SVERIGES RIKSBANK
WORKING PAPER SERIES 198

Technology Shocks and the Labour-Input Response: Evidence from Firm-Level Data

Mikael Carlsson and Jon Smedsaas

MAY 2006

WORKING PAPERS ARE OBTAINABLE FROM

Sveriges Riksbank • Information Riksbank • SE-103 37 Stockholm
Fax international: +46 8 787 05 26
Telephone international: +46 8 787 01 00
E-mail: info@riksbank.se

Technology Shocks and the Labor-Input Response: Evidence from Firm-Level Data*

Mikael Carlsson† and Jon Smedsaas‡

Sveriges Riksbank Working Paper Series
No. 198
May 2006

Abstract

We study the relationship between technology shocks and labor input on Swedish firm-level data using a production function approach to identify technology shocks. Taking standard steps yields a contractionary contemporaneous labor-input response in line with previous studies. This finding may, however, be driven by measurement errors in the labor-input variable. Relying on a unique feature of our data set, which contains two independently measured firm-specific labor input measures, we can evaluate the potential bias. We do not find any evidence supporting that this bias would conceal any true positive contemporaneous effect. The results thus point away from standard flexible-price models and towards models emphasizing firm-level rigidities.

Keywords: Technology Shocks; Labor Input; Business Fluctuations; Micro Data
JEL classifications: C33; D24; E32

*We would like to thank Malin Adolfson, Susanto Basu, Oskar Nordström-Skans and seminar participants at Sveriges Riksbank for useful comments and Björn Andersson and Magnus Lundin for help with preparing the data. Part of this project was conducted while Carlsson was visiting the Robert Schuman Center for Advanced Studies at the European University Institute, whose hospitality is much appreciated. Carlsson gratefully acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank.

†Corresponding Author: Research Department, Sveriges Riksbank, SE 103 37 Stockholm, Sweden. E-mail: mikael.carlsson@riksbank.se.

‡Department of Economics, Uppsala University.

1 Introduction

Recently a series of papers have argued that technology improvements have a contractionary short-run effect on labor input (e.g., Basu, Fernald, and Kimball (2004), Carlsson (2003), Galí (1999), Francis and Ramey (2003) and Alexius and Carlsson (2005)). This finding is both important and controversial, since it provides evidence on the relative empirical importance of different types of business cycle models. In the real business cycle (RBC) approach to modeling business cycles, the central source of macroeconomic fluctuations is technology shocks; see King and Rebelo (1999) for a recent exposition. The main implication of the above finding is that there is no empirical support for the key mechanism in the canonical RBC model where technology, output and hours all rise together to generate realistic business cycle patterns. Instead e.g. Basu, Fernald, and Kimball (2004) and Galí (1999) argue that the negative correlation is due to nominal frictions, which are central to New Keynesian models.[1]

All papers cited above rely on data that is aggregated at various levels. In this paper, we take a micro-data approach to address the short-run relationship between technology and labor input. There are three main motivations for a study on firm-level data. First, any results regarding the relationship between technology and labor input stemming from aggregate data are necessarily conditioned on the specific monetary policy regime during the sample period (see e.g. Dotsey (2002) for a discussion). However, when using micro data, concerns about the exact behavior of monetary policy, or shifts in the same during the sample period, are no longer an issue. This is due to the fact that the central bank can only react to the common component in micro-level technology shocks. Thus, if this component is important for the overall variation in firm-level technology shocks, it is possible to instead restrict the attention to the idiosyncratic variation in technology.

[1] This finding has been questioned by other studies. For example, Christiano and Vigfusson (2003) argue that the result in Galí (1999) crucially depends on how the labor-input series (used in the structural VAR approach) is filtered. See also the response to this criticism in Galí (2004), Fernald (2005) and Francis and Ramey (2003).

The second motivation, which is related to the first, is that if the common component is of no importance for the overall variation in firm-level technology shocks, or if we focus on the idiosyncratic part only, this shock should not give rise to any general equilibrium effects. Thus, this type of study provides direct evidence on the extent of adjustment frictions specific to the firms' environment. Note that although we need to think differently about the effects of an aggregate and a firm-specific technology shock, if Basu, Fernald, and Kimball (2004) and Galí's (1999) notion of price stickiness as the key mechanism driving the observed negative aggregate relationship is correct, we should also see a negative contemporaneous relationship between firm-specific technology shocks and the firms' employment decisions. In fact, since there are no general equilibrium effects (including policy interventions) at work in this case, the mechanism giving rise to a negative correlation under price stickiness is much less involved. That is, if firms cannot immediately adjust their relative prices, their demand will not change immediately either, and since demand can be met using less labor (due to the improvement in technology), we expect a negative contemporaneous effect on labor input from an idiosyncratic technology shock.

