UPPSALA
UNIVERSITET

Department of Economics
Working Paper 2018:10

DEEP DYNAMICS

Nils Gottfries, Glenn Mickelsson and Karolina Stadin

Department of Economics
Uppsala University
Box 513
751 20 Uppsala
Sweden

Working Paper 2018:10
June 2018
ISSN 1653-6975

DEEP DYNAMICS

NILS GOTTFRIES, GLENN MICKELSSON AND KAROLINA STADIN

DEEP DYNAMICS*

Nils Gottfries[#], Glenn Mickelsson[##] and Karolina Stadin[###]

June 7, 2018

Combining micro and macro data, we construct demand-side shocks, which we take to be exogenous for individual firms. We estimate a reduced-form model to describe how firms adjust their production, employment, capital stock, and inventories in response to such shocks. Then, we chose the structural parameters of a theoretical model so that the theoretical model can match the impulse-response functions from the estimated reduced-form model. Firms' reactions to demand-side shocks are well explained by a model where firms have modest market power, face convex adjustment costs and where they can vary utilization flexibly. The stock-out motive helps to explain inventory dynamics.

Keywords: capacity utilization, production factor, labor hoarding, labor productivity, inventory holdings, returns to scale, production function, Solow residual
JEL codes: E22, E23, E24, E32

[#] Department of Economics, Uppsala University, UCLS, IZA and CESifo, nils.gottfries@nek.uu.se
[##] Department of Economics, Uppsala University, glenn.mickelsson@nek.uu.se
[###] Ratio Institute, Stockholm, and UCLS, Uppsala University, karolina.stadin@nek.uu.se
*This paper builds on Stadin (2014) and Mickelsson and Stadin (2016). We have received helpful comments from many colleagues, including Martin M. Andreasen, Mikael Carlsson, Daria Finocchiaro, Peter Fredriksson, Georg Graetz, Bertil Holmlund, Per Krusell, Jesper Lindé, Florin Maican, Ingvar Strid, Karl Wallentin, Erik Öberg, and seminar participants at the Greater Stockholm Macro Group, the Ratio Institute, Stockholm School of Economics and the Universities of Lund, Stockhom, Uppsala, and Örebro. Financial support from the Wallander-Hedelius-Browaldh Foundation and the Marianne and Marcus Wallenberg Foundation is gratefully acknowledged.

1. Introduction

Understanding how firms in general react to shocks is important for understanding business cycles and the role of stabilization policy. Alternative theories provide potential explanations of key stylized facts, such as the pro-cyclicality of investment, labor input, labor productivity, and inventory holdings. In this paper, we investigate the relevance of some of the explanations that have been proposed by estimating a structural model of a manufacturing firm using panel data.

Pro-cyclical labor and factor productivity has been documented in many studies, and the Solow residual has been used to measure technology in the real business cycle literature (Prescott, 1986) but many researchers have questioned the interpretation of the Solow residual as a technology shock.[1] Hall (1988) considers variations in labor input and production that are related to shocks (military spending, oil prices, and the political party of the president) that should be uncorrelated with technology shocks. He shows that the variations in labor input that are associated with these shocks lead to more than proportional changes in production and he interprets this as evidence of increasing returns to scale. With increasing returns, firms will make losses if the price is equal to the marginal cost, and since firms typically do not make losses, not even in periods of low demand, Hall concludes that firms must have very substantial market power. An alternative explanation is that there are costly variations in factor utilization (Burnside, Eichenbaum, Rebelo, 1993; Sbordone, 1996, 1997). If we do not properly account for the cost of increasing utilization, the marginal cost will be underestimated, and the markup will be overestimated.

Another stylized fact is that inventory holdings are pro-cyclical. If inventories were held solely to smooth production in the face of demand-side shocks, we would expect inventories to decrease in periods of high demand. Hence, some researchers have viewed pro-cyclical inventory investment as an indication that the cost of producing must be low in boom periods, perhaps due to technology shocks, increasing returns or positive externalities.[2] Alternatively, a stock-out avoidance motive can explain pro-cyclical inventory holdings in the face of demand-side shocks (Kahn 1987, 1992; Bils and Kahn 2000). The basic idea is that

[1] Hart and Malley (1999), Baily, Bartelsman and Haltiwanger (2001), and Field (2010) document pro-cyclical productivity for different countries and time periods. The literature on the “paradox of short run increasing returns to labor” goes back many years; see Fay and Medoff (1985) and Biddle (2014) for reviews.
[2] See Bliner and Maccini (1991) for a review of the earlier literature on inventories.

firms need to have stocks of a variety of finished goods on the shelves in order to sell, and to satisfy higher demand, they need to have more goods on the shelves. More recent studies finding evidence in line with the stock-out avoidance theory are Galeotti, Maccini and Schiantarelli (2005), Wen (2005) and Kryvtsov and Midrigan (2013).

In this paper, we use Swedish firm-level data to estimate firms´ responses to specific demand-side shocks and we build a theoretical model that matches those responses. Our study does not directly address the question of the relative importance of supply and demand shocks for business cycle fluctuations, but we obtain estimates of structural parameters that can be used as benchmarks in the construction of macroeconomic models. Our analysis proceeds in four steps:

First, we construct a product market *demand index* that varies across firms because the shares of production that are used for domestic consumption and investment vary across industries and because the share of the firm's production that is sold in the export market varies across firms. This approach is similar to Hall (1988) and Kryvtsov and Midrigan (2013) in that we try to measure demand-side shocks that should be unrelated to technology shocks and cost shocks that affect individual firms or industries.

Second, we try to capture the empirical responses of firms by estimating a *reduced-form model* using panel data for manufacturing firms. The empirical model includes production, the capital stock, the number of employees, the inventory stock, and the demand index. The endogenous variables depend on their own lags and on the demand index, which we take to be exogenous for the individual firm. We include firm and time fixed effects in the estimation.[3] As explained in Section 2, the basic idea is that the effects of omitted state variables are "mopped up" by the lags of the variables that we can observe. We find that firms react strongly to the demand shocks that we have constructed. While production and inventory holdings respond quickly to the demand shock, registered inputs respond with very substantial lags. This implies positive responses of factor and labor productivity (as they are normally measured) to demand-side shocks.

Third, we construct a *theoretical model* that incorporates many of the explanations for demand-driven fluctuations in labor productivity and inventory holdings that have been suggested in the literature. We assume that firms have market power and that they face adjustment costs and implementation lags in hiring and investment. We allow firms to vary the utilization of both labor and the capital stock at a cost. Furthermore, workers can spend

[3] The model is a reduced-form model of the firm in the sense that all variables on the right-hand side are assumed to be either predetermined or strictly exogenous relative to the firm's decision variables.

time on activities such as organizing and training, which increase “organizational capital” and future production, but do which not increase production in the current period. Inventories of finished goods are needed to prevent stock-outs, but they can also be used to smooth production, and part of the inventory stock consists of inputs.

In the fourth step, we investigate the relevance of the different theoretical mechanisms by estimating the *deep structural parameters* of our model. We follow the approach of Christiano, Eichenbaum and Evans (2005) by choosing the structural parameters to match the estimated impulse-response functions from our empirical reduced-form model. Confidence intervals are calculated by bootstrapping, i.e., resampling from the population of firms with replacement and re-estimating the parameters.

We find that our theoretical model can explain the estimated response to the shock very well. A strong response of production to a demand shock is explained by firms having market power and firm-level demand being very sensitive to the demand shock as we measure it. Slow adjustment of labor and capital is explained by convex adjustment costs and implementation lags (time to build). Production increases rapidly with the demand shock, and most of the short-run adjustment is achieved by increasing utilization. According to our estimates, increasing returns to scale in production play a small role. Inventory investments respond positively to demand-side shocks, and this “accelerator effect” on inventory investments is explained partly by the stock-out motive, which affects the holdings of finished-goods inventories, and partly by the fact that a substantial fraction of the inventory stock consists of intermediate goods and goods in process, which are necessary for production.

As far as we know, this is the first paper to estimate a structural model of the joint dynamics of production, capital stock, employment, and inventory holdings using panel data for firms. These decisions are intimately linked, so it makes sense to model them jointly. Our approach to identification follows Hall (1988) in that we try to isolate movements in the endogenous variables that are caused by specific demand-side shocks, which should be orthogonal to productivity shocks and cost shocks. In terms of the estimation, we follow Christiano, Eichenbaum and Evans (2005), estimating the structural parameters by matching empirical impulse-response functions. But contrary to these studies, we use micro data instead of macro data.

