital is given by one minus the sum of the revenue shares of labour and materials, which we measure from the data. Using data on factor compensation, changes in output, and changes in inputs, we can thus calculate the residual Δa_{jt}, which provides an estimate of changes in the physical Solow residual. This might not accurately measure technology shocks (η^a), due to varying factor utilisation, inventories, or truly idiosyncratic factor prices, but the SVAR will filter out true technology shocks from equation (12) as long as η^a is the only factor that permanently shifts A_{jt}. Material inputs are deflated using three-digit sectoral price indices, which implies that we allow not only for an arbitrary set of transitory factor price shocks, but also for permanent input price shocks in the manufacturing sector as long as these are shared with other similar (at the three-digit level) firms. Summary statistics of the Solow residual are found in the third row in table 2.

---

[16]Our monopolistic competition model implies pure economic profits. However, similar to U.S. evidence discussed in Basu, Fernald, and Shapiro (2001), we find a very small average (1968-93) share of economic profits (−0.001) when relying on the aggregate Swedish manufacturing data from Carlsson (2003). This finding thus supports the commonly used approximation in the literature of measuring (average) cost shares by (average) revenue shares, which is also used here. For simplicity, however, we do not complicate the cost structure in our model, to accommodate explicitly the absence of economic profits in the data.

[17]Treating cost/revenue shares as constant over time is sensible given that the approach we take is not sensitive to transitory variation. And, trying to account for time variation in output elasticities is complicated, because observed factor payments might not be allocative period-by-period, for example, because of implicit contracts.

15

To compute $\Delta wnulc_{jt}$ we rely on cost minimization and use C_N as the estimate of α, letting it vary by two-digit industry. The rest of the components of $\Delta wnulc_{jt}$ are directly observed in the firm-level data. The fourth row in table 2 shows the distribution of $\Delta wnulc_{jt}$.

The computation of Δwnd_{jt}, requires an estimate of the demand elasticity σ. We obtain this by estimating the demand equation (2) using the Solow residual to instrument the firm idiosyncratic price, as in Foster, Haltiwanger, and Syverson (2008). The instrument is consistent with our initial assumptions, because the Solow residual is expected to affect firm-level sales only through firm-level prices. The results of this procedure suggest an elasticity of substitution equal to 3.306 (s.e. 0.075). The estimate of σ is well in line with standard calibration exercises (see, e.g., Erceg, Henderson, and Levin 2000) as well as Swedish micro-evidence provided by Heyman, Svaleryd, and Vlachos (2013). As robustness checks, we also show that the main results are robust to using sector-specific estimates of σ and a very wide span of assumed values of σ. The ensuing measure of Δwnd_{jt} is provided in the final row in table 2.

3.2 Data on Labour Flows

To analyse the impact of the shocks on the use of labour and flows into and out of the firms, we link a longitudinal employer-employee database (Statistics Sweden's register-based labour market statistics, or RAMS) to the firm-level data. These data are based on tax records and include the identity of all employees in each firm. Following Carlsson, Messina, and Nordström-Skans (2016) and others, we measure employment in November each year and restrict the (main) analysis to full-time employees in their main jobs.[18] In the end, we are able to match shocks and labour flows for 40,451 firm-year observations in 6,125 firms. The final sample covers nearly two-thirds of all manufacturing employees.[19] For completeness, we further study how the use of

---

[18]More precisely, our raw annual data include information on all employment spells (even with annual earnings corresponding to a few kronors) and include information on the identity of the employer, the employee, total annual earnings, as well as the first and last remunerated months during the year. Following Nordström Skans, Edin, and Holmlund (2009), and similar in spirit to Song et al. (2019), we only include employment spells where average monthly earnings exceed 75 percent of the monthly minimum wage and the spell covers November. In auxiliary exercises, we separately study the dynamics of all those employment spells that are excluded due to these restrictions.

[19]The employment data that were used to construct the variables in the VAR were obtained from a different source (IS) than the employment, hiring, and separation data used in the final analysis (which were obtained from RAMS). This insulates the analysis from the threat of joint measurement

16

marginal workers (i.e., employees who do not satisfy these criteria) changes in response to the shocks.

We measure employment in logs or, when decomposing the results into flows, we measure net employment changes using the metrics proposed by Davis, Haltiwanger, and Schuh (1996). Thus, 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 employment flow equation holds, that is, $Employment_t = Employment_{t-1} + Hires_t - Separations_t$.

Table 3: *Summary Statistics: Worker Data*

	Category	(1) Mean	(2) S.d.	(3) p(25)	(4) p(75)	(5) Firms	(6) Observations
Net employment rate	Overall	0.012	0.208	-0.062	0.089	6,125	40,451
	Within		0.195				
Hiring rate	Overall	0.150	0.151	0.063	0.200	6,125	40,451
	Within		0.127				
Separation rate	Overall	0.138	0.152	0.061	0.174	6,125	40,451
	Within		0.131				
ST separation rate	Overall	0.061	0.082	0	0.083	6,125	40,451
	Within		0.065				
Marginal net emp. rt.	Overall	0.009	0.353	-0.069	0.082	6,125	40,451
	Within		0.334				

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

We do not observe the contract type in the data, but to explore the role played by the (potential) flexibility provided by marginal workers, we use two additional flow margins. We (i) define the short-tenure (ST) separation rate as the number of separations of short-tenure (fewer than three years, constituting 23 percent of all worker-year observations) workers divided by average employment across the two years, and (ii)

errors in the calculation of the shocks and the employment adjustment analysis. However, the estimates of the impact of the shocks on overall employment are very similar using the two data sources, suggesting that the issue is of minor importance.