The third advantage of using micro-level data is that cleansing effects of recessions will not be an issue either, in contrast to aggregate studies. That is, if inefficient firms are driven out of business in recessions, aggregate technology will improve in times of low economic activity. Thus, the aggregate finding of a contractionary effect of technology on labor input may actually be driven by reversed causation. However, this aggregation effect will not be an issue here since we will focus on a balanced panel of firms.

A central issue in this type of study is identifying technology. Here, we employ a robust production-function approach, as used in e.g. Basu, Fernald, and Kimball (2004) and Carlsson (2003), to estimate technology as the residual from a production function. To our knowledge, the only previous micro-data study is Marchetti and Nucci (2005) which relies on a very similar approach when studying the relationship between technology and labor input on Italian firm-level data. Following their steps, we also find that labor input falls contemporaneously in response to a positive technology shock.

Relative to Marchetti and Nucci (2005) this paper provides two important contributions. First, a major concern related to this type of study regards measurement errors in the labor input measure, because any such measurement errors will end up in the technology measure with a negative sign. Thus, any observed negative correlation between technology and labor input may actually be driven by measurement errors. In this paper, we investigate whether this type of measurement errors is important for the results. One approach to study this issue is to use an alternative independently measured labor-input measure, instead of the original labor-input measure in the regression of technology on labor input. Since measurement errors in the two labor-input series are expected to be uncorrelated, the latter results would not suffer from the negative bias. A unique feature of the Swedish data set used is that it actually contains two independently measured firm-specific measures of labor input; first, a survey-based measure from the Industry Statistics Survey (IS) and second, a measure based on government register data drawn from the Register Based Labor Market Statistics (RAMS) data base. This allows us to explore the effects of measurement errors in the labor-input variable.

Our second contribution is that we have access to a firm-level producer price index. This is important considering the concern raised by Klette and Griliches (1996) regarding biased returns to scale estimates stemming from using sectoral price deflators when computing firm-level real gross output.[2] This, in turn, is important since the returns to scale constitute a key parameter when computing the technology measure.

The main findings in this paper are (i) that the common component of technology shocks is very small, which leads us to only focus on the effects of a firm-level technology shock in this paper. (ii) We find a contemporaneous contractionary effect on labor input from a positive technology shock when using

[2]It should be noted that Marchetti and Nucci (2005) take the Klette and Griliches (1996) approach and include a control variable, motivated by an assumption about the shape of the demand-function faced by the firm, to address this problem.

the same labor-input measure as when constructing the technology measure. (iii) When evaluating this finding using our alternative labor-input data, we find no evidence supporting that the bias stemming from measurement errors conceals any true significantly positive contemporaneous effect, although the true contemporaneous effect may very well be zero. (iv) We find lagged positive effects on labor input when using both labor input measures. Overall, our results are in line with the view that there is some impediment to immediate adjustment at the firm level. The non-expansionary contemporaneous effect of technology shocks on labor input followed by lagged positive effects points away from the standard propagation mechanism in flexible-price models and towards models emphasizing firm-level rigidities.

This paper is organized as follows: Section (2) outlines the empirical specification to estimate technology shocks. Section (3) discusses the data set used and the estimation procedure together with intermediary results. Section (4) presents the results for the labor-input response, including our findings when using the alternative labor-input measure, and finally Section (5) concludes.