An alternative method is to estimate the structural equations directly on the data, which requires specific assumptions about the unobserved shocks. Galeotti, Maccini and Schiantarelli (2005) take this approach. In a closely related paper, they estimate a model of inventories, employment and hours worked on industry data. In our view, there are three advantages of our approach compared to the direct estimation of a structural model. First, we focus on movements in the endogenous variables that are caused by exogenous shocks rather than all the variation in the endogenous variables. Second, by estimating a reduced form model, we can study the effects of a particular exogenous shock while remaining agnostic about what the other shocks and unobserved state variables are. In this way, we allow the data to speak more freely compared to if we were to estimate a tight structural model directly on the data. The third advantage is that by comparing the impulse-response functions in the theoretical model to those estimated using the unconstrained reduced form model, we can see clearly why one version of the structural model can explain the dynamics, while other versions fail. For example, we find that a model with increasing returns to scale in production but no variation in utilization can explain the “excess” response of production compared to the labor input at one horizon but not the whole profile of the impulse-response functions (see Section 6).[4]

Our approach to the estimation is further motivated in Section 2. In Section 3, we present the data and the estimated reduced-form model. A theoretical model of the firm is presented in Section 4, and Section 5 explains how we estimate the structural parameters. The estimated structural model is presented in Section 6, and in Section 7, we relate our results to previous research. Section 8 concludes.

2. Using a reduced-form model to find firms’ responses

Studying firm dynamics is difficult. To fully characterize a firm’s dynamic optimization problem, we would need to observe a large set of state variables that are relevant to the firm’s decisions. The problem is that we cannot observe all the relevant state variables, and estimating the decision rules without some of the state variables will lead to biased estimates.[5] In this study, we attempt to represent the relevant set of state variables using lagged values of the variables that we can observe. By estimating a reduced-form model of the firm, with product demand modeled as a separate stochastic process, we determine how firms respond to demand shocks that are constructed to be exogenous relative to the firm and the industry.

[4] Other differences compared to Galeotti, Maccini and Schiantarelli (2005) are that we treat the price and the capital stock as endogenous and that we model the effect of stock-outs on sales explicitly.

[5] The same argument applies to Euler equation estimations. If future and past employment help to explain current employment for a given wage, we may interpret it as an indication of adjustment costs, but this result may also be due to some omitted state variable. Consequently, it is very likely that the estimation of Euler equations leads to overestimation of the adjustment costs.

Aggregate state variables are "mopped up" by time dummies. By choosing the structural parameters so that the firm's response in the theoretical model matches the empirical response, we obtain estimates of the parameters in our theoretical model.

To see how this might work (or not work), consider *as an example* a standard model of a firm with quadratic adjustment costs related to changes in labor and capital. The firm faces a downward-sloping demand curve, and production is given by the production function: $y_t = (1-\alpha) n_t + \alpha k_{t-1} + a_t$, where y_t is production, n_t is employment and k_t is the capital stock at the end of the period *t*, and a_t is factor productivity (all variables are logs and firm-specific). The parameter α cannot be directly inferred from accounting data because the markup is unknown. Assume that there are two exogenous state variables that matter for the firm: a demand shifter d_t, which we can observe, and factor productivity a_t, which we cannot observe. Assume that the logs of the exogenous state variables follow AR(1) processes: $d_t = \rho_d d_{t-1} + \varepsilon_{dt}$. and $a_t = \rho_a a_{t-1} + \varepsilon_{at}$ where $E(\varepsilon_{dt}\varepsilon_{at}) = 0$.

The approximate solution to the firm's dynamic optimization problem consists of two log-linear decision rules relating current employment and the capital stock at the end of the period to the initial levels of capital and employment as well as demand and factor productivity. We also add white noise shocks to the decision rules:

$$n_t = b_{11} n_{t-1} + b_{12} k_{t-1} + b_{13} d_t + b_{14} a_t + \varepsilon_{nt} \tag{1}$$

$$k_t = b_{21} n_{t-1} + b_{22} k_{t-1} + b_{23} d_t + b_{24} a_t + \varepsilon_{kt}\,. \tag{2}$$

Now, we cannot estimate these decision rules because we do not observe a_t. However, we can use the equation for the productivity process to substitute for current productivity and then the production function in period t-1 to substitute for lagged productivity. By doing so, we obtain a reduced-form model with shocks that are serially uncorrelated and demand as an exogenous driving force:

$$n_t = b_{11} n_{t-1} + b_{12} k_{t-1} + b_{13} d_t + b_{14} \rho_a \left(y_{t-1} - (1-\alpha) n_{t-1} - \alpha k_{t-2} \right) + b_{14} \varepsilon_{a,t} + \varepsilon_{n,t} \tag{3}$$

$$k_t = b_{21} n_{t-1} + b_{22} k_{t-1} + b_{23} d_t + b_{24} \rho_a \left(y_{t-1} - (1-\alpha) n_{t-1} - \alpha k_{t-2} \right) + b_{24} \varepsilon_{a,t} + \varepsilon_{k,t} \tag{4}$$

$$\begin{aligned} y_t &= (1-\alpha)\left[b_{11} n_{t-1} + b_{12} k_{t-1} + b_{13} d_t + \varepsilon_{nt} \right] + \alpha k_{t-1} \\ &\quad + \left(1 + (1-\alpha) b_{14}\right)\left[\rho_a \left(y_{t-1} - (1-\alpha) n_{t-1} - \alpha k_{t-2} \right) + \varepsilon_{at} \right] \end{aligned} \tag{5}$$

$$d_t = \rho_d d_{t-1} + \varepsilon_{dt}. \tag{6}$$

Thus, we have effectively "mopped up" the effect of the initial firm-specific productivity level by including lagged values of production and capital on the right hand side.[6] If productivity follows an AR(2) process, we can account for this in the same way by including additional lags of production, capital and labor input. This reduced-form model can be estimated and used to trace out the effects of the demand shock.

In general, there are many unobserved state variables, so linear combinations of observed state variables will be imperfect representations of the unobserved state variables. Thus, the "mopping up" will be less than perfect, but we can still can hope that our reduced-form model captures firms' dynamic responses to the demand shocks in a rough way.

An alternative would be to estimate a standard vector-autoregressive model with the endogenous variables and then use d_t as an instrument for the shocks (see Gertler and Karadi 2015; Ramey 2017). The advantage of including the exogenous variable d_t explicitly in the system is that the demand shocks are not mixed up with other shocks.[7]

3. Data and empirical model

In this section, we present the firm-level data and the construction of the demand index followed by the presentation of the empirical model and the estimated impulse-response functions.

3.1 Firm-level data

The firm-level panel consists of yearly data from Statistics Sweden for all firms in Sweden. As described below, our main sample consists of firms with at least ten employees in the manufacturing sector from 1997-2008. Firms merge, split and buy plants from each other, and it is not obvious when a firm is different enough that it should be regarded as a new firm. In this study, we are interested in how established firms react to changes in product demand, so we want to diminish the noise caused by firms merging or buying and selling establishments. We therefore use the FAD units from Statistics Sweden to identify firms. The FAD units are based on legal organizational numbers, but the FAD number changes when there are mergers

[6] Note that the coefficient relating current production to lagged production $\left(1+\left(1-\alpha\right)b_{14}\right)\rho_a$ reflects the autoregressive character of the productivity shock *and* the indirect effect of productivity on hiring.

[7] Levinsohn and Petrin (2003) and Ackerberg, Caves and Frazer (2015) invert the input demand function $m_t = f\left(n_t, k_{t-1}, a_t\right)$ and use the result to substitute for the unobserved productivity shock. The production function is then estimated in a two-step procedure as in Olley and Pakes (1996). This procedure requires input prices to be the same for all firms and does not allow for variations in utilization. Our main focus is on modeling dynamics. As it turns out, we are not very successful in estimating the parameters of the production function.

or splits affecting more than 50 percent of the workforce, even if the legal organizational number remains the same.[8] When we say "firm" below, we refer to the FAD identity.

Real production (*Yr*) is the value of the firm's total production deflated by the producer price for the industry. As a robustness check, we instead use real value added deflated by the value-added deflator for the industry (*VAr*) to measure production. The real inventory stock (*Zr*) is the value of the firm's inventories at the end of the year deflated by the producer price for the industry. We do not use firm-level prices because such prices are only available for a small subset of the firms and we want to include as many firms as possible in our estimation.

N is the number of persons employed by the firm, reported as "full-year equivalents". This measure takes into account people working part-time or who are employed for only part of the year, but it does not consider overtime. This means that variations in "utilization" in our theoretical model will reflect variations in official and unofficial overtime as well as other forms of factor utilization such as effort per hour. Still, we think that employment data are closely related to official hours worked. For manufacturing as a whole, the correlation between yearly changes in registered hours worked and the number of people employed is 0.91 (see *Figure A1* in the Appendix).[9]

The real capital stock (*Kr*) consists of machines and buildings. Generous depreciation allowances imply that the book value is much lower than the economic value. Therefore, we constructed economic capital stocks using the perpetual inventory method. For a detailed description, see the Appendix.