17

measure the change in the number of marginal workers, who are defined as individuals who are employed during the year but not in November, and hence are not included in the stock of end-of-year employees.

Descriptive statistics are presented in table 3. The average hiring rate during the observation period is 15 percent, and the average separation rate is 14 percent, of which slightly less than half (6 percent) are separations of short-tenure workers.

3.3 Estimation and Validation of the VAR

To derive the shocks of interest, we estimate an SVAR on the three variables defined in table 1, Δa_{jt}, $\Delta wnulc_{jt}$, and Δwnd_{jt}, and the fourth residual variable, Δy_{jt}. Details on the estimation, distribution of the shocks, and impulse responses are provided in appendix B. We find that the standard deviation of the demand shock is about 60 percent larger than that of the technology shock (16.2 and 10.1, respectively). The impulse responses of the four variables to the two shocks display fairly limited dynamics, generally converging to the long-run equilibrium within a year. This is particularly the case for the Solow residual, which suggests that much of the dynamics in standard measures of Solow residuals that use sectoral prices to deflate output may be due to the dynamics of idiosyncratic prices (see Carlsson and Nordström-Skans (2012) for direct evidence on relative price dynamics).

Appendix B.2 provides internal support for the interpretation of the shocks based on theory-consistent signs for the three unrestricted responses within the VAR system. The estimated VAR model does not impose any restrictions on how technology shocks affect WNULC and WND, but, as predicted from the model, WNULC falls permanently and WND increases in response to a (permanent) positive technology shock. And in line with the model, positive factor price shocks lead to a reduction in WND.

In appendix B.3, we further validate the interpretation of the derived shocks by showing that they have the expected qualitative impacts on firm-specific prices and output. From theory, we know that technology and demand shocks should affect output. The response of prices instead depends on the nature of the shock. A positive technology shock lowers the cost of production, so firms need to lower their prices to increase their sales along a fixed demand curve. Instead, demand shocks shift the firm-specific demand curve, allowing the firm to sell more at constant prices. The appendix validates these predictions: a one standard deviation technology (demand)

18

shock increases output by 6 (10) percent in the long run. Moreover, as expected, prices decrease significantly due to technology shocks, but they increase marginally when hit by demand shocks. These results are not imposed from the construction of our variables. In particular, prices could well (from a pure measurement standpoint) respond in either direction to structural innovations in technology and demand.

4 Shocks and Employment Adjustments

4.1 Main Results

Figure 1 shows impulse responses of log employment with bootstrapped confidence bands from the VAR when using log employment change as the fourth variable. The results show that idiosyncratic demand shocks have a substantially greater impact than the corresponding technology shocks on firm-level labour adjustments. A positive demand shock of one standard deviation increases employment by slightly more than 6 percentage points, whereas the impact of an equivalent technology shock has a very limited impact on employment. It is also evident from figure 1 that the dynamics of labour adjustments are fairly limited. More than 90 percent of the long-run adjustments in response to the permanent shocks occur within the first year.

The objective of our analysis is to illustrate how job and worker flows respond to permanent shifts in idiosyncratic production functions and demand curves. To explore departures from linearity and potential asymmetries in the response margins (see section 5), it is useful to extract the measures of structural shocks from the SVAR and relate them in standard regression frameworks to different outcomes. Thus, this is how we proceed in most of the analyses we present in the paper. Our point of departure is the following baseline specification:[20]

$$Outcome_{jt} = \eta_{jt}^a \delta_1 + \eta_{jt}^\omega \delta_2 + \rho_t \beta_\rho + \mu_j + \xi_{jt,} \tag{13}$$

where $Outcome$ denotes employment (or some measure of labour flows) for firm j at time t. The coefficients δ_1 and δ_2 capture the impact of the two structural shocks.[21]

---

[20]In appendix B.5, we derive the long-run structural representations of the relationship between the shocks and our main outcomes under different structural assumptions.

[21]Formally, the inference is exposed to potential generated regressor bias leading to attenuation and underestimated standard errors. However, all the key results are quantitatively unchanged when estimating them internally in the VAR or relying on an instrumental variable strategy (see below),

19

Figure 1: *Employment Responses*

Impulse Responses To a Technology Shock   Impulse Responses To a Demand Shock

Note: The values are impulse responses to a one stanrard deviation shock expressed in percentage points. The x-axis denotes years since the shock. Lines depict the mean of the bootstrap distributions. Shaded areas depict the bootstrapped 95 percent confidence intervals calculated from 1,000 replications.

Moreover, we include time, ρ_t, and firm fixed effects, μ_j. The latter allows for drift terms in accordance with the stochastic process for the shocks postulated in section 2.1. The time fixed effects ensure that identification is driven by idiosyncratic, rather than aggregate, shocks.[22]

The long-run effects are estimated by adding additional lags of the structural shocks and summing the estimated coefficients. Using the variables included in the VAR as outcomes, the linear equation (13) augmented with the lags recovers the VMA parameters in equation (9). In practice, as we discussed in the context of the SVAR estimation and elaborate further below, the dynamic responses to permanent idiosyncratic shocks are fairly limited, and adding only one lag summarizes the long-run effect.[23]

Our baseline specification, following equation (13), is presented in table 4. Column 1 focuses on changes in net employment (defined as in Davis, Haltiwanger, and Schuh

---

both of which are insensitive to generated regressor biases. This is because the VAR is estimated with considerable precision using panel data, which minimizes the effects of the generated regressor biases.