2 Empirical Specification

The idea behind the production function approach is that technology can be measured as the residual from a production function, taking increases in production factors and the intensity to which they are used into account. We start by postulating the following production function for firm i:

$$Y_{it} = F(Z_{it}K_{it}, E_{it}N_{it}, V_{it}, M_{it}, A_{it}), \tag{1}$$

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

Taking the total differential of the log of (1) and invoking cost minimization, we arrive at:

$$\Delta y_{it} = \eta[\Delta x_{it} + \Delta u_{it}] + \Delta a_{it}, \tag{2}$$

where Δy_{it} is the growth rate of gross output, η is the overall returns to scale, Δx_{it} is a cost share weighted input index defined as $C_{iK}\Delta k_{it} + C_{iN}\Delta n_{it} + C_{iV}\Delta v_{it} + C_{iM}\Delta m_{it}$, $\Delta u_{it} = C_{iK}\Delta z_{it} + C_{iL}\Delta e_{it}$ and C_{iJ} is the cost share of factor J in total costs.[3] Thus, given data on factor compensation, changes in output, input and utilization, and an estimate of the returns to scale η, the resulting residual Δa_{it} provides a times series of technology change for the firm. Notice that Δa_{it} reduces to a gross-output Solow residual if $\eta = 1$, $\Delta u_{it} = 0$ and there are no economic profits.[4] Hence, Δa_{it} is a Solow residual purged of the effects of non-constant returns, imperfect competition and varying factor utilization.

The main empirical problem associated with (2) is that capital and labor utilization are not observed. A solution to this problem is then to include proxies for factor utilization in (2). Here, we follow the approach taken by e.g. Burnside, Eichenbaum, and Rebelo (1995) who use energy consumption as a proxy for the flow of capital services. This procedure can be legitimized by assuming that there is a zero elasticity of substitution between energy and the flow of capital services. This, in turn, implies energy and capital services to be perfectly correlated. Adding the assumption that labor utilization is constant, we arrive at the following specification by Burnside, Eichenbaum, and Rebelo (1995):

$$\Delta y_{it} = \eta\Delta\widetilde{x}_{it} + \Delta a_{it}, \tag{3}$$

where input growth, $\Delta\widetilde{x}_{it}$, is defined as $(C_{iK} + C_{iV})\Delta v_{it} + C_{iN}\Delta n_{it} + C_{iM}\Delta m_{it}$. In a closely related paper, Carlsson (2003) experiments with using various proxies for labor utilization when estimating production functions like equation (3)

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

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

on Swedish manufacturing industry data.[5] However, including these controls has no discernible impact on the results.

Note also that we do not need any direct measure of the firm's capital stock. This is a very useful feature of this specification, since the latter variable is very difficult to measure with only a short span of investment data in the time dimension.

When empirically implementing the specification (3), we take an approach akin to the strategy outlined by Basu, Fernald, and Shapiro (2001). First, the specifications are regarded as log-linear approximations around the steady state growth path. Thus, the products ηC_{iJ} (i.e. the output elasticities) are treated as constants. Second, the steady-state cost shares are estimated as the time average of the cost shares for the three-digit industry to which the firm belongs. Third, to calculate the cost shares, we assume that firms make zero profit in the steady state.[6] Taking total costs as approximately equal to total revenues, we can infer the cost shares from factor shares in total revenues. The cost share of capital and energy is then given by one minus the sum of the cost shares for all other factors. Finally, the growth rate of technology, i.e. technology change, is modeled as $\Delta a_{it} = \alpha_i + v_t + \varepsilon_{it}$, where α_i is the firm-specific deterministic growth rate in technology (drift), v_t is a common technology disturbance across all firms and ε_{it} is an idiosyncratic technology disturbance. Inserting this into (3) yields:

$$\Delta y_{it} = \eta \Delta \widetilde{x}_{it} + \alpha_i + v_t + \varepsilon_{it}. \tag{4}$$

3 Data and Estimation

The data set we use is primarily drawn from the Industry Statistics Survey (IS) and contains annual information for the years 1989-1996 on inputs and output

[5] That is, hours per employee, overtime per employee and the frequency of industrial accidents per hour worked.

[6] This seems to be a reasonable assumption judging from direct evidence from aggregate Swedish industry data (using the data underlying Carlsson, 2003).

for all Swedish manufacturing plants with 20 employees or more and a sample of smaller plants (see Appendix A for details on the data set we use). Since there is one plant per firm for about 80 percent of the observations we can, in practice, consider the data as representing firms.

An important feature of the data set is that it includes a firm-specific producer price index constructed by Statistics Sweden. Thus, the concern raised by Klette and Griliches (1996) regarding biased returns to scale estimates stemming from using sectoral price deflators when computing real gross output should not be an issue here. This is important, since the returns to scale constitute a key parameter when computing the technology measure.