3.2 Sample selection

In this study, we are interested in profit-maximizing firms that produce differentiated products using labor and capital and that have substantial inventory stocks consisting primarily of goods they have produced and inputs, rather than goods which are only traded. For this reason, we chose to study firms in the manufacturing sector (industries 15-36 according to SNI92). We include only firms that have at least ten employees in all their years of existence. We take this approach partly because export data are missing for many small firms, and partly because the dynamic responses of very small firms are likely different from those of medium-sized and large firms; log changes in employment can easily become very large when firms

[8] Further information on the definitions can be found in the document "Företagens och arbetsställenas dynamik (FAD)" from Statistics Sweden.
[9] Yet another problem is that firms may use labor that is provided by staffing companies but we cannot take account of this because of lack of data.

are small.[10] Publicly owned firms are dropped because they may have different objectives than privately owned firms. With these exclusions, we obtain a sample of 44-55 thousand observations, depending on what variable we consider. *Table 1* shows some descriptive statistics for this sample. The first columns show the statistics for the levels of the variables, and the latter columns show some ratios.

To create a sample of reasonably homogenous firms and to address measurement errors, we exclude *firms* that *in some period* had “extreme” levels of production per worker, inventory stock relative to production or capital stock relative to production. With one exception, “extreme” is defined as being in the bottom or top 5 percent of the sample of all firm-year observations. The exception is the lower limit for the inventory ratio, which we set to the 25th percentile (6.35 percent of yearly production) because we want to study firms which have substantial inventory stocks.[11] These cutoff limits are shown in *Table 1*. In the baseline estimation, we include only firms for which we have no missing observations 1997-2008 and this leads to a balanced panel with 818 firms and 9816 observations. We use only firms that have a complete data for all years in order to reduce the “Nickell bias” in the estimation with fixed effects for firms.[12]

3.3 The firm-specific demand index

We construct a firm-specific demand index, $D_{i,t}$, as a weighted average of domestic and international demand for the relevant industry using the firm’s average export share. The demand index is constructed to be as exogenous as possible to the firm and the industry by using only data for components of aggregate demand, data for foreign activity, and weights that do not vary over time.[13]

---

[10] By including only firms which have at least ten employees in all years we reduce attrition bias.

[11] Log changes can become very large if stocks are low. Firms with very large inventory stocks may be involved in extensive trading in addition to producing and storing their own products. Firms with very large capital stocks may be involved in property investment.

[12] The estimation method is OLS. Nickell bias means that the estimated coefficient for the first lag of the dependent variable tends to be underestimated because some of the variation is instead picked up by the firm fixed effects when there are few observations in the time dimension. We tried to deal with the Nickell bias by conducting diff-GMM estimation (Arellano-Bond), but we were unable to find an instrument set that is both valid and sufficiently relevant to provide a good identification.

[13] Similar approaches have been used by Lundin et al. (2009), Carlsson, Eriksson and Gottfries (2013), and Eriksson and Stadin (2017). The demand index used here was constructed by Stadin (2015).

Table 1. Descriptive statistics for firm-year observations, manufacturing

Total sample	**Yr**	**N**	**Zr**	**Kr**	**Yr/N**	**Zr/Yr**	**Kr/Yr**	**VAr/N**	**Zr/VAr**	**Kr/VAr**	**VA/Y**	**Kb/Y**
Mean	230000	113	29700	104000	1456	0.1471	0.4428	506	0.5023	1.1649	0.3760	0.2116
Std. d.	1510000	512	159000	727000	1126	0.1953	2.8488	325	10.2011	25.4607	3.7632	0.9830
1%	5804	10	0	65	380	0	0.0040	108	0	0	0.0605	0
5%	8862	12	123	966	***540***	0.0071	***0.0425***	***252***	0.0116	***0.1118***	0.1620	0.0109
25%	18500	18	1556	5013	827	***0.0635***	0.1680	359	***0.1554***	0.4414	0.2936	0.0603
50%	*40600*	*32*	*4952*	*13900*	*1149*	*0.1227*	*0.3305*	*447*	*0.3317*	*0.8746*	*0.3922*	*0.1485*
75%	112000	73	15600	41500	1714	0.1923	0.5730	576	0.5660	1.4813	0.4934	0.2778
95%	699000	365	89200	274000	***3346***	***0.3643***	***1.1454***	***945***	***1.2439***	***3.0839***	0.6477	0.5903
99%	3040000	1208	394000	1610000	5517	0.6057	1.9851	1571	2.3665	5.9384	0.7862	1.0147
Observations	49289	54818	54035	49156	49289	49286	43718	50909	54817	45783	49994	49994

Note: Full panel with all private firms in manufacturing in Sweden with at least 10 employees for all years of their existence. The industries included are SNI 15-36, and the years included are 1996-2008. X% denotes the Xth percentile, Yr is real production (output), N is full-time equivalent employees, Zr is real inventory stock, Kr is the real capital stock, and VAr is real value added. Real values are in thousands of SEK in prices as of year 2000. The PPI for the two-digit industry is used to deflate Y and Z, and the value-added deflator is used to calculate the real value added. The calculation of the real capital stock is described in the text. Kb/Y is the nominal book value of the capital stock relative to the nominal value of production. Boldface numbers are the limits used to delineate the sample used for the baseline estimation.

To motivate the demand index, let us consider an economy where goods produced in *J* different industries (indexed *j*) are used for consumption and investment. Let aggregate investment be a Cobb-Douglas aggregate of composite goods produced in different industries, where the latter are CES aggregates of goods produced for investment by different firms within the industry:

$$I = \prod_{j=1}^{J} I_j^{\theta_j^I} \qquad \text{where} \qquad I_j = \left(\sum_{i \in j} \left(Q_i^I \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \qquad \text{and} \qquad \sum_{j=1}^{J} \theta_j^I = 1. \tag{7}$$

Q_i^I denotes the amount of goods produced by firm *i* and used for investment . Let P_i be the price charged by firm *i*. Investors minimize the cost of a given investment *I* subject to these constraints. Maximizing

$$L = -\sum_{i=1}^{N} P_i Q_i^I + \lambda_I \left(\prod_{j=1}^{J} I_j^{\theta_j^I} - I \right) + \sum_{j=1}^{J} \lambda_j \left(\left(\sum_{i \in j} \left(Q_i^I \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} - I_j \right) \tag{8}$$

with respect to I_j and Q_i^I, we obtain first-order conditions $\lambda_j I_j = \theta_j^I \lambda_I I$ and $\mathrm{Q}_i^I = \left(P_i / \lambda_j \right)^{-\eta} I_j$, where λ_I is the relevant price index of aggregate investment, $\lambda_I I$ is total investment expenditures, and λ_j is the relevant price index for goods produced in industry *j*. Applying the same reasoning to aggregate consumption, we obtain the total demand for goods produced by firm *i* in industry *j*:

$$\hat{D}_i = Q_i^I + Q_i^C = \frac{\theta_j^I \lambda_I I}{\lambda_j} \left(\frac{P_i}{\lambda_j} \right)^{-\eta} + \frac{\theta_j^C \lambda_C C}{\lambda_j} \left(\frac{P_i}{\lambda_j} \right)^{-\eta} = \left(\theta_j^I \lambda_I I + \theta_j^C \lambda_C C \right) P_i^{-\eta} \lambda_j^{\eta-1}. \tag{9}$$

Taking logs on both sides and linearizing with respect to $\ln C,$ $\ln I,$ and $\ln \lambda_j$ we obtain

$$\begin{aligned} \ln \hat{D}_i &= \ln\left(\theta_j^I \lambda_I e^{\ln I} + \theta_j^C \lambda_C e^{\ln C} \right) - \eta \ln P_i + (\eta - 1) \ln \lambda_j \\ &\approx \ln \bar{D}_i + \frac{\theta_j^I \bar{\lambda}_I \bar{I} \left(\ln \mathrm{I} - \ln \bar{\mathrm{I}} \right) + \theta_j^C \bar{\lambda}_C \bar{C} \left(\ln C - \ln \bar{C} \right)}{\theta_j^I \bar{\lambda}_I \bar{I} + \theta_j^C \bar{\lambda}_C \bar{C}} - \eta \left(\ln P_i - \ln \bar{P}_i \right) + (\eta - 1)\left(\ln \lambda_j - \ln \bar{\lambda}_j \right) \qquad (10) \\ &= \ln \bar{D}_i + \frac{\bar{I}_j}{\bar{I}_j + \bar{C}_j} \left(\ln \mathrm{I} - \ln \bar{\mathrm{I}} \right) + \frac{\bar{C}_j}{\bar{I}_j + \bar{C}_j} \left(\ln C - \ln \bar{C} \right) - \eta \left(\ln P_i - \ln \bar{P}_i \right) + (\eta - 1)\left(\ln \lambda_j - \ln \bar{\lambda}_j \right) \end{aligned}$$

where we used the fact that $\theta_j^I \lambda_I I = \lambda_j I_j$ and $\theta_j^C \lambda_I C = \lambda_j C_j$ and where bars denote steady-state values. We see that the weights are the steady-state shares of production in industry *j* that are used for investment and consumption. The same logic can be applied to government expenditures and sales in different countries. Based on this reasoning, we construct a variable that shifts demand for products produced by firm *i* as

$$\ln D_{i,t} = (1-\delta_i)\left[\phi_j^C \ln C_t + \phi_j^G \ln G_t + \phi_j^I \ln I_t + \left(1-\phi_j^C-\phi_j^G-\phi_j^I\right)\ln EX_t\right] + \delta_i\left(\sum_m \omega_{j,m} \ln Y_{j,m,t}^F\right).^{14} \quad (11)$$

