Economics Letters 247 (2025) 112161

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

The heterogeneous effects of technology shocks. Evidence from the Czech Labour market☆

Monika Junicke [a], Jakub Matějů [b], Haroon Mumtaz [c],*, Angeliki Theophilopoulou [d]

[a] University of Economics and Business, Prague, Czech Republic
[b] Czech National Bank, Czech Republic
[c] School of Economics and Finance Queen Mary University of London, United Kingdom
[d] Brunel University of London, United Kingdom

ARTICLE INFO

JEL classification:
C32
E32
Q54

Keywords:
Technology shocks
FAVAR

ABSTRACT

This paper uses administrative labour market data from Czechia to investigate the heterogeneous effects of technology shocks. Using a FAVAR, the shock is identified using medium run restrictions à la Uhlig (2004b). The shock has a positive effect on hours for workers with wages at and above the median, while there is some evidence that workers on low wages reduce their hours. Analysis of industrial and demographic groups indicates that the former group is likely to consist of males, to be educated or to work in services.

1. Introduction

A large empirical literature has investigated the prediction of real business cycle theory that technology shocks are important for business cycle fluctuations. Earlier papers in this literature report results that are at odds with this assertion. For example, in a seminal contribution (Galí, 1999) identifies the technology shock as the only disturbance in a VAR that can affect labour productivity in the long-run. He finds that technology shocks are associated with a decline in hours worked. Similar results are reported in Francis and Ramey (2005). More recent papers have highlighted the drawbacks of VARs with long-run restrictions. For example, Uhlig (2004a) proposes a scheme based on medium run restrictions that are more computationally robust and finds a mild positive response of hours. Using sign restrictions Dedola and Neri (2007) also report that technology shocks lead to an increase in hours worked.[1]

One unifying feature of this literature is the focus on the effect of the shock on *aggregate* hours. In contrast, evidence on how the distribution of hours changes after the shock is scarce. Our paper fills this gap in the literature and investigates the possible heterogeneity in the effects of this shock across workers. In particular, we use administrative data from Czechia to show that technology shocks are associated with an increase in hours for workers towards the middle and right of the wage distribution. However, hours for low-wage workers decline after the shock. We find evidence that the positive response of hours reflects the effect of the shock on male workers and those in service industries. In contrast, the response of hours at the left tail of the wage distribution is related to the response of workers who only have primary education, those in agriculture and construction or female workers. The paper is related to Saijo (2019) who show using US data that technology shocks lead to an increase in hours for stock-holders and a decline in hours for non stock-holders. The focus of the current paper is broader as we explore the heterogeneous response of hours along the wage distribution and for demographic groups defined by gender, industry, age and education. Moreover, the micro-data in our study has substantially more comprehensive coverage than a survey.[2]

While our findings are based on the Czech data, their relevance likely extends to other small, open, and developed economies. This is because Czechia shares characteristics with these economies, including an industrial focus, consistent GDP growth, global market integration,

☆ The authors gratefully acknowledge the financial support of the Czech Science Foundation under grant No. GA CR 21-15530S for the project “Winners and losers: who benefits from the Czech national bank actions? The heterogeneous impact of monetary policy across firms and across households.” We thank an anonymous referee and the Editor Eric Young for useful comments.

* Corresponding author.
E-mail addresses: monika.junicke@vse.cz (M. Junicke), jakub.mateju@cnb.cz (J. Matějů), h.mumtaz@qmul.ac.uk (H. Mumtaz), angeliki.theophilopoulou@brunel.ac.uk (A. Theophilopoulou).

[1] For a comprehensive review of the literature see Ramey (2016).