Another important quality of our data set is that, besides the survey based firm-specific (annual average) employment measure available in IS (on which we rely when estimating technology growth), we also have access to a firm-specific measure of the number of employees in November from the Register Based Labor Market Statistics data base (RAMS). Whereas the IS employment data is based on a survey collected by Statistics Sweden, the RAMS employment data is based on the income statements that employers are, by law, required to send to the Swedish Tax Authority. The key aspect here is that the IS and the RAMS measure of labor input are measured independently from each other. Therefore, it is very unlikely that any measurement errors are the same in these two measures of labor input. Thus, we can use the RAMS labor-input measure to explore the effects on the observed short-run relationship between technology and labor input from measurement errors.[7]

Given the availability of data and after standard cleaning procedures, we are left with a balanced panel of $1,516$ firms observed over the years 1990-1996 (once more, see Appendix A for details).

Since the firm is likely to consider the current state of technology when making its input choices, we need to resort to an instrumental variable technique when estimating (4). Following Marchetti and Nucci (2005), we use a difference-

[7] The correlation between Δn and Δn^{RAMS} is 0.41 (allowing for fixed effects).

GMM estimator developed by Arellano and Bond (1991). Under the null of no serial correlation in the technology residual, we can use suitably lagged $\Delta\widetilde{x}$ as instruments. Guided by the Hansen test of the over-identifying restrictions, we use $\Delta\widetilde{x}_{it-s}$, $s \geq 3$ as instruments.[8] Moreover, to avoid including irrelevant instruments, we truncate the instrument set at the fifth lag.

Table 1: Returns to Scale Regression

η	1.174 (0.182)
Time Dummies	Yes
AR(2)	1.44 [0.151]
AR(3)	0.48 [0.628]
Hansen	4.08 [0.850]

Sample 1990-1996 with 1,516 plants (i.e. 9,096 observations). Difference GMM second-step estimates with robust Windmeijer (2005) finite-sample corrected standard errors in parenthesis. The GMM-type instruments used are lags (3-5) of the independent variable. P-values for diagnostic tests inside brackets.

In Table 1, we present the estimation results for equation (4). The estimate of the returns to scale, η, equals 1.17, but is somewhat imprecisely estimated (s.e. of 0.18). For this reason, it is reassuring to see that the point estimate of η is very similar to estimates reported by earlier studies. For example, Carlsson (2003) reports an overall estimate of the returns to scale for the Swedish manufacturing sector of 1.16 for the sample period $1968 - 1993$. Moreover, the estimate of η is economically sensible in that it implies that firms price their output higher than marginal cost, that is, since η also equals the average markup of price over

[8] Given that we use a difference-GMM estimator, the second and higher ordered lags of $\Delta\widetilde{x}$ should be valid instruments under the null of no serial correlation in the residual. However, when including the second lag in the instrument set, the Hansen test of the over-identifying restrictions is significant, although the AR(2) test of the differenced residuals cannot reject the null of no serial correlation in the residual. We proceed with caution, however, and drop the second lag from the instrument set.

marginal cost under the assumption of zero profits on average (see Basu and Fernald (1997) for a discussion).

We have also tested for different returns to scale in the durables and the non-durables sector. However, the null of equal returns to scale cannot be rejected on any reasonable level of significance. Moreover, using the full set of available instruments yields a very similar estimate of the returns to scale (η equal to 1.20 with a s.e. of 0.176) and the results presented in the tables below are robust to instead using this estimate of η.

We also see in Table 1 that the AR(2) test of the differenced residuals indicates that there is no serial correlation in the idiosyncratic technology series ε. Moreover, the Hansen test of the over-identifying restrictions cannot reject the joint null hypothesis of a valid instrument set and a correctly specified model.

As discussed in the introduction, it is interesting to see to what degree firm-level technology shocks are common across firms. To this end, we put the common effects back into the technology measure and run a regression on this measure on a full set of time dummies. The R^2 from this regression implies that the common component only explains about half a percent of the total variation in firm-level technology shocks. Thus, it is highly unlikely that the firm-level technology shocks have any noticeable general equilibrium effects or that monetary policy could have any significant influence on the results presented here. This result is in line with the findings of Marchetti and Nucci (2005) who also report a modest common component of technology shocks in Italian firm-level data. Thus, in practice, we do not need to distinguish between firm-level technology shocks and idiosyncratic firm-level technology shocks and following Marchetti and Nucci (2005), we will focus on the effects of the former concept throughout the rest of the paper.