The subscript *i* denotes the firm, *j* denotes the industry, *t* denotes the year, and *m* denotes the country. The weight δ_i is the mean export share for the firm over the sample period, and the weights ϕ_j^C, ϕ_j^G and ϕ_j^I are industry-specific shares calculated on the two-digit level (SNI92) using input-output tables from Statistics Sweden for 2005. The weights are kept fixed over time to make the demand variable as exogenous as possible. ϕ_j^C is the private domestic consumption of production in industry *j* as a share of final demand excluding direct exports, and ϕ_j^G and ϕ_j^I are the corresponding shares for public consumption and investment. The remaining share, $1-\phi_j^C-\phi_j^G-\phi_j^I$, is the share of final demand excluding direct exports going *indirectly* to exports for the relevant industry, that is, the share used as intermediate inputs into domestic products that are eventually exported. C_t is real private consumption, G_t is real public consumption, I_t is real gross fixed investment, and EX_t is real exports; all are aggregate values in fixed prices from Statistics Sweden's table for the gross national product from the user side.

The weight $\omega_{j,m}$ is the share of industry *j's* direct exports that goes to country *m*. The countries included to represent the export market are Sweden's main trading partners: Germany, Norway, the United Kingdom, Denmark, Finland, the USA, France, the Netherlands, Belgium, Italy, and Spain. The variable $Y_{j,m,t}^F$ is real value added for industry *j* in country *m* from the OECD STAN database. It is meant to capture the demand in country *m* for goods produced by industry *j*.

In order to not introduce spurious correlations due to industry-specific shocks, we do not use industry prices (λ_j) to construct the demand index. To the extent that industry prices respond to industry demand, we can view the effect on an individual firm of an industry-specific demand shift as the combined effect of an exogenous shift in industry demand *and* the industry price response. Both factors should increase the demand for goods produced by an individual firm.

---

[14] In the theoretical model below, *potential sales* by firm *i* are determined by $\hat{D}_{i,t} = \Phi D_{i,t}^{\Sigma} P_{i,t}^{-\eta}$, where D_{it} is the firm-specific demand shifter and P_{it} is the price set by the firm.

Figure 1. Log changes of firm-level variables, industry averages for the four industries with the largest number of firms, baseline panel