[2] We also use identification schemes for the technology shock that have superseded the long-run restrictions employed by Saijo (2019).

https://doi.org/10.1016/j.econlet.2024.112161
Received 14 November 2024; Received in revised form 22 December 2024; Accepted 30 December 2024
Available online 11 January 2025

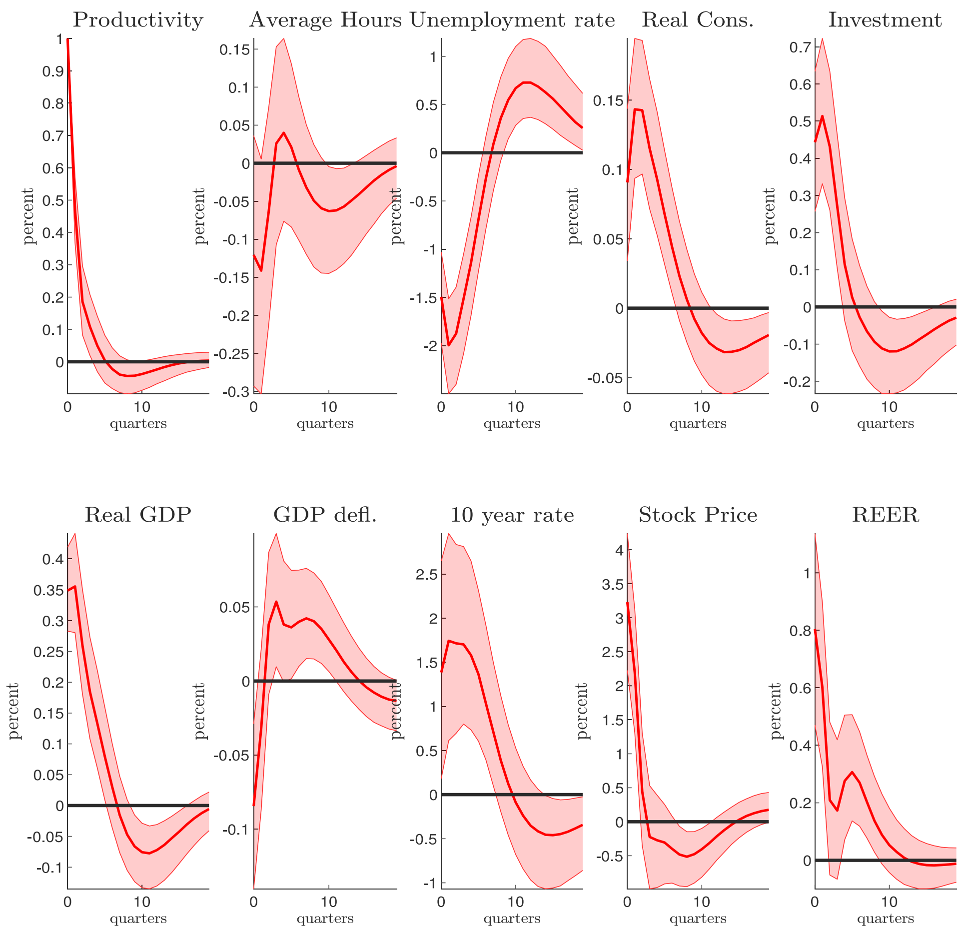

Fig. 1. Response of selected aggregate variables to a productivity shock. Real Cons. is real consumption, GDP defl. is GDP deflator inflation and REER denotes the real effective exchange rate. The shaded area displays the 68% error bands.

and EU membership. In terms of within-country income distribution, income inequality in Czechia has mirrored the broader trend in developed countries. The Czech National Bank's inflation-targeting strategy and institutional framework resemble those of other central banks.

The paper is organised as follows: Section 2 describes the empirical model and the data. The main results are presented in Section 3. Section 4 concludes.

2. Empirical model and data

In order to estimate the effects of the technology shock on the distribution of hours, we employ a Factor Augmented VAR (FAVAR) (see Bernanke et al. 2005). The model is defined by the VAR:

$$Y_t = BX_t + u_t, \quad (1)$$

$$u_t \sim N(0, \Sigma) \quad (2)$$

where $Y_t = \begin{pmatrix} Z_t \\ F_t \end{pmatrix}$. Z_t is a measure of productivity for Czechia, while F_t denotes a set of common factors extracted from both aggregate and individual-level data. The vector $X_t = [Y'_{t-1}, \ldots, Y'_{t-P}, 1]'$ is $(NP+1) \times 1$ and defines the regressors in each equation and B denotes the $N \times (NP+1)$ matrix of coefficients $B = [B_1, \ldots, B_P, c]$.