4 The Labor-Input Response

To investigate the labor response to technology, we regress technology shocks, Δa, on labor-input growth, Δn. In Table 2, we present the results from this regression. As seen in Table 2, the contemporaneous effect on labor input from a technology shock is significantly negative. In the second column of Table 2, we have also included two lags of the technology shock.[9] We see that the contemporaneous effect is almost unchanged. Moreover, there is evidence of an initially upward sloping impulse-response function in the *level* of labor input,[10] that is, the contraction is followed by a significantly positive *growth* effect in the second period. We also see that this recovery continues in the second period. However, the coefficients for lags beyond the second are statistically insignificant. This can either be interpreted as a lack of effects beyond the second lag or, perhaps more likely, that the data is not very informative about long-run effects.[11]

Table 2: Technology Shocks and Labor Input

ε_{it}	-0.143^*	-0.124^*
	(0.011)	(0.015)
ε_{it-1}	–	0.050^*
		(0.014)
ε_{it-2}	–	0.058^*
		(0.014)

Superscript * denotes significantly different from zero at the five-percent level. Fixed-effect regressions with robust standard errors (clustered on plant). Sample 1990/1992-1996 with 1,516 plants (i.e. 10,612/7,580 observations).

The finding of a negative contemporaneous effect of technology shocks on labor input followed by a lagged positive effect is in line with Marchetti and Nucci (2005). This, in turn, is also in line with a sticky price explanation, that is, when prices are sticky, firm-level demand cannot change instantaneously and

[9] A caveat is in place since we treat ϵ as data, although ϵ is a generated regressor. However, as noted by Pagan (1984), using unlagged residuals does not affect the inference.

[10] By initially we mean between period 0 and 1, where period 0 denotes the impact period of the shock.

[11] Only using the idiosyncratic part of technology yields similar results, the difference being that although the point estimate for the first lag is still positive, it is not statistically significantly so.

since demand can be met using less inputs when technology improves, labor input initially falls. But as prices eventually adjust so does demand, and hence labor input rise, just as we would expected to have happened contemporaneously in a frictionless model.[12]

4.1 Measurement Errors In the Labor-Input Measure

An important concern associated with the finding of a negative contemporaneous effect of technology shocks on labor input is that the result may actually be driven by potential measurement errors in the labor-input variable. Since measurement errors will end up in the technology measure with a minus sign in front, and since we then regress this measure on the miss-measured labor-input variable, the measurement error will create a negative bias. Basu, Fernald, and Kimball (2004), for example, discuss this problem at length.[13]

To see this problem, consider:

$$\Delta n^{IS,Observed} = \Delta n^{True} + \epsilon^{IS,ME}, \tag{5}$$

where $\epsilon^{IS,ME}$ is a measurement error in the IS labor-input variable. Then the residual is equal to:

$$\begin{aligned} \Delta y - \eta \Delta \widetilde{x}^{IS,Observed} &= \Delta y - \eta \Delta \widetilde{x}^{True} - \eta C_N \epsilon^{IS,ME} && (6) \\ &= \epsilon^{True} - \eta C_N \epsilon^{IS,ME} && (7) \\ &= \epsilon^{IS,Observed}. && (8) \end{aligned}$$

Note though that the IV-procedure still gives a consistent estimate of η as long as the measurement error is uncorrelated with the instruments. Then, it is easy

[12] One objection to this explanation is that when summing the point estimates, we see that the employment level has not returned to its pre-shock value even after two years (although the confidence interval for the level impulse-response covers zero after two years). Thus, the degree of price stickiness needed to explain these results is potentially implausibly high. However, this result might be due to the contemporaneous effect being downward biased, thus biasing down the entire *level* impulse-response function. This is something we turn to next.

[13] Basu, Fernald, and Kimball (2004) also present calculations that suggest that their finding is not driven by measurement errors in the labor-input variable they use.

to see that we will have:

$$cov(\Delta n^{IS,Observed}, \epsilon^{IS,Observed}) = cov(\Delta n^{True}, \epsilon^{True}) - \eta C_N var(\epsilon^{IS,ME}). \quad (9)$$

Thus, the negative contemporaneous effect reported in Table 2 might actually be spurious.