20 Wood and products of wood (54 firms)

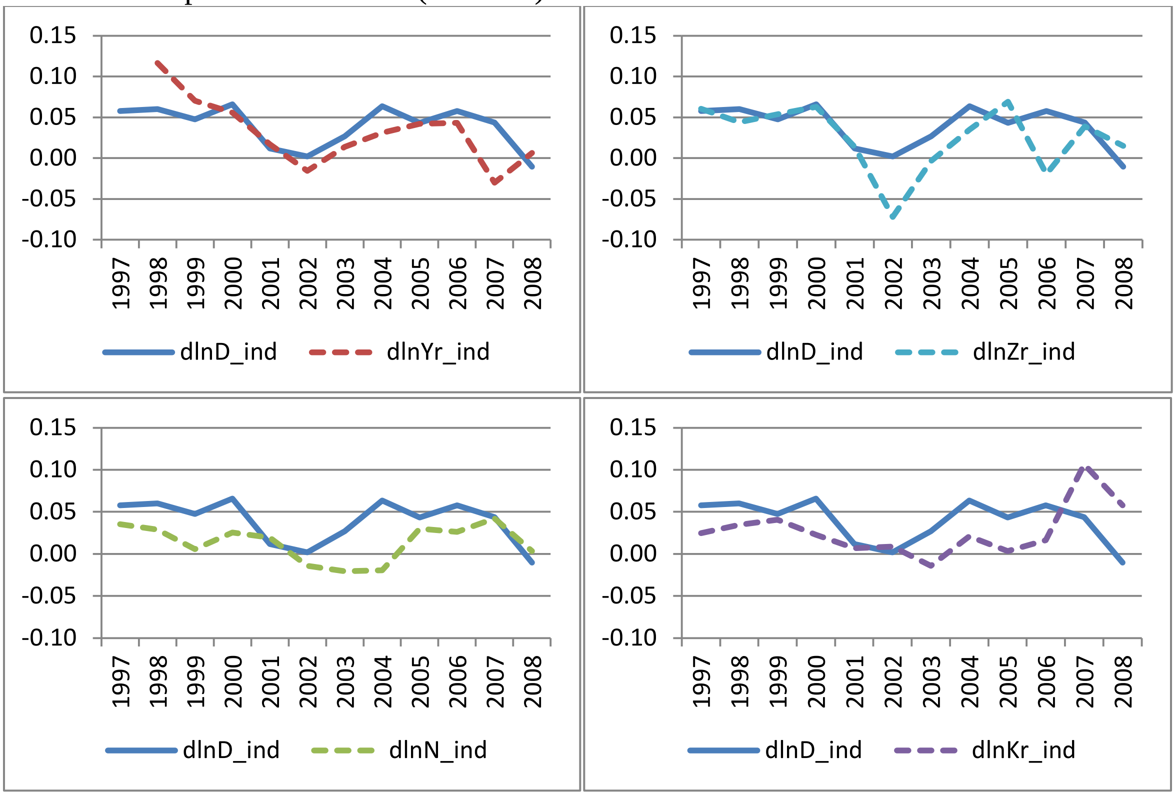

25 Rubber and plastics (65 firms)

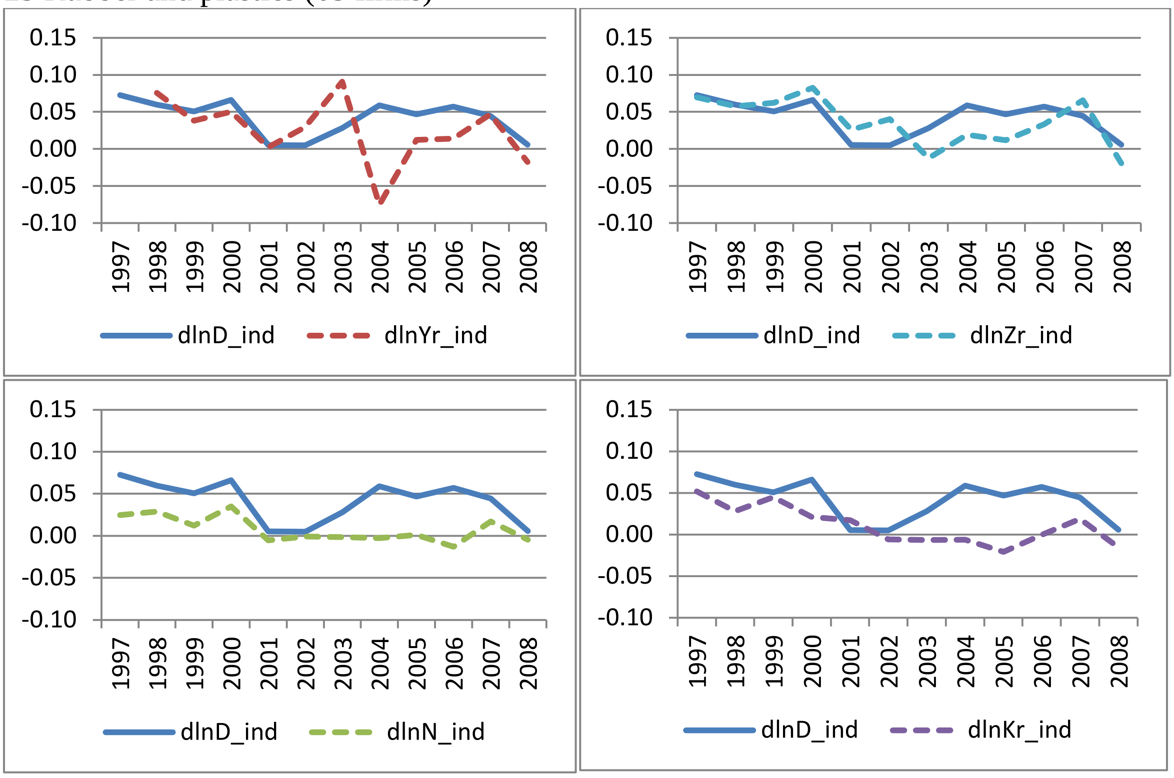

28 Fabricated metal products except machinery and equipment (189 firms)

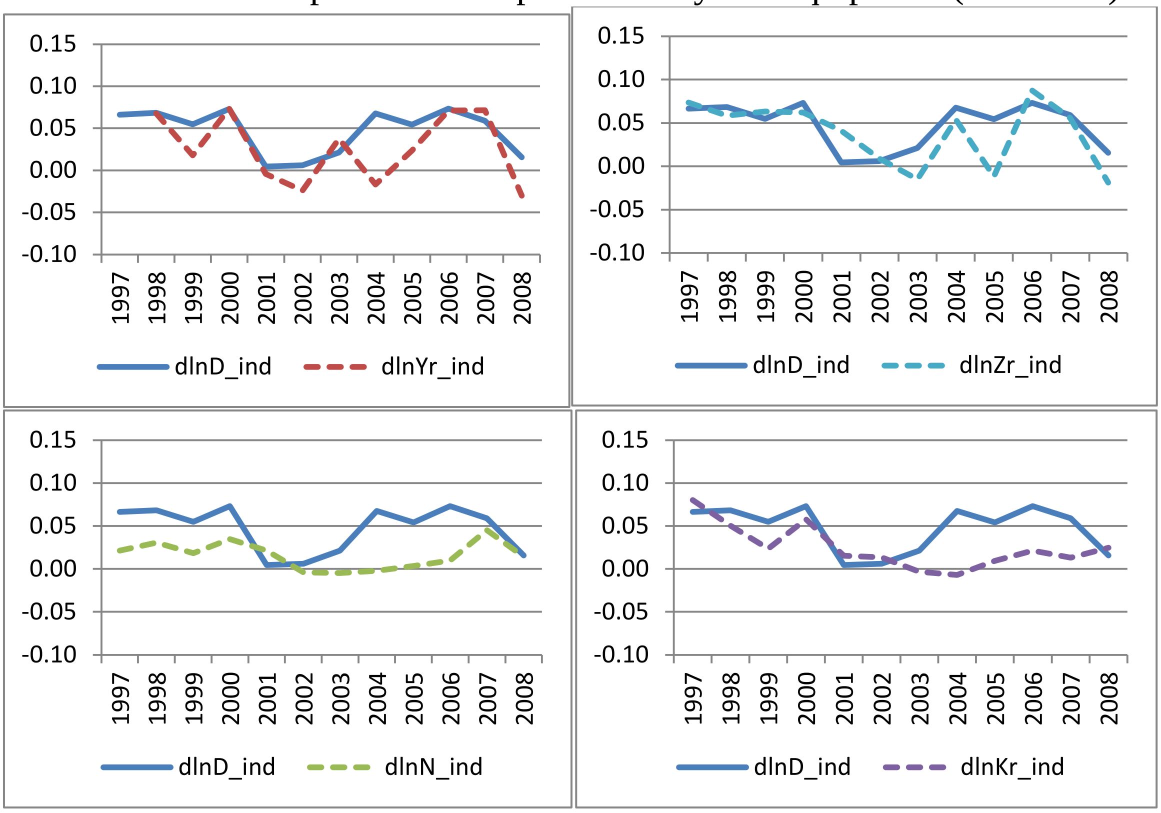

29 Machinery and equipment n.e.c. (168 firms)

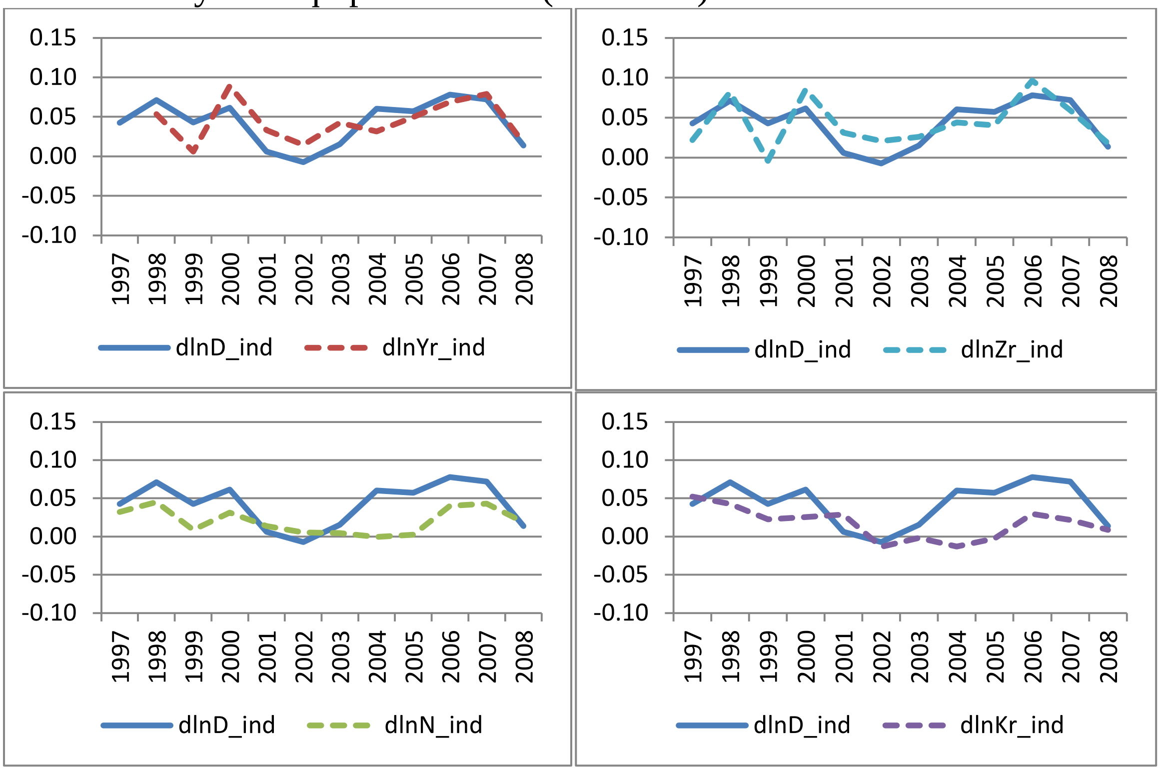

Note: Two digit industries (SNI92). Baseline panel: private firms with at least 10 employees and no extreme observations or missing values. *Yr* is real production, *Zr* is real inventory holdings, *N* is full-time equivalent employees, and *Kr* is the real capital stock.

What type of variation does the firm-specific demand variable represent? Since we include time dummies in our estimated model, these will mop up fluctuations in demand that are common to all firms. Still, there will be a business cycle element in our demand variable because investment varies more than consumption over the cycle. Additionally, our demand variable will reflect structural changes in the composition of domestic demand and differences in economic developments between Sweden and foreign markets, which will affect firms differently depending on their presence in the different markets.

For a detailed list of the number of firms per industry, see the *Table A1* in the Appendix.[15] *Figure 1* illustrates the data for the four industries with the largest number of firms in the baseline panel. For each year, we have calculated log changes of the firm-level variables and taken the average for the industry. In most cases, production and inventories co-vary strongly with the demand index with little or no lag. The changes in the capital stock and the number of workers also co-vary with demand but with a substantial lag.