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

[23]That is, the coefficients on additional lags are small and insignificant.

[1996]). As expected, the results are very similar to those presented in figure 1. The effect of a normal (one standard deviation) technology shock is 0.11 (not statistically different from 0). If we add one lag of the shocks to the regression and calculate the long-run employment responses (column 4), we find that a one standard deviation technology shock increases employment by 0.4 percentage points, the effect being statistically significant. Instead, demand shocks are the main drivers of employment adjustments: a positive one standard deviation shock to the demand curve increases employment by 5.6 (6) percentage points in the short (long) run. The small differences between the short- and long-run effects corroborate the limited dynamics found in the SVAR framework. Table 4, panel B, instead shows estimates in the form of elasticities (see appendix B on the computation of the elasticities). The conclusions are very similar.

Table 4: *Contemporaneous and Long-Run Effects on Labour Flows*

	SHORT RUN			LONG RUN		
	(1)	(2)	(3)	(4)	(5)	(6)
	NER	HR	SR	NER	HR	SR
A) Responses to a one standard deviation shock:						
Technology (η^a)	0.115	-0.050	-0.165*	0.412*	-0.093	-0.504**
	(0.119)	(0.075)	(0.078)	(0.163)	(0.116)	(0.128)
Demand (η^ω)	5.609**	2.906**	-2.703**	6.009**	3.125**	-2.884**
	(0.173)	(0.096)	(0.120)	(0.228)	(0.156)	(0.186)
B) Elasticities:						
Technology (η^a)	0.011	-0.005	-0.016*	0.0409*	-0.009	-0.050**
	(0.012)	(0.007)	(0.008)	(0.016)	(0.011)	(0.013)
Demand (η^ω)	0.347**	0.180**	-0.167**	0.371**	0.193**	-0.178**
	(0.011)	(0.006)	(0.007)	(0.014)	(0.010)	(0.012)
Observations	40,451	40,451	40,451	34,414	34,414	34,414
Firms	6,125	6,125	6,125	6,116	6,116	6,116

Note: Robust standard errors are in parentheses. NER: Net employment rate; HR: Hiring rate; SR: Separation rate. Hiring and separation rates are measured as the flow between the end points of two years divided by the average employment across these two points in time. The net employment rate is the difference between the hiring rate and the separation rate. The regressions include time dummies and firm fixed effects. The long-run impact is based on the sum of the contemporary effect and the effect of the first lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

We proceed by estimating equation (13) for hires and separations. The results (table 4, columns 2 and 3) show that a one standard deviation demand shock increases

21

the hiring rate by 2.9 percentage points and reduces the separation rate by 2.7 percentage points in the short run (slightly more in the long run, as shown in columns 5 and 6). These numbers should be compared with average hiring and separation rates of about 14 to 15 percent each, as shown in table 3. Thus, the estimates imply that 52 percent of the net employment adjustment is obtained using the hiring margin, and 48 percent using the separation margin. On average, firms rely as much on variations in separations as on variations in hires when responding to the shocks. The results further imply that the low response of net employment to technology shocks does not mask any substantive counteracting responses of gross worker flows. Rather, idiosyncratic technology shocks appear to have a limited impact on hiring and separation rates in the short and long run.

Next, we isolate the analysis of separations of short-tenure workers, which constitute 23 percent of the worker-year observations in the sample. The results in table 5, column 2, show that these make up slightly more than one-third of the total separation response to demand shocks in the short run (column 1) and even less in the longer run (column 4 versus column 5). The lower relative contribution of short-tenure separations in the long run is consistent with a reduction in contemporary hires, which reduces the number of short-tenure workers who can be released in the next period. As a final analysis, we document the responses in terms of "marginal workers," defined as short-term workers who are hired within the year but are not present in November, and thus were not classified as regular workers.[24] The results, presented in table 5, show that the adjustments in marginal workers are very similar to the adjustments in regular employees in the sense that most of the adjustment is due to demand shocks. We also see some evidence of overshooting, in the sense that the short-run response in the use of marginal workers (3.8 percent of the number of full-time employees) is larger than the long-run adjustment (3 percent).

Overall, our main results show that (i) the responses of employment and labour flows are much stronger to permanent demand shocks than to permanent technology shocks; (ii) most labour adjustments happen within the year; (iii) hires and separations are equally important as adjustment margins; and (iv) short-tenure separations and adjustments of marginal workers follow similar adjustment patterns as regular employees, but with a somewhat larger initial response.

---

[24]We measure the *number* of marginal employees and do not address the intensity with which these are used.