To see whether our finding is driven by measurement errors, we use an independently measured firm-level measure of labor input drawn from the RAMS (Register Based Labor Statistics) data base. Since the IS and the RAMS measures are independently measured, it is very unlikely that any measurement errors are the same in the IS and the RAMS measure of labor input. Note that when using the RAMS measure of labor input, we will have that:

$$cov(\Delta n^{RAMS,Observed}, \epsilon^{IS,Observed}) = cov(\Delta n^{RAMS,\ True}, \epsilon^{True}). \quad (10)$$

Thus, measurement errors will not give rise to the negative bias in this case. Thus, we can replace Δn with Δn^{RAMS} in the regression of technology shocks on labor-input growth to evaluate the effect of potential measurement errors.

However, we also need to note that the RAMS labor measure corresponds to the employment level in November each year. Thus, in terms of dynamics, it does not give us the "annual average" effect, but instead the "end of the year" effect. Thus, the interpretation of the estimated impulse responses is not exactly the same, since part of the intra-year dynamics have already taken place when we look at the "end of the year" effect.

In Table 3, we present the results from replacing Δn with Δn^{RAMS} in the regression of technology shocks on labor input. Relative to the results from using the IS labor-input measure, we see that the point estimate of the contemporaneous effect of a technology shock is now close to zero and statistically insignificant (both with and without including lags). The lagged growth effects are significantly positive and numerically close to what we saw when using the IS labor-input measure, however.[14]

[14] Only using the idiosyncratic part of technology yields similar results, the difference once more being that although the point estimate for the first lag is still positive, it is not statistically significantly so.

Table 3: Technology Shocks and the RAMS Labor-Input Measure

ε_{it}	-0.007	0.015
	(0.011)	(0.014)
ε_{it-1}	–	0.056*
		(0.015)
ε_{it-2}	–	0.065*
		(0.013)

Superscript * denotes significantly different from zero at the five-percent level. Fixed-effect regressions with robust standard errors (clustered on plant). Sample 1990/1992-1996 with 1,516 plants (i.e. 10,612/7,580 observations).

So what do these results tell us about the effect of technology shocks on labor-input? Well, first of all the "end of the year effect" is zero. But this does not necessarily contradict that the true contemporaneous "annual average" effect is negative. In fact, under the null that the level impulse-response function is upward sloping, as indicated in Tables 2 and 3, we would expect the point estimates for the contemporaneous "end of the year" effect from RAMS to be more positive than the contemporaneous "annual average" effect, regardless of any measurement errors. What these results then indicate is that the true "annual average" effect is zero or negative.[15] Thus, although a significant part of the estimated negative contemporaneous "annual average" effect may be driven by measurement errors in the labor input variable, we find no evidence supporting that the bias stemming from measurement errors hides any significantly positive contemporaneous "annual average" effect, as expected in a frictionless model.

Moreover, in both Table 2 and Table 3, we find lagged positive effects on labor input. Once more, this is not expected, unless there is some impediment to immediate adjustment.

Overall, our results are in line with the view that there is some impediment

[15] If we only looked at the idiosyncratic part of firm-level technology shocks, we would conclude the contemporaneous "annual average" effect to be zero since we found no evidence supporting that the *level* impulse-response function is initially upward sloping in this case.

to immediate adjustment at the firm-level, as also found by Marchetti and Nucci (2005) – although it can always be debated whether the contemporaneous effect measured on annual data is negative or zero.[16] Interestingly, Marchetti and Nucci (2005) corroborate their view that these results are driven by price stickiness by presenting evidence of an inverse relationship between the degree of price rigidity and the size of the effect of technology shocks on labor input. Since there is no obvious reason to believe that the variance in measurement errors in labor input is systematically related to the degree of price stickiness, this finding thus supports the view that price rigidity does seem to play a part in explaining the results.[17]

5 Conclusions

We have estimated technology shocks on firm-level data using a robust production function approach and have studied the empirical relationship between technology shocks and labor input. When evaluating the effect on labor input from a firm-level technology shock, using the same labor input measure as used when generating the technology shock, we find that firms initially reduce their labor input in response to a technology improvement.