3.4 Empirical model and identification

To capture the effect of demand shocks on real production (*Yr*), the real capital stock (*Kr*), employment (*N*) and real inventory holdings (*Zr*), we estimate a reduced-form model of the firm with two lags of the endogenous variables and the firm-specific product demand index as an exogenous variable:

$$\begin{aligned}\ln Yr_{i,t} = {} & \beta_1^{Y}\ln Yr_{i,t-1} + \beta_2^{Y}\ln Yr_{i,t-2} + \beta_3^{Y}\ln N_{i,t-1} + \beta_4^{Y}\ln N_{i,t-2} + \beta_5^{Y}\ln Kr_{i,t-1} + \beta_6^{Y}\ln Kr_{i,t-2} \\ & + \beta_7^{Y}\ln Zr_{i,t-1} + \beta_8^{Y}\ln Zr_{i,t-2} + \beta_9^{Y}\ln D_{i,t} + \beta_{10}^{Y}\ln D_{i,t-1} + \varepsilon_{i,t}^{Y}\end{aligned} \tag{12}$$

$$\begin{aligned}\ln N_{i,t} = {} & \beta_1^{N}\ln Yr_{i,t-1} + \beta_2^{N}\ln Yr_{i,t-2} + \beta_3^{N}\ln N_{i,t-1} + \beta_4^{N}\ln N_{i,t-2} + \beta_5^{N}\ln Kr_{i,t-1} + \beta_6^{N}\ln Kr_{i,t-2} \\ & + \beta_7^{N}\ln Zr_{i,t-1} + \beta_8^{N}\ln Zr_{i,t-2} + \beta_9^{N}\ln D_{i,t} + \beta_{10}^{N}\ln D_{i,t-1} + \varepsilon_{i,t}^{N}\end{aligned} \tag{13}$$

$$\begin{aligned}\ln Kr_{i,t} = {} & \beta_1^{K}\ln Yr_{i,t-1} + \beta_2^{K}\ln Yr_{i,t-2} + \beta_3^{K}\ln N_{i,t-1} + \beta_4^{K}\ln N_{i,t-2} + \beta_5^{K}\ln Kr_{i,t-1} + \beta_6^{K}\ln Kr_{i,t-2} \\ & + \beta_7^{K}\ln Zr_{i,t-1} + \beta_8^{K}\ln Zr_{i,t-2} + \beta_9^{K}\ln D_{i,t} + \beta_{10}^{K}\ln D_{i,t-1} + \varepsilon_{i,t}^{K}\end{aligned} \tag{14}$$

[15] More than 80% of the firms are in the same industry throughout the sample period. A firm that changes its industry is assigned to the industry to which it belonged for the longest period of time. Typically, a firm does not change its production entirely but simply passes a threshold in the composition of goods that leads to a change in industry classification.

$$\begin{aligned}\ln Zr_{i,t} &= \beta_1^Z \ln Yr_{i,t-1} + \beta_2^Z \ln Yr_{i,t-2} + \beta_3^Z \ln N_{i,t-1} + \beta_4^Z \ln N_{i,t-2} + \beta_5^Z \ln Kr_{i,t-1} + \beta_6^Z \ln Kr_{i,t-2} \\ &\quad + \beta_7^Z \ln Zr_{i,t-1} + \beta_8^Z \ln Zr_{i,t-2} + \beta_9^Z \ln D_{i,t} + \beta_{10}^Z \ln D_{i,t-1} + \varepsilon_{i,t}^Z\end{aligned} \tag{15}$$

$$\ln D_{i,t} = \rho_1 \ln D_{i,t-1} + \rho_2 \ln D_{i,t-2} + \varepsilon_{i,t}^D \tag{16}$$

We estimate these equations using OLS with fixed effects for each firm (FAD number), and we include time dummies to control for common unobserved macro shocks and trends.

As explained in Section 2, we consider the shocks in the first four equations as technology- and cost shocks plus other shocks that we cannot measure. The key identifying assumption is that these shocks are uncorrelated with the demand variable after we have eliminated shocks that are common to all firms (e.g. common TFP-shocks) by including time dummies. To see when this can be problematic, consider two industries that both sell in the home market, where the first industry produces investment goods while the second industry produces consumption goods. Comparing firms in these industries, we identify effects of demand shocks from the fact that an increase in investment raises demand only for firms in the first industry. However, suppose there is a technology or cost shock that affects a large fraction of the firms that produce investment goods. Such a shock will reduce the price of investment goods, which will, most likely, lead to an increase in aggregate investment.[16] This shock will affect firms that produce investment goods directly as well as the demand variable as we measure it, so there will be a correlation between demand and the shocks in the equations for the endogenous variables, leading to biased estimates. To provide some idea whether shocks of this type were important in this period, we plot the ratio of investment to consumption together with the ratio of the corresponding deflators in *Figure A2* in the Appendix. While we observe a clear cyclical pattern in investment relative to consumption, the relative price varies much less, and the two variables are slightly positively correlated. This means that if anything, we underestimate the effects of demand shocks on production.[17]

[16] For a macroeconomic analysis of the effects of investment-specific technology shocks, see Greenwood, Hercowitz and Krusell (2000).

[17] A similar argument could be made with respect to domestic and foreign demand. A productivity or cost shock that affects firms which produce for the domestic market, but leaves exporters unaffected, may lead to a change in domestic demand relative to foreign demand. However, it is hard to imagine a large shock that would affect firms differentially in this way.

3.5 Impulse-response functions

Figure 2 shows how firms on average respond to an exogenous demand shock.[18] The effect of the shock on demand itself is slightly hump-shaped and quite persistent; the half-time is 9 years. Since we have time dummies in the model, we are not capturing the general business cycle but rather shocks that are more persistent. Production responds immediately to a change in demand and more than the demand shock itself, which may be because the demand for goods produced in manufacturing is more sensitive to shocks and hence more volatile than aggregate consumption and investment, which enter our constructed demand index.[19]

Figure 2. Effects of a firm-specific demand shock in empirical model

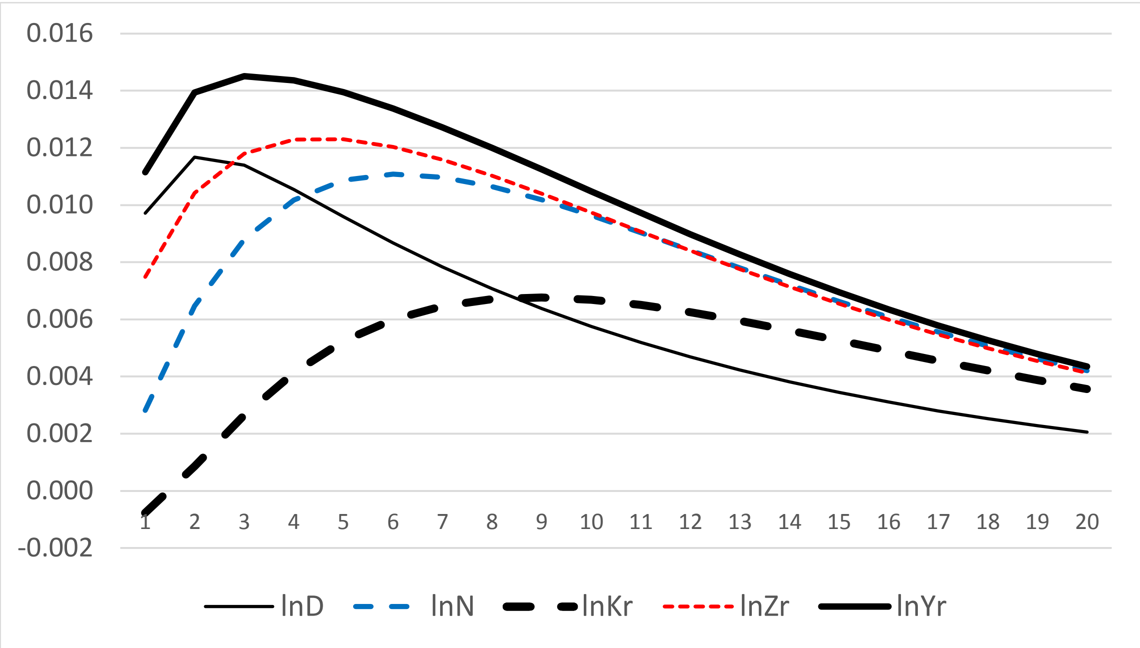

Note: Impulse-response functions from reduced-form model estimated on a balanced panel consisting of private manufacturing firms with at least ten employees, no extreme observations, and no missing values 1997-2008 (818 firms). Firm and time fixed effects are included. The variables are in logs and the time units on the horizontal axis are years. With two lags, the number of observations included in the estimations is 8180. Corresponding figures for alternative samples and specifications can be found in the Appendix.

Production peaks in period 3, one year after the peak in demand. The inventory stock responds positively and peaks two years after the peak in production. Employment reaches its peak 3 years after the peak in production, while the capital stock peaks 6 years after the peak in production. Note that employment almost catches up with production, but the response of the capital stock is much weaker. The first-year effect on the number of workers is 25 percent

[18] For a detailed presentation of the regression results, see *Table A2* in the Appendix. Confidence intervals for the impulse-response functions are shown in *Figure 5*.
[19] See further discussion in Section 6.

of the effect on output, and the capital stock does not respond at all in the first year. Thus, we see that firms are able to satisfy demand with a relatively small increase in registered inputs, which implies very strong responses of labor and factor productivity as they are commonly measured.