22

Table 5: *Contemporaneous and Long-Run Effects on Short-Tenure Separations and Marginal Workers*

	SHORT RUN			LONG RUN		
	(1)	(2)	(3)	(4)	(5)	(6)
	SR	STSR	MNER	SR	STSR	MNER
Technology (η^a)	-0.165*	-0.117**	0.110	-0.504**	-0.177**	0.482
	(0.078)	(0.038)	(0.162)	(0.128)	(0.066)	(0.248)
Demand (η^ω)	-2.703**	-1.010**	3.796**	-2.884**	-0.416**	3.019**
	(0.120)	(0.052)	(0.213)	(0.186)	(0.076)	(0.278)
Observations	40,451	40,451	40,451	34,414	34,414	34,414
Firms	6,125	6,125	6,125	6,116	6,116	6,116

Note: Values are the effect of a one standard deviation shock. SR: Separation rate; STSR.: Short-tenure separation rate, measured as the number of separations of short-tenure ($<$ 3 years) workers; MNER: marginal net employment rate, refers to workers who do not fulfill the criteria for full-time primary employment, but are employed by the firm at time t. All rates are measured as the flow between the end points of two years divided by the average (full-time primary) employment across these two points in time. The regressions include time dummies and firm fixed effects. The long-run impact is based on the sum of the contemporary effect and the effect of the first lag. Robust standard errors are in parentheses. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

4.2 Transitory Shocks

The focus of the analysis so far has been on how firms adjust employment, hires, and separations when hit by *permanent* idiosyncratic shocks. Here we instead derive an alternative measure of demand and technology shocks that also includes transitory disturbances. For technology shocks, we simply use the observed (physical) Solow residuals. For demand shocks, we use the residuals from estimation of a log-linearized version of the demand equation (2). This regression includes time dummies to control for aggregate shocks and firm fixed effects to eliminate between-firm permanent heterogeneity. Because prices are endogenous in the regression, we use the Solow residuals as instruments. Since the ensuing residuals of the estimated demand equation represent changes in sales without price adjustments, they serve as a measure of demand shocks. This strategy to derive demand and technology shocks is similar to that of Foster, Haltiwanger, and Syverson (2008). Thus, we label these shocks "FHS".

In contrast to our SVAR filter, the (static) FHS procedure does not differentiate between permanent and transitory shocks, and the processes do not account for factor

price shocks. The correlation between the FHS demand shocks and our baseline SVAR demand shocks is 0.538. The standard deviation of the FHS demand shocks is considerably higher than in the baseline SVAR (0.24 versus 0.16). Thus, the two demand shock series appear to contain a substantial common component without being identical. The correlation between the FHS demand shocks and the factor price component of the SVAR is considerably smaller (-0.25), but it is statistically significant. As expected, the FHS demand shocks are uncorrelated with the SVAR technology shocks. And, as expected given the limited dynamics observed in the physical Solow residual series, the physical Solow residual is highly correlated with the SVAR technology shocks (0.98), and only marginally related to our SVAR demand shocks (correlation of 0.02) and SVAR factor price shocks (correlation of 0.06).

Table 6: *Baseline Estimates versus Solow Residuals and FHS Demand Shocks for Log Employment*

	SHORT RUN			LONG RUN		
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	FHS	Transitory	Baseline	FHS	Transitory
Technology (η^a)	0.153	0.333*	-0.157	0.504*	0.993**	0.0754
	(0.159)	(0.168)	(0.164)	(0.214)	(0.250)	(0.216)
Demand (η^ω)	5.986**	3.406**	0.674**	6.357**	4.061**	0.863**
	(0.233)	(0.183)	(0.136)	(0.310)	(0.252)	(0.217)
Observations	40,451	40,451	40,451	34,414	34,414	34,414
Firms	6,125	6,125	6,125	6,116	6,116	6,116

Note: Values are the effect of a one standard deviation shock. In the FHS columns the technology shock is the Solow residual, and the demand shock is FHS demand, as defined in the main text. The transitory shocks are calculated as the residual component of the FHS series. Robust standard errors are in parentheses. The regressions include time dummies and firm fixed effects. The long-run impact is based on the sum of the contemporary effect and the effect of the first lag. ** and * denote statistical significance at the 1 and 5 percent levels, respectively.

Table 6 shows how these measures relate to labour flows. The estimates with our SVAR shocks are reproduced in columns 1 and 4 for comparison. Clearly, the main findings hold when using the FHS series (columns 2 and 5): the short-run impact of the demand shocks is 10 times that of the technology shock in the short run, and about four times in the long run. But it is also noticeable that the estimated impact of the demand shocks is about half as large when using FHS demand as when using the SVAR demand shock.

24

To see what drives the difference, we proceed by purging the FHS series of our permanent structural shocks. To this end, we run a regression with the FHS demand as the dependent variable and use our SVAR shocks (demand, technology, and factor prices) as regressors and then repeat this for the Solow residual. We label the residuals of this exercise *transitory* demand and technology shocks. Because these residuals are measured in the same units as the composite FHS demand and technology shocks, we can directly compare their impacts on employment adjustments.[25]

The results, in table 6, columns 3 and 6, show that the ensuing transitory demand shocks have a much more muted impact on employment than the SVAR and composite FHS shocks. This reinforces the idea that our SVAR strategy captures the most relevant determinants of labour adjustments. The result holds for the short- and long-run responses. That the long-run response to transitory demand shocks does not revert back when the lag is introduced suggests that the transitory series may still contain a persistent component.[26] With this caveat in mind, that the part of the demand series that is certified to be permanent has a much larger effect suggests that firms' employment adjustment depends on the time-series properties of the shocks, as in Franco and Philippon (2007); Roys (2016); and Guiso, Schivardi, and Pistaferri (2005).