An important concern associated with the finding of a negative contemporaneous effect of technology shocks on labor input is, however, that this result may be driven by potential measurement errors in the labor-input variable. An important advantage of the data set we use is that it includes two independently measured firm-level labor-input measures, which allows us to see whether our

[16] Note that the size (and in its extension – even the sign) of the expected effect from a sticky price model depends on the degree of price stickiness and on which frequency we observe the data.

[17] Note, however, that this argument does not provide any evidence on the true sign of the contemporaneous effect in the Marchetti and Nucci (2005) study, only on the relative effect. That is, since the estimates for both their subsamples of firms (with stickier prices/less sticky prices) should be approximately equally negatively biased, if there are measurement errors in the labor-input variable.

results are driven by measurement errors in the labor-input variable.

When evaluating the above results using the alternative labor-input measure, we find no evidence supporting that the bias stemming from measurement errors hides any significantly positive contemporaneous effect, although the true contemporaneous effect may very well be zero. Moreover, we find lagged positive effects on labor input using both labor input measures.

Since the common component across firms in the firm-level technology component is very small, the results in this paper are not affected by the monetary policy regime in effect during the sample period. It also implies that the general equilibrium effects from this shock are negligible – thus the evidence on adjustment frictions found in this paper directly pertains to adjustment impediments in the firms' economic environment.

Moreover, since we study a balanced panel, aggregation effects working through the cleansing of inefficient units in recessions are of no concern for our results.

Overall, our results are in line with the view that there is some impediment to immediate micro-level adjustment, as also argued by Marchetti and Nucci (2005). The non-expansionary contemporaneous effect of technology shocks on labor input followed by lagged positive effects point away from the standard propagation mechanisms in frictionless models. Interestingly, Marchetti and Nucci (2005) present evidence of a link between the degree of price rigidity and the strength of the contractionary effect of technology shocks on labor input. Thus, a potentially more rewarding explanation of the findings in this paper is likely to include the notion of price stickiness.

References

ALEXIUS, A., AND M. CARLSSON (2005): "Measures of Technology and the Business Cycle," *Review of Economics and Statistics*, 87, 299–307.

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., AND J. FERNALD (1997): "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 105, 249–283.

BASU, S., J. FERNALD, AND M. KIMBALL (2004): "Are Technology Improvements Contractionary?," *Forthcomming in American Economic Review.*

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.

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

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

CHRISTIANO, L. EICHENBAUM, M., AND R. VIGFUSSON (2003): "What Happens After a Technology Shock," *Federal Reserve Board International Discussion Papper 768.*

DOTSEY, M. (2002): "Structure From Shocks," *Federal Reserve Bank of Richmond Economic Quarterly*, 88, 37–47.

FERNALD, J. (2005): "Trend Breaks, Long Run Restrictions, and the Contractionary Effects of Technology Improvements," *FRB of San Francisco Working Paper No. 2005-21.*

Francis, N., and V. Ramey (2003): "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," *forthcoming in Journal of Monetary Economics*.

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

——— (2004): "On the Role of Technology Shocks as a Source of Business Cycles: Some New Evidence," *Journal of the European Economic Association*, 2, 372–380.

King, R. G., and S. T. Rebelo (1999): "Resuscitating Real Business Cycles," in *Handbook of Macroeconomics Vol. 1B*, ed. by J. Taylor, and M. Woodford, Amsterdam. North-Holland.

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

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

Pagan, A. (1984): "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, 1, 221–247.

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

Appendix

A Data

The data set we use is primarily drawn from the Industry Statistics Survey (IS) and contains annual information for the years 1989-1996 on inputs and output for all Swedish manufacturing plants with 20 employees or more and a sample of smaller plants. For about 80 percent of the observations, there is one plant per firm so, in practice, we can think of the data as representing firms.

Our measure of real output, Y, is the value of total sales taken from the Industry Statistics Survey (IS) deflated by a firm-specific producer-price index. The firm-specific price index is compiled by Statistics Sweden using a combination of plant-specific unit values and detailed disaggregate producer-price indices. The data used to compute the firms' price index varies depending on a judgement of the quality of the unit-value data. The producer-price index for the relevant class of goods is used if e.g. the price change, as implied by the unit value data, is outside a certain acceptable range or if the data required to construct unit values is not available. The price changes for the different goods are then used to construct a price index for the individual firm. Labor input, N, is measured as the average number of employees during the year and is taken from the IS. For the Swedish manufacturing sector, Carlsson (2003) reports that the growth rate of hours per employee is acyclical. Thus, we are not likely to leave out any important variation in labor input by looking at the growth rate of the number of employees instead of total hours. Real intermediate inputs, M, are measured as the sum of costs for intermediate goods and services collected from the IS deflated by a three-digit producer price index collected by Statistics Sweden.[18] Moreover, energy, V, is measured as the plants' electricity consumption in MWh taken from the IS. Finally, when computing the cost shares, we also need a measure of the firms' labor cost, which is measured as total labor cost including e.g. collective fees available in the IS.