In the Appendix, we analyze the robustness of these estimation results with respect to changes in the sample and the specification. We conclude that the qualitative results are reasonably robust. For all specifications, production, inventories, labor and the capital stock respond positively to the demand shock and in most cases, the order of the response is the same: production responds quickly, followed by inventories and labor, while the capital stock responds with the longest lag. *Table A3* in the Appendix summarizes the robustness checks.[20]

4. Theory

There appear to be some adjustment lags or costs that slow down the adjustment of capital and labor input, but production and inventories respond quickly to the demand shock. Below, we present some features of our theoretical model that can potentially explain these empirical responses. First, we discuss adjustment costs, implementation lags, increasing returns to scale and factor utilization. Then, we specify the relation between output and value added, price rigidity and our model of inventory holdings. Finally, the firm's maximization problem is presented in Section 4.7.

4.1 Adjustment costs and implementation lags

We assume that firms can hire workers at a given wage and buy capital at a given price, but subject to adjustment costs. We include *quadratic adjustment costs* for labor and capital as a simple representation of various types of adjustment costs that affect the average response of firms. The adjustment costs are equal to $c_N\left(H_t - \delta_n \bar{N}\right)^2/2 + c_K\left(I_t - \delta_k \bar{K}\right)^2/2$, where δ_k is the rate at which capital depreciates, and δ_n is an exogenous separation rate for labor. $\bar{N}$ and $\bar{\mathrm{K}}$ denote the steady-state levels of N_t and K_t, so there are quadratic costs for hiring

[20] To check whether demand is exogenous, we ran a regression with the demand variable as dependent variable and two lags of all variables on the right hand side. The coefficients for lagged production turned out to be barely statistically significant, but all coefficients (except those for lags of demand itself) were found to be lower than 0.003.

and investing more than the steady-state levels $\delta_n \bar{N}$ and $\delta_k \bar{K}$.[21] These costs take the form of reduced production due to disruptions in the production process.

We also include an *implementation lag* in investment by assuming that some given fractions of the investment that is *decided* in year t are *implemented* in years *t*, *t*+1 and *t*+2:

$$K_t = (1-\delta_k) K_{t-1} + \lambda_{k1} I_t + \lambda_{k2} I_{t-1} + (1-\lambda_{k1}-\lambda_{k2}) I_{t-2}. \tag{17}$$

K_t is the capital stock at the end of the period, and I_t is the investments *decided* in period t. This approach is similar to "time to build" (Kydland and Prescott 1982) and it is more flexible than assuming either no lag or a one-period implementation lag, as in Burnside-Eichenbaum-Rebelo (1993). Similarly, we assume that hiring is implemented in the current and the coming year:

$$N_t = (1-\delta_n) N_{t-1} + \lambda_n H_t + (1-\lambda_n) H_{t-1}. \tag{18}$$

N_t is employment during period t, and H_t is the hiring *decided* in period *t*.

4.2 Increasing returns to scale in production

Our estimated impulse-response functions show that firms can increase production in the short run with much smaller percentage increases in the registered inputs of capital and labor. One possible explanation is that there are *increasing returns to scale* so that changes in inputs lead to proportionally larger changes in production (Hall 1988). To model this, we assume that the capital stock consists of two components. First, there is a flexible part $\hat{K}_t$ that enters a CES production function with constant returns to scale, and second, there is a fixed amount of capital F_k that the firm must have in order to produce at all. Thus, the total observed capital is given by $K_t = F_k + \hat{K}_t$. Similarly, we distinguish between fixed and flexible employment: $N_t = F_n + \hat{N}_t$.

4.3 Factor utilization and organizational capital

Looking at the dynamic response in *Figure 2*, we see that production increases much more than observed inputs in the first year, but after a few years, employment has almost caught up with production. It is unlikely that increasing returns can explain the whole picture; firms appear to have some form of excess capacity that they can use to meet demand. A standard

[21] The specification implies that adjustment costs are zero in the steady state; this helps to solve analytically for the steady state. This is necessary for the estimation.

way to model this is to allow for variable *utilization* of the factors of production (Burnside, Eichenbaum, Rebelo, 1993; Sbordone 1996, 1997).

The key question, then, is why the firm did not make full use of its resources for production before the shock occurred. As Bean (1990) notes, hiring and firing costs cannot explain why firms operate *within* their production frontier during recessions. This suggests that there must be some cost of increasing resource utilization, or else the firm would always make full use of its resources in production.[22] We allow for variations in the utilization (u_t) of both factors of production at a cost given by $\Phi_u\left(u_t-1+\left(c_u/2\right)\left(u_t-1\right)^2\right)\hat{N}_t$. The variable u_t enters multiplicatively in the production function below; it can represent effort or overtime, either of which increases the use of both labor and capital. In the latter interpretation, the convex cost may reflect an overtime premium that may be part of an explicit or implicit contract.[23]

Several authors have noted that workers spend substantial amounts of time on activities that do not contribute to *current* production but that increase *future* production (see e.g. Fay and Medoff, 1985; Bean, 1990; Kim and Lee, 2007). There are many such activities we can think of, including cleaning and maintenance, reorganizing, and training. To capture this, we include another element in the model that we call *organizational capital*. We assume that the firm has a stock of organizational capital Ω, and the larger this stock is, the more it can produce with given inputs. Workers spend a share x_t of their time on activities that increase current production and a fraction $1-x_t$ of their time accumulating organizational capital that increases future production. Thus, we write the production function for value added

$$F\left(u_t,\Omega_{t-1},\hat{K}_{t-1},x_t,\hat{N}_t,H_t,I_t\right)=Au_t\left(a-\frac{a-1}{\Omega_{t-1}^{\xi}}\right)\left(\alpha\hat{K}_{t-1}^{\frac{\sigma-1}{\sigma}}+\left(1-\alpha\right)\left(x_t\hat{N}_t\right)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \qquad (19)$$
$$-\frac{c_N}{2}\left(H_t-\delta_n\bar{N}\right)^2-\frac{c_K}{2}\left(I_t-\delta_k\bar{K}\right)^2$$

where $a>1$ and where $\hat{K}_t$ and Ω_t denote the stocks of flexible capital and organizational capital at the end of period *t*. *Figure 3* illustrates the function $a-\left(a-1\right)/\Omega_{t-1}^{\xi}$ for different values of the parameters *a* and ξ. We normalize so that $\Omega=1$ in the steady state, and thus the

[22] Alternatively, the markup is very large (Hall, 1988) or prices are very sticky (Rotemberg and Summers, 1990) – see the discussion in Section 7.
[23] As mentioned in Section 3, the number of workers is measured as "full-time equivalent," and this measure does not take into account variations in registered overtime, so variations in u_t will partly reflect variations in overtime.

function value is one in the steady state. As organizational capital increases, the function value increases asymptotically toward a, and if organizational capital falls to $(a-1)/a$, the function value falls to zero. Roughly speaking, the parameter *a* determines the slope of the function, while ξ determines its concavity.

We assume that the accumulation of organizational capital is governed by

$$\Omega_t = (1-\delta_\omega)\Omega_{t-1} + \chi(1-x_t) \tag{20}$$

where δ_ω is the depreciation rate of organizational capital, and $1-x_t$ is the fraction of time spent accumulating organizational capital. The parameter χ is set to be consistent with the normalization $\Omega = 1$ in the steady state. The basic idea behind this specification is that when there is temporarily high demand, the firm will tell workers to increase the fraction of their time spent on current production and to spend less time organizing and training.

Figure 3. Productivity contribution of organizational capital

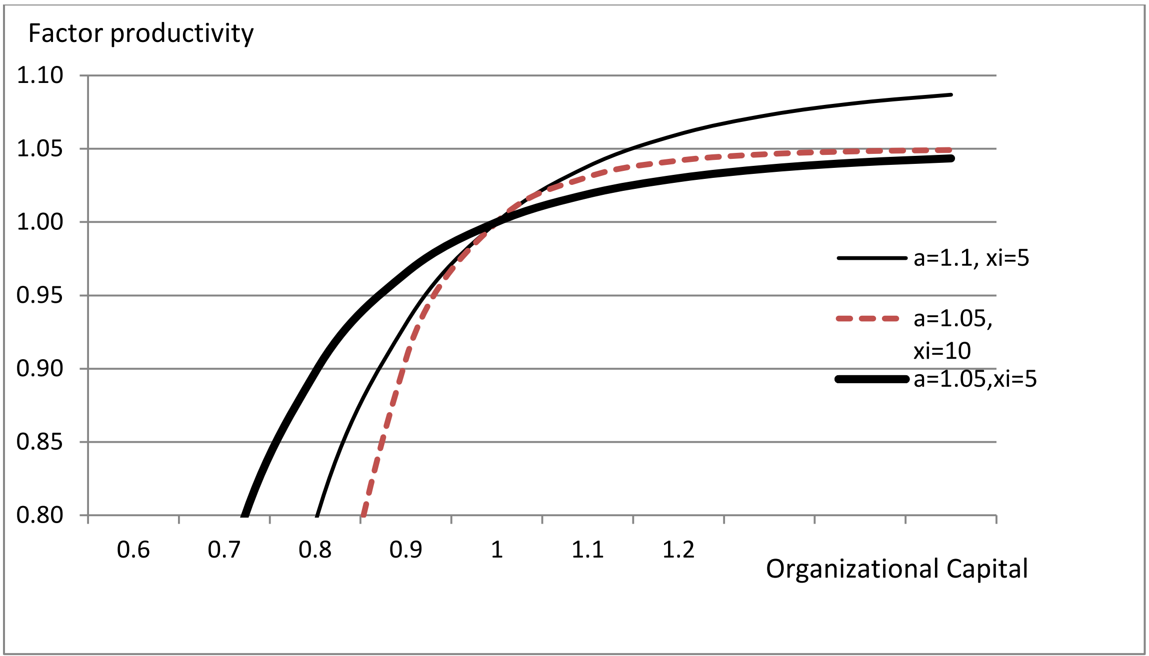

Note: The function shows the contribution of organizational capital to factor productivity $\left(a-(a-1)/\Omega_{t-1}^{\xi}\right)$ as a function of Ω for different values of the parameters *a* and ξ.