The observation equation of the model is defined as:

$$\begin{pmatrix} Z_t \\ \tilde{X}_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & \Lambda \end{pmatrix} \begin{pmatrix} Z_t \\ F_t \end{pmatrix} + \begin{pmatrix} 0 \\ v_t \end{pmatrix} \quad (3)$$

$\tilde{X}_t$ is a $M \times 1$ vector of variables that include quarterly aggregate measures of macroeconomic and financial conditions. $\tilde{X}_t$ also contains measures of hours constructed using individual-level data. These series include average hours in groups defined by each wage decile, groups defined by level of education, age, gender and industry. These disaggregated series are available at an annual or bi-annual frequency.

Finally, v_t is a $M \times 1$ matrix that holds the idiosyncratic components which are assumed to be autocorrelated and follow an $AR(q)$ process. Note that the idiosyncratic components corresponding to the hours data can be interpreted as shocks that are specific to those groups and also capture possible measurement error in the individual-level data. In contrast, the shocks to Eq. (1) represent macroeconomic shocks that are of interest in this exercise. This ability to estimate the impact of macroeconomic shocks while accounting for idiosyncratic disturbances is a key advantage of the FAVAR over a VAR where these two sources of fluctuations may be conflated (see Giorgi and Gambetti 2017). Moreover, by incorporating a large data set, the FAVAR reduces the problem of information deficiency (see e.g. Forni and Gambetti 2014). In addition, the model allows us to easily incorporate lower frequency data on hours that are only available annually before 2012 and twice a year, thereafter. As described in the appendix, we assume that these observations are averages of quarterly hours that are treated as additional latent states variables in the FAVAR.

2.1. Identification of the productivity shock

Uhlig (2004a) shows that an identification scheme based on medium-run restrictions performs better than long-run identification schemes (Galí, 1999) in recovering the productivity shock (see also Uhlig (2004b)). This long-run scheme has the drawback that the infinite