4.3 Robustness

We have carried out a an extensive battery of checks to assess the robustness of our claims that (i) firm-level demand shocks are more important in the determination of labour adjustments than firm-level technology shocks, and (ii) employment adjustment to the permanent shocks is very rapid with limited short-term dynamics. To conserve space, we defer discussions about basic specification checks such as variations in the demand elasticity, sample selection, using an alternative fourth variable, and variations based on alternative VAR specifications to our online appendix E. Here, we focus on those exercises that contain more economic intuition.

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

---

[25]The decomposition resembles that in Guiso, Schivardi, and Pistaferri (2005), which extracts the permanent component of firm-level value added using high-order polynomials of lags as instruments. Although the mechanics of the methods differ, the underlying logic is similar.

[26]If we partial out the series with the residual shock of the SVAR $\left(\eta_{jt}^y\right)$, the estimates are about half the size and the long-run estimate is nonsignificant.

25

and Nordström-Skans (2016), we estimate RTS separately for the durables and non-durables sectors among Swedish manufacturing firms, obtaining 1 for durables and 0.9 for non-durables. In both cases, we cannot reject the null of constant RTS. These results are very similar to what Basu, Fernald, and Kimball (2006) report for the United States. What matters is the long-run RTS, which implies that the theoretical case for assuming constant RTS becomes stronger. To assess the robustness of the results to this assumption, the model can be altered to accommodate increasing or decreasing RTS. This affects the measures that are fed into the SVAR (for details, see appendix B.4) and hence also the estimated magnitudes of employment adjustments. However, the main conclusions from the baseline analysis are not altered. Appendix table E1, column 2, shows that a positive technology shock of one standard deviation raises employment by 1 percentage point in the short run (1.4 in the long run, see column 5) when the Solow residual is constructed using 0.9 RTS. But this estimate remains far below the estimated impact of a demand shock: an increase of 6.1 percentage points in the short run and 6.3 percentage points in the long run. If instead we impose an RTS coefficient of 1.1, the results change in the other direction (the impact of technology turns negative), but the main message for the strong relative importance of demand remains unaltered.[27]

Sectoral heterogeneity. The dynamic panel approach used for estimation took advantage of our large-N, small-T panel setting to estimate the VAR system with considerable precision. This is a key advantage relative to standard SVAR estimations in the macro literature. However, a potential cost is that the underlying dynamic processes are assumed to be equal across different types of firms. To address this concern, we have allowed for separate dynamics for each two-digit industry, and the employment adjustment results remain unchanged (see appendix table E2, column 6). Demand shocks have a much larger impact on employment adjustment than technology shocks do.

Firm exit. A possible concern with the analysis is that we disregard the firm exit

---

[27]An alternative robustness exercise would be to estimate production functions with one of the many proposed estimators in the literature (for example, following Olley and Pakes (1996)). However, this would only matter in practice if the output elasticities were very different from the cost shares, and this would certainly be a cause for concern since it would require that the overall returns to scale fell outside reasonable bounds. This is why we choose to use the cost shares (measured as revenue shares) directly and instead provide robustness exercises over the returns to scale. Given that the overall results are robust to these variations, relying on any sensible production function estimates would not change our results.

process. Firms are likely to exit in response to severe negative demand or technology shocks, and this process may impact labour dynamics. Furthermore, we lose coverage in the IS firm-level data when firms shrink below 10 employees. To address these concerns, we analyse the employment impact of the shocks using a two-period specification instead of the one-period baseline. In the baseline, we evaluate employment changes between t and $t-1$ divided by average employment over the two years. The new specifications change the numerator, which is now defined over changes between $t+1$ and $t-1$. Since labour flows are defined even if all workers exit the year after the shock, and because we can measure employment also when firms fall below 10 employees in the RAMS data, we can calculate the impact of the shocks excluding or including the firms that fall below the threshold (or exit entirely) in $t+1$. Reassuringly, the results are insensitive to whether we include or exclude these observations (see appendix table E6).[28]

Endogenous demand elasticities. The difference in employment responses between permanent demand and technology shocks should be understood in a context where the two types of shocks have similar empirical relationships to output, and where technology affects prices much more forcefully than demand shocks do. Qualitatively, these patterns are thus all well in line with what is expected from our motivating model presented in section 2. But in the model, the magnitudes of the employment, output, and price responses to the two shocks are all tightly determined by a (constant) value of σ (see appendix B.5 for the full Jacobian). Unsurprisingly, the magnitudes of the empirical (long-run) responses of employment, prices, and output to those shocks do not concur with a single constant value of σ. For example, employment responses to technology and demand shocks in the model are related by a factor of $\frac{1}{\sigma-1}$, which, given the long-run responses of employment (shown in figure 1), would imply a value of σ of about 1.1. Instead, the output and price responses to technology shocks (shown in appendix figure B5) suggest a value of $\sigma = 3.3$ (see the discussion in appendix B.5). Although we show in the appendix that we can choose any reasonable number for σ without affecting the results, we cannot satisfy the full set of responses in output, prices, and employment with any single value of σ. The data thus seem to ask for a model that is richer in its description of product market responses to the shocks,

---

[28] We have also analysed the explicit relationship between the shocks and the probability of firm exit from the sample. The main driver of firm exits is large negative demand shocks, which is well in line with the results for the United States in Foster, Haltiwanger, and Syverson (2008).