[18] In some rare **instances,** we had to resort to a two-digit producer price index.

Beside the survey based firm-specific measure of average annual employment available in IS, we also have access to a firm-specific measure of the number of employees in November from the Register Based Labor Market Statistics data base (RAMS). Whereas the IS employment data is based on a survey collected by Statistics Sweden, the RAMS employment data is based on the income statements that employers are, by law, required to send to the Swedish Tax Authority. Note that the IS and the RAMS measures of labor input are independently measured and it is thus very unlikely that any measurement errors are the same in the two measures of labor input.

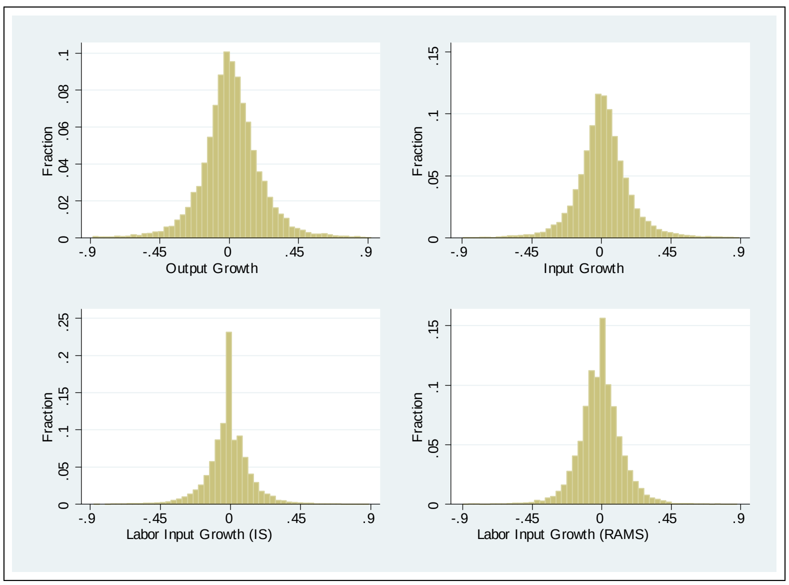

Figure 1: Data Distributions

Although we have removed obviously erroneous observations, the data set still contains very large observations. To avoid that our results are affected by plants subject to episodes of extreme conditions, these observations are removed (see below). This is likely to be especially relevant here, since Sweden experienced a boom-bust episode during the late 1980s and early 1990s. Moreover, this procedure will mitigate potential problems with large measurement errors.

In Figure 1, the data distributions are plotted for the relevant variables (truncated at ± 0.9 in the log-difference space). Since the main mass of the data seems to be well captured in the interval ± 0.45 for all variables, we limit the data set to contain firms with observations only within this interval.[19] Note that e.g. $dy = 0.45$ corresponds to an annual increase of almost 60 percent in real output.[20]

All in all, this leaves us with a balanced panel of $1,516$ plants observed over the years 1990-1996, i.e. $10,612$ observations.

[19]This leads to a reduction in sample size of 36 percent. Note, however, that we are still left with a large sample of firms $(1,516)$.

[20]The chosen intervals are slightly more limiting with respect to the distribution for dy and dx. However, making a small increase in these two intervals yields very similar results, relative to those presented in Tables 1 to 3 in the main text.

Earlier Working Papers:

For a complete list of Working Papers published by Sveriges Riksbank, see www.riksbank.se

Sveriges Riksbank
Visiting address: Brunkebergs torg 11
Mail address: se-103 37 Stockholm

Website: www.riksbank.se
Telephone: +46 8 787 00 00, Fax: +46 8 21 05 31
E-mail: registratorn@riksbank.se

ISSN 1402-9103 Riksbank Printing Office Stockholm 2006