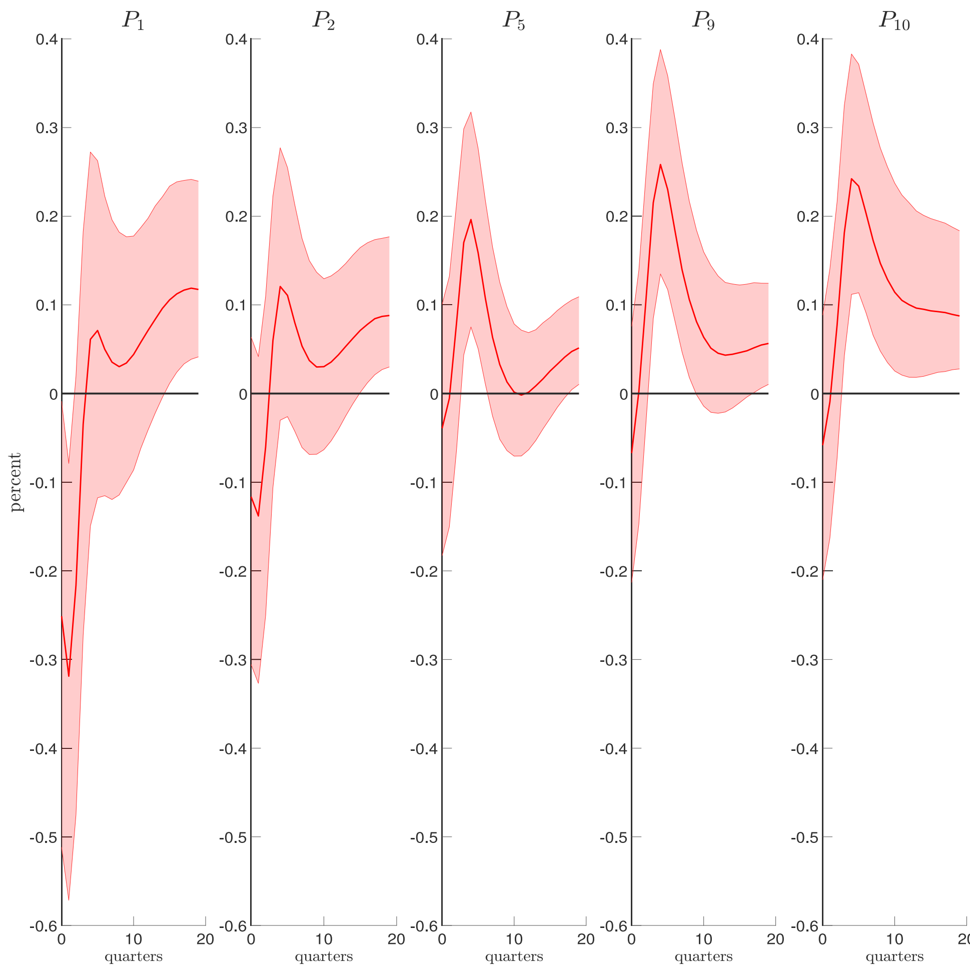

Fig. 2. Response of hours in wage decile groups. For e.g. P_1 denotes the first decile, while P_{10} is the last decile. The shaded area displays the 68% error bands.

horizon impulse response has to be estimated and this can be challenging using a short span of data (see Erceg et al. 2005). The method of Uhlig (2004a) is less susceptible to this computational issue as it works with medium horizons. We adopt this strategy as our benchmark approach.

The structural shocks are defined as $\varepsilon_t = A_0^{-1} u_t$ where $A_0 A_0' = \Sigma$. A_0 is not unique and the space spanned by these matrices can be written as $\tilde{A}_0 Q$ where Q is an orthonormal rotation matrix such that $Q'Q = I$. The productivity shock is identified by imposing the restriction that this shock makes the largest contribution the forecast error variance (FEV) of Z_t at the one year horizon. The appendix provides details of this calculation and also shows that our main results do not depend on the identifying scheme.

2.2. Data and estimation

The data set $\tilde{X}$ consists of 84 aggregate series. In addition, we include 23 series on hours constructed from the labour market data described below. The aggregate series (listed in the appendix) cover the main sectors of the economy: Real activity, inflation, money, credit and finance. The series are quarterly from 2002Q1 to 2019Q4. The source of the data is the FRED database and Global Financial database. All non-stationary series are transformed by taking log-differences.

Labour market data at the individual level is obtained from the administrative ISPV dataset, which provides rich contract-level information at annual frequency from 2002 to 2011 and bi-annual frequency thereafter to 2020.[3] The data covers a large proportion of all Czech labour market contracts. The key variables include wage per-hour, hours worked and employee characteristics, such as gender, age, and education level. We merge the data with the RES Business Register database, which provides information on the business sector in which the employer operates.[4]

We construct average hours in groups defined by a number of characteristics. First, we consider 10 groups defined by the deciles of wage that are denoted by $P_1, \ldots, P_{10}$. In addition, we construct groups based on the following characteristics:

1. Education: The averages are calculated from individuals with primary, secondary or tertiary level education.
2. Age: We consider three age groups: individuals less than 35 years of age, between the ages of 35 and 50 and older than 50 years.
3. Sector of employment: We include average hours in Agriculture, Manufacturing, Construction and Services.
4. Gender: We construct average hours for males and females.

We also construct average hours across all respondents as a measure of economy-wide hours. We include these series in logs. However, as discussed below, the results are robust to using log differences.

The model is estimated using a Gibbs sampling algorithm that is described in details in the appendix along with the prior distributions

[3] Informační systém o průměrném výdělku (Average Earnings Information System)

[4] Registr ekonomických subjektů (Business Register). See the on-line appendix for detailed description of the data