27

although the assumption of a constant σ is one we share with most of the literature.[29] Fortunately, this apparent anomaly can be resolved by a straightforward generalization of the model, allowing the elasticity of demand (and thereby the markup) to be a function of the shocks, that is, allowing for $\sigma = \sigma(A_{jt}, \Omega_{jt})$, in the spirit of Kimball (1995). This replaces equation (2) by

$$Y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\sigma(A_{jt},\Omega_{jt})} Y_t \Omega_{jt}, \ \ \sigma(A_{jt}, \Omega_{jt}) > 1 \text{ and } \sigma(\bar{A}_{jt}, \bar{\Omega}_{jt}) = \sigma, \qquad (14)$$

where an upper-bar denotes an average across firms. The only change relative to the measurement equations outlined in table 1 of section 2 is that WND acknowledges that σ is no longer constant. The modified model-expression is thus

$$WND = \psi(A_{jt}, \Omega_{jt}) Y_t P_t^{\sigma(A_{jt},\Omega_{jt})} A_{jt}^{\sigma(A_{jt},\Omega_{jt})} \left(P_{jt}^F\right)^{-\sigma(A_{jt},\Omega_{jt})} \Omega_{jt}. \qquad (15)$$

This extension is fully consistent with our long-run restrictions. The extension is discussed in detail and quantified in appendix B.5. The quantification shows that the derivatives of σ to reconcile the joint responses of output, prices, and employment to the two shocks are quite modest. The implied values of σ in response to an interval of ± 1 standard deviation technology shocks is $[2.7, 3.9]$ and for demand shocks the corresponding interval is even tighter, at $[3.2, 3.4]$.

5 Asymmetries and Worker Flows

This section provides an analysis of how firm-level employment adjustments of different signs and magnitudes in response to permanent demand shocks translate into worker flows.[30] This analysis is similar in spirit to that of Abowd, Corbel, and Kramarz (1999) and Davis, Faberman, and Haltiwanger (2012), who provide decomposition exercises on the relative contributions of various worker flows to positive and negative changes in net employment within French and U.S. firms, respectively. In contrast to these decompositions, we analyse adjustments due to a well-defined permanent shock. This provides a causal layer to the analysis, which is potentially important

---

[29]Including, for example, Foster, Haltiwanger, and Syverson (2008, 2016) and Pozzi and Schivardi (2016).

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

28

since short-run employment fluctuations may be caused by worker flows (i.e., exits or unsuccessful hiring attempts may lead to fluctuations in employment). Using our structural demand shocks as instruments for employment adjustments ensures that the analysis is immune to such reverse causality.

First, however, we characterize the potential asymmetries in net employment adjustments in response to the shocks by replacing the linear terms η_{jt}^a and η_{jt}^ω in equation (13) with two second-order polynomials, one for positive values and one for negative values.[31] Figure 2 documents how the shocks affect the change in employment ($\Delta Employment_{jt}$) using this functional form.[32] For completeness, we show the responses to technology and demand shocks but focus our attention on the demand shock responses. The demand shocks have a similar, almost linear, relationship to employment adjustments on both sides of zero. The net change in employment in response to a one standard deviation positive demand shock (6 percentage points) is reasonably close to the response to a one standard deviation negative shock (-7 percentage points) in absolute values. The differences at the endpoints of very large (± 2 standard deviations) shocks are somewhat more pronounced (9 versus -13 percentage points).[33]

Next, we analyse the potentially asymmetric use of hiring and separations as adjustment margins in response to positive and negative changes in net employment due to demand shocks. To this end, we use an instrumental variables (IV) approach where demand shocks serve as instruments for employment adjustments. We thus replace η_{jt}^a and η_{jt}^ω in equation (13) with a two-sided second-order polynomial of employment changes and instrument these by a corresponding polynomial for the demand shocks (see appendix C). The IV strategy allows us to see the relative contribution of hiring (versus separation) to employment adjustments as a function of the sign and magnitude of the induced employment adjustment. Due to the underlying identity ($\Delta Employment_{jt} = Hiring_{jt} - Separations_{jt}$), we only show estimates on the impact on *hiring rate* (the separation response is the exact mirror image) and explicitly calculate the share of adjustments arising from hiring as a function of the sign and

---

[31]Using shorthand for the indicator function $I_x^+ = I(x_{jt} > 0)$, we include in the regressions $G(x_{jt})$ $= g_1 I_x^+ x_{jt} + g_2 I_x^+ x_{jt}^2 + g_3 (1 - I_x^+) x_{jt} + g_4 (1 - I_x^+) x_{jt}^2$ for each of the shocks $x_{jt} = (\eta_{jt}^a, \eta_{jt}^\omega)$.

[32]A nonparametric description of the data suggests that the functional form is reasonable. Figures are available from the authors on request.

[33]Appendix figure E1 shows estimates of the responses to the *transitory* demand shocks derived in section 4.2. As expected, their impacts are substantially lower than the impact of the permanent shocks, regardless of the sign or magnitude of the shock.

29

Figure 2: *Shocks and the Net Employment Rate*

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

magnitude of the shock. For the direct impact of the shocks on hiring and separations, see Appendix E.