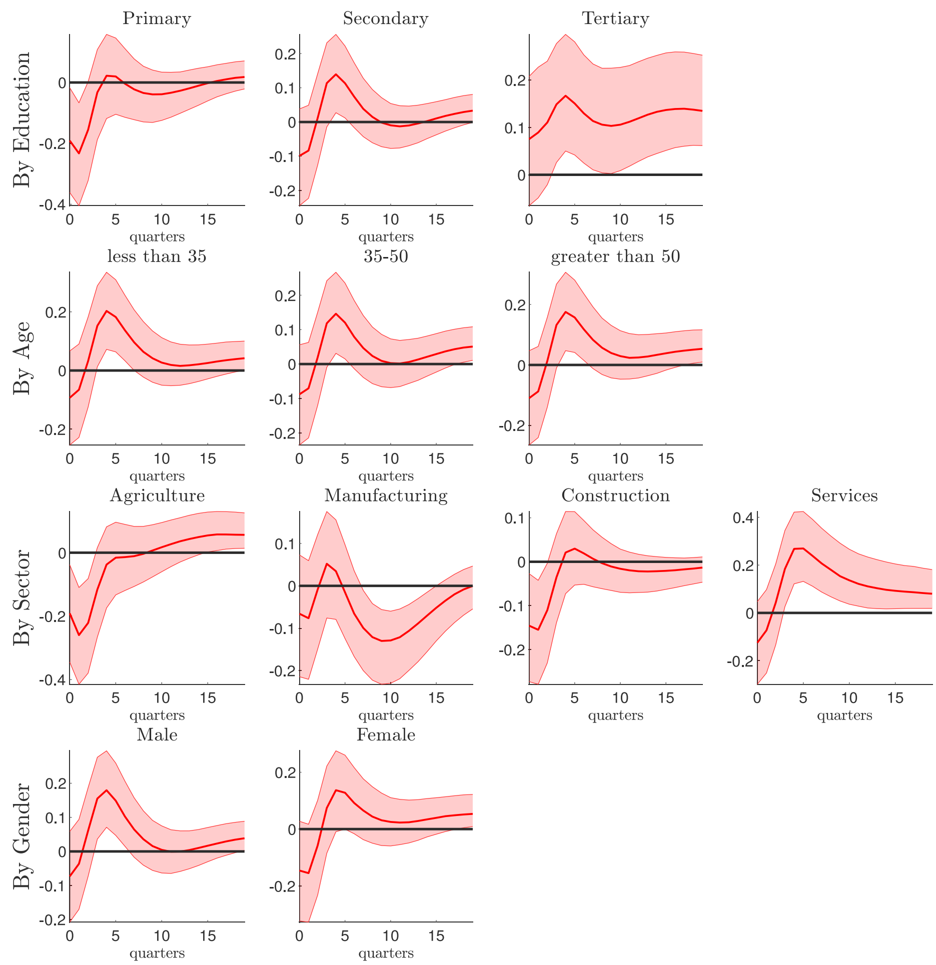

Fig. 3. Response of hours for different demographic groups and sectors. The shaded area displays the 68% error bands.

and convergence diagnostics. We set the number of factors to 6 based on the IC_{p2} criterion of Bai and Ng (2002). The lag lengths in Eq. (1) is fixed at 4 while the idiosyncratic components follow an AR(1) process.

3. Results

Before moving on to the distributional response of hours, it is instructive to consider the response of aggregate variables shown in Fig. 1. The productivity shock has an ambiguous effect on average hours: the median response is negative over the medium horizon and at odds with real business cycle theory, albeit with large error bands.[5] The shock is expansionary and increases real activity, stock prices and long-term interest rates, while pushing down inflation in the short term. The shock is associated with a real exchange rate appreciation supporting the findings of Enders et al. (2011) reported for the US.

3.1. Heterogeneous effects of the shock

Fig. 2 presents our key result. The figure shows the response of hours in selected wage deciles on the left and right tail of the wage distribution. The response of hours for individuals that earn wages below the median resembles the aggregate hours response shown in Fig. 1 with the median showing a decline at short and medium horizons. In contrast, hours increase towards the right tail of the wage distribution and the response is statistically different from zero about one year after the shock.[6]

Fig. 3 investigates this heterogeneity further by exploring the response of hours by demographic groups and sectors. The top panel shows the negative response of hours is more evident for workers with primary education who are concentrated at the left tail of the wage distribution (see Figure 1 in the online appendix). The second row of the Figure shows that there is limited heterogeneity across the age distribution. In contrast, the response clearly varies across sectors. Hours decline for workers in agriculture and construction. The response for manufacturing is negative at medium and long-horizons, while hours display an increase in the services sector that is statistically different from zero at the one year horizon. As services dominate the right tail of the wage distribution while manufacturing and agriculture is more prevalent towards the middle and left tail (see Figure 4 in the appendix), these impulse responses are consistent with the distributional results in Fig. 2. The final row of the figure shows that it is male workers that increase hours after the shock. As Figure 1 in the appendix shows, male workers are substantially more likely to be high wage earners.