The results presented in figure 3 show evidence of a strongly asymmetric use of the two possible adjustment margins. There is a clear positive and essentially linear relationship between net employment adjustments and hires when employment is growing, but a very modest relationship when employment is shrinking. Trivially, this implies that the separation-response must have the opposite structure. This becomes clear from the right-hand panel in figure 3, which shows the share of employment adjustment that takes place through hires (versus separations), as a function of demand-induced net employment changes. The share of adjustments through hiring jumps up from 20 to 95 percent when employment adjustments become positive instead of negative. This implies that the share of adjustments through separations jumps down from 80 to 5 percent instead.[34] To complete the picture, appendix E shows the reduced-form impact of the shocks directly on hires and separations, confirming the impression that firms adjust to positive shocks through increased hires and to negative shocks through

---

[34]In contrast to figure 2 (where zero referred to the absence of an idiosyncratic shock), zero here refers to the state when net employment adjustment is predicted to be zero based on the first stage (i.e., based on the combination of the shock polynomials, year dummies, and firm fixed effects).

Figure 3: *Hiring Rate and Net Employment Changes: IV Results*

Note: Left-side panel: contemporaneous hiring rate in percentage units as a (nonlinear) function of employment adjustment in percentage units. Employment adjustments are instrumented by demand shocks. The shaded area depicts 95 percent confidence intervals. Right-side panel: implied fraction of employment adjustment achieved through changes in hirings as a function of the size and magnitude of the employment adjustment.

increased separations.

5.1 Actual and Possible Hiring Responses

Obviously, there are limits to how much a firm *can* reduce its employment by not hiring. We therefore let figure 4 show a set of accounting exercises where we contrast the actual responses from figure 3 (but focusing on only negative adjustments) with assessments of how much firms *could* have reduced employment without inducing separations. The first assessment scenario, denoted "hypothetical homogeneous," imposes the *average empirical separation rate of firms* with unchanged employment within our data (10 percent) on all firms. In this case, as long as the need for adjustment is 10 percent or less, reduced hires could fully accommodate the necessary employment reduction. If the shock is 20 (30) percent instead, the firm could instead accommodate half (one-third) of the adjustment through reduced hires. The curve shows that firms could have accommodated much larger shocks without induced separations if the separation rate remained at 10 percent at all times.

But assuming a fixed exit rate of 10 percent is clearly not a valid assumption for

Figure 4: *Actual (IV) and Simulated Hiring Responses to Employment Changes*

Note: The figure shows the actual (estimated from data) and hypothetical maximum (simulated) fraction of negative employment adjustments (in percentage units) achieved through changes in hirings. Employment adjustments are instrumented by demand shocks. "Hypothetical homogeneous" assumes that the same fraction of workers always leaves the firm. "Hypothetical heterogeneous" imposes a random individual quit rate on the actual firm size distribution.

small firms, as their number of exits will differ between years for stochastic reasons. We therefore also provide a second benchmark, assuming instead that the *individual* probability of leaving a firm is 10 percent. We randomly allocate quits across all workers in our full sample and then aggregate the data to the firm level to get a firm-level counterfactual distribution of quit rates. Within this distribution, some (primarily small) firms will not experience any quits at all, which means that they cannot accommodate even the smallest employment reduction through reduced hires. Other firms will experience many random separations, allowing them to accommodate large employment reductions through reduced hires. The curve denoted "hypothetical heterogeneous" displays the simulated frontier of adjustments with random individual quits within the data. As is evident, the observed (actual) employment adjustments are far from this benchmark. The actual share of adjustment through reduced hires is much lower than a hypothetical strict reliance on hires would allow for. The shaded area between the heterogeneous hypothetical curve and the actual behaviour of the firm could be interpreted as a region of flexibility, because it depicts the amount of negative labour adjustments through induced separations (i.e., separations above the

32

random rate) that could have been accomplished through reduced hires instead.

The main take-away from this exercise is that there is substantial scope for firms to rely on a symmetric use of hiring-adjustments (without induced separations) across fairly large positive and negative shocks. But instead, firms choose to let the adjustment margins be an asymmetric function of the sign of the shock, where they primarily vary their separation rates when shocks are negative and instead (almost) only adjust their hiring rates in response to positive shocks.[35] This finding implies that firing costs (broadly defined to include any impediments to separations, for example regulatory restrictions, buyouts, or loss of morale from stayers) do not appear to prevent firms from continuing to hire when they are reducing net employment (see Abowd, Corbel, and Kramarz (1999) for a discussion of the role of firing costs in this context).

6   Conclusions

This paper has analysed how firms adjust their labour inputs in response to permanent idiosyncratic firm-level shocks to technology and demand. We identify the shocks by imposing a set of long-run restrictions in an SVAR estimated on firm-level data. The restrictions are derived from a stylized model of monopolistically competitive firms. The SVAR is estimated using dynamic panel data methods, allowing us to identify the parameters of the reduced form with considerable precision. To estimate the model, we rely on a very rich data set that merges information about inputs, outputs, and prices of Swedish manufacturing firms with a linked employer-employee data set.

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

Our labour adjustment results show that both the nature (as argued by Foster, Haltiwanger, and Syverson (2008)) and time-series properties (as argued by Guiso, Schivardi, and Pistaferri (2005)) of the shocks matter. Permanent demand shocks, which affect output but not relative prices, have a pronounced impact on employment. In contrast, technology shocks have relatively limited employment effects despite affecting output and relative prices. A similar limited employment response is found

---

[35]In appendix D we show evidence from heterogeneity analyses related to this process.

33

for transitory demand shocks, which may explain why we find a larger employment response from demand-side disturbances compared with Pozzi and Schivardi (2016), who do not distinguish between permanent and transitory shocks.[36]

Further results suggest that employment adjustments in response to permanent shifts in the product demand curve are fast and symmetric. By far the largest part of employment adjustment takes place within a year. Almost as much of the employment adjustment is through changes in separation rates as through changes in hiring rates, suggesting that both margins should be considered endogenous at the firm level.

Finally, we provide the first analysis of the asymmetric impact on worker flows (hires versus separations) when employment needs to adjust because of the relevant (i.e., permanent, demand) shocks. The results show that the sign of the shock determines the primary margin of adjustment: firms primarily adjust through separations if shocks are negative and primarily through hires if shocks are positive. The fact that negative shocks are accommodated by increased separations rather than reduced hires implies that the employment adjustments cause excessive worker flows.

The speed of adjustment, symmetry between hires and separations as adjustment margins, and continued recruitment of workers in the face of negative shocks jointly suggest that labour market rigidities play a very limited role in hampering firm-level labour adjustments in the face of permanent idiosyncratic demand shocks. However, the adjustments with respect to transitory shocks are much more muted. Thus, firms accommodate the impact of permanent shocks, but may hoard labour and refrain from hiring when hit by transitory shocks.

The conclusion that the nature of the shocks (technology versus demand) and the time-series properties (permanent versus transitory) of these shocks matter for job and worker reallocation suggests that cross-country comparisons of labour flows need to be careful in accounting for the types of shocks that hit these economies. Building on this notion, our empirical approach also suggests a route forward in trying to understand the forces behind the declining rates of labour adjustments observed, for example, in the United States. Essentially, our empirical approach provides a tool for assessing whether this development is due to the changing nature of the underlying firm-level shocks or the reduced impact of these shocks on labour reallocation. Although this is beyond the scope of this paper, it serves as an example of the type of questions that

---

[36]Our finding that the impact of technology shocks on employment is small is more similar to Pozzi and Schivardi (2016), which is reasonable since almost all within-firm disturbances in TFPQ are permanent, which makes the distinction between permanent and transitory shocks superfluous.

34

can be answered by combining data on labour flows with well-identified, firm-level structural shocks.

Uppsala University, UCLS, and Sveriges Riksbank

Inter-American Development Bank, and IZA

Uppsala University, IFAU, UCLS, and IZA

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(2), 170–187.

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

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

BASU, S., J. FERNALD, AND M. KIMBALL (2006): "Are Technology Improvements Contractionary?," *American Economic Review*, 95, 1418–1448.

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

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

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

CABALLERO, R., E. ENGLE, AND J. HALTIWANGER (1997): "Aggregate Employment Dynamics: Building from Microeconomic Evidence," *American Economic Review*, 87, 115–137.

CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, 74(2), 323–364.

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

CARLSSON, M., J. MESSINA, AND O. NORDSTRÖM-SKANS (2016): "Wage Adjustment and Productivity Shocks," *Economic Journal*, 126, 1739–1773.

36

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.

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

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

DAVIS, S. J., J. C. HALTIWANGER, AND S. SCHUH (1996): *Job Creation and Job Destruction*. MIT Press, Cambridge, MA.

ERCEG, C., D. HENDERSON, AND A. LEVIN (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46, 281–313.

ESLAVA, M., AND J. HALTIWANGER (2018): "The Life-Cycle Growth of Plants: The Role of Productivity, Demand and Distortions," Available at SSRN: https://ssrn.com/abstract=3177289 or http://dx.doi.org/10.2139/ssrn.3177289.

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

ESLAVA, M., J. HALTIWANGER, A. KUGLER, AND M. KUGLER (2010): "Factor Adjustments after Deregulation: Panel Evidence from Colombian Plants," *Review of Economics and Statistics*, 92, 378–391.

FORLANI, E., R. MARTIN, G. MION, AND M. MUÛLS (2016): "Unraveling Firms: Demand, Productivity and Markups Heterogeneity," CEP Discussion Papers 1402, London School of Economics.

37

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 (2016): "The Slow Growth of New Plants: Learning and Demand," *Economica*, 83(329), 91–129.

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

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 (2013): "Competition, Takeovers and Gender Discrimination," *Industrial and Labor Relations Review*, 66(2), 409–432.

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

HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131, 1291–1364.

KIMBALL, M. (1995): "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, 27, 1241–1277.

LISE, J., C. MEGHIR, AND J.-M. ROBIN (2016): "Matching, Sorting and Wages," *Review of Economic Dynamics*, 19, 63–87.

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

MICHELACCI, C., AND D. LOPEZ-SALIDO (2007): "Technology Shocks and Job Flows," *Review of Economic Studies*, 74, 1195–1227.

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

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

OLLEY, G. S., AND A. PAKES (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64(6), 1263–1297.

POZZI, A., AND F. SCHIVARDI (2016): "Demand or Productivity: What Determines Firm Growth?," *Rand Journal of Economics*, 47, 608–630.

ROYS, N. (2016): "Persistence of Shocks and the Reallocation of Labor," *Review of Economic Dynamics*, 22, 109–130.

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

SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): "Firming Up Inequality," *Quarterly Journal of Economics*, 134, 1–50.

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

39