Broadly speaking, these results are consistent with skill biased technological change whereby the shock disproportionally increases the productivity and demand for high skilled workers. Hours may increase for these workers if they take advantage of higher returns to skill. The heterogeneity may also be driven by stock holding as in Saijo (2019).

[5] This negative response is consistent with either the presence of aggregate price stickiness and non-accommodative monetary policy (Galí, 1999) or real rigidities (Francis and Ramey, 2005). Slanicay (2016) provides evidence that prices are more sticky in Czechia as compared to the Euro-Area.

[6] The online appendix presents the results for all 10 decile groups.

If high earners hold stocks, then they may increase hours worked to benefit from the wealth effect generated by the technology shock.

3.2. Robustness

We carry out a number of robustness checks that are presented in detail in the on-line appendix:

1. Identification: We identify the productivity using the sign restrictions methodology in Dedola and Neri (2007) and Francis et al. (2003). Figure 8 in the appendix shows that the distributional pattern of the hours response is similar to benchmark. The results are less precise when long-run restrictions are used to identify the productivity shock (Figure 9 in the appendix). This unsurprising given the short span of our data. However, the median responses accord well with our benchmark results.
2. Specification: Figure 10 in the appendix shows that the results are very similar to benchmark when the number of factors is increased to 8. We also estimate the model using hours in log-differences. The results in this case are supportive of the benchmark and show an increase in hours towards the right tail of the wage distribution, in services, for males and those with higher than primary eduction.

4. Conclusions

Using administrative labour market data for Czechia, this paper shows that technology shocks have a heterogeneous effect on hours worked. Hours increase for high earners and decline for workers on low wages. The former group appears to consist of more educated individuals, male workers and those employed in services.

Data availability

Data will be made available on request.

References

Bai, J., Ng, S., 2002. Determining the number of factors in approximate factor models. Econometrica 70, 191–221.

Bernanke, B.S., Boivin, J., Eliasz, P., 2005. Measuring the effects of monetary policy: A factor augmented vector autoregressive (FAVAR) approach. Q. J. Econ. 120, 387–422.

Dedola, Luca, Neri, Stefano, 2007. What does a technology shock do? A VAR analysis with model-based sign restrictions. J. Monetary Econ. 54 (2), 512–549.

Enders, Zeno, Müller, Gernot J., Scholl, Almuth, 2011. How do fiscal and technology shocks affect real exchange rates?: New evidence for the United States. J. Int. Econ. 83 (1), 53–69.

Erceg, Christopher, Guerrieri, Luca, Gust, Christopher, 2005. Can long-run restrictions identify technology shocks? J. Eur. Econom. Assoc. 3 (6), 1237–1278.

Forni, Mario, Gambetti, Luca, 2014. Sufficient information in structural VARs. J. Monetary Econ. 66 (C), 124–136.

Francis, Neville, Owyang, Michael T., Theodorou, Athena T., 2003. The use of long-run restrictions for the identification of technology shocks. Review 85 (Nov), 53–66.

Francis, Neville, Ramey, Valerie A., 2005. Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. J. Monetary Econ. 52 (8), 1379–1399.

Galí, Jordi, 1999. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? Amer. Econ. Rev. 89 (1), 249–271.

Giorgi, Giacomo De, Gambetti, Luca, 2017. Business cycle fluctuations and the distribution of consumption. Rev. Econ. Dyn. 23, 19–41.

Ramey, V.A., 2016. Chapter 2 - Macroeconomic shocks and their propagation. In: Taylor, John B., Uhlig, Harald (Eds.), Handbook of Macroeconomics. vol. 2, Elsevier, pp. 71–162.

Saijo, Hikaru, 2019. Technology shocks and hours revisited: Evidence from household data. Rev. Econ. Dyn. 31, 347–362.

Slanicay, Martin, 2016. Analysis of structural differences and asymmetry of shocks between the czech economy and the euro area. Stat. Stat. Econ. J. 96, 34–49.

Uhlig, Harald, 2004a. Do technology shocks lead to a fall in total hours worked? J. Eur. Econom. Assoc. 2 (2–3), 361–371.

Uhlig, Harald, 2004b. What Moves GNP? Econometric Society 2004 North AmericanWinter Meetings 636, Econometric Society.