4.4 Output and value added

We assume that value added and materials inputs are combined in a Leontief production function

$$Y_t = \min\left\{F\left(u_t, \Omega_{t-1}, \hat{K}_{t-1}, x_t, \hat{N}_t, H_t, I_t\right), M_t / \mathrm{m}\right\} \tag{21}$$

where Y_t is the quantity produced, $F(\cdot)$ is as defined above, and M_t is the quantity of intermediate inputs used. Cost minimization then implies that $Y_t = F(\cdot) = M_t / m$. Normalizing the price of intermediate inputs to one, the cost of inputs is mY_t: a fixed amount of cloth is needed to make a shirt.

Since there is no substitutability between value added and materials, materials inputs and total output will always be proportional to value added, and it should not matter whether we measure production by output or value added. Clearly, one could allow for some substitutability between materials and other inputs, but as seen in the Appendix, the impulse-response functions are fairly similar when we use value added to measure production. Note that we allow for increasing returns to scale by including fixed costs in terms of capital and labor, but not in terms of materials.[24]

4.5 Price rigidity

Another factor that can prevent firms from always optimally utilizing their resources is *price rigidity*. If demand falls and the firm cannot (or does not want to) reduce its prices, quasi-fixed resources will become less utilized (Rotemberg and Summers, 1990). We incorporate price rigidity in a simple way by including a quadratic adjustment cost for prices: $\theta\left(P_\tau / P_{\tau-1} - 1\right)^2 / 2$. We assume that firms always satisfy demand, which makes sense if firms have sufficient market power.

4.6 Inventory model

The estimated responses show that firms increase their inventory stocks when demand increases. This result is the opposite of the production-smoothing idea that by drawing down inventories in periods of high demand, firms can stabilize production. To explain the observed pattern, we follow the ideas of Kahn (1987, 1992) and Bils and Kahn (2000) and assume that inventories of finished goods are needed in order to sell the good. Below, we present a very

[24] See section VI in Basu (1996) and Basu-Fernald (1997) page 255 for discussions of these issues.

stylized model that provides a reasonable functional form that we can include in our estimated model.

Consider a firm that sells goods, e.g., steel bars, that come in *M* different varieties, which we will call sizes. Let us assume that a customer will only buy a steel bar if he/she finds the right size. The firm has a sales department and a production department, and the sales department sends an order to the production department *T* times per year to replenish the inventory stock. For concreteness, we can think of the case when *T*=12, so inventories are replenished every month.[25]

Let $\hat{D}$ be the *potential sales* of all varieties during a year (we omit the time index here). $\hat{D}$ is what the firm would sell if it would never stock out, and we assume that $\hat{D} = \Phi D^{\Sigma} P^{-\eta}$, where Φ is a constant, *D* is a demand shifter, and *P* is the price set by the firm. To make the model as simple as possible, we assume that D is known, that *P* is set at the beginning of the year, and that both are constant over the year. Demand in a particular month for a particular size is assumed to be $\lambda \hat{D}$, where λ is a stochastic variable that is uniformly distributed between λ_1 and λ_2. The supports of the distribution are given by $\lambda_1 = (1-\Psi)/(TM)$ and $\lambda_2 = (1+\Psi)/(TM)$, where *T* is the number of inventory periods (months), and *M* is the number of sizes. The parameter Ψ has a value between zero and unity; it reflects the degree of uncertainty about the demand for individual varieties.

Since demand is assumed to be symmetrically and independently distributed across sizes, the sales department will ensure that they stock up with the same quantity of each size whenever they replenish inventories. Let *z* be the inventory stock of a specific size held at the beginning of a month. It follows immediately that $\lambda_1 \hat{D} < z < \lambda_2 \hat{D}$. With a smaller inventory stock, the firm would always stock out, and there is no reason to hold a larger inventory stock than the maximum possible sales of a particular size. If the realization of λ is such that $\lambda \hat{D} \leq z$, sales of that specific size will be $\lambda \hat{D}$, and if $\lambda \hat{D} > z$, sales of that specific size will be *z*. Letting $\hat{\lambda}$ be the critical value of λ where the firm runs out of stock $\left(\hat{\lambda}\hat{D} = z\right)$, we obtain the expected sales of a particular size in a given month as

[25] Contrary to the Ss model, *T* is taken as exogenous, so our model differs from the Ss model in the same way as the Taylor model of wage/price setting differs from state-contingent pricing. This type of model is called "periodic review" in the inventory literature; see Urban (2004).

$$s=\int_{\lambda_1}^{\hat{\lambda}}\lambda\hat{D}\frac{TM}{2\Psi}d\lambda+\int_{\hat{\lambda}}^{\lambda_2}z\frac{TM}{2\Psi}d\lambda=\frac{TM}{2\Psi}\left(\frac{\hat{\lambda}^2\hat{D}}{2}-\frac{\lambda_1^2\hat{D}}{2}+\lambda_2 z-\hat{\lambda}z\right)=\frac{1+\Psi}{2\Psi}z-\frac{(1-\Psi)^2}{4\Psi TM}\hat{D}-\frac{TM}{4\Psi}\frac{z^2}{\hat{D}} \tag{22}$$

We assume that $\hat{D}$ is observed at the beginning of the year and is constant throughout the year, so the total expected sales of storable finished goods during the year are:

$$TMs=\frac{1+\Psi}{2\Psi}T\hat{Z}-\frac{(1-\Psi)^2}{4\Psi}\hat{D}-\frac{T^2\hat{Z}^2}{4\Psi\hat{D}} \tag{23}$$

where $\hat{Z}$ is the total stock of finished goods at the beginning of the month: $\hat{Z}=Mz$. To keep the model simple, we assume that the firm sells a large number of varieties, so we can view this function as a deterministic function that determines sales. This function has several natural properties:

i) For a given inventory stock, the maximum sales are $T\hat{Z}$, and sales approach that limit as $\hat{D}$ goes to $\hat{Z}T/(1-\Psi)$; for a lower level of demand, there will be some varieties that will not sell out.

ii) For a given level of demand, the inventory stock that maximizes sales is $\hat{Z}=(1+\Phi)\hat{D}/T$. To ensure that the required sizes are always available, the firm needs finished goods inventories of each variety that correspond to the maximum possible demand during an inventory period. (In the model, the costs of financing, depreciation and storage make the optimal inventories smaller than this amount.)

iii) Starting from a low *T*, a higher *T* will increase sales.

Figure 4 shows the sales of finished storable goods (TMs) as a function of the inventory stock, when potential sales per year are 1200, *T*=12 so the stock is replenished every month, and $\Psi=1$ so that the demand for a particular size is uniformly distributed between zero and twice the expected demand for that size. In this case, sales are equal to $12\cdot\hat{Z}-0.03\cdot\hat{Z}^2$ in the relevant interval. To never stock out, the firm needs to have an inventory stock that is 200, twice as large as the potential monthly sales. With a smaller stock, it will sell less because some sizes will run out. If the stock is replenished more seldom, this will reduce the sales for a given stock.[26]

[26] We do not explicitly address adjustments during the year, but the following very stylized timing assumptions can serve to motivate the specification: i) Inventory stocks are replenished at the beginning of the month, and customers buy the good at the end of the month. ii) Firms learn the level of demand for the coming year at the start of the year. iii) Workers dislike putting in more effort during a year, but they are indifferent to how effort is allocated during the year. This means that firms adjust the level of inventories at the beginning of the year.

Figure 4. Sales of storable goods as a function of the inventory stock

$\left(\Psi = 1, \hat{D} = 1200\right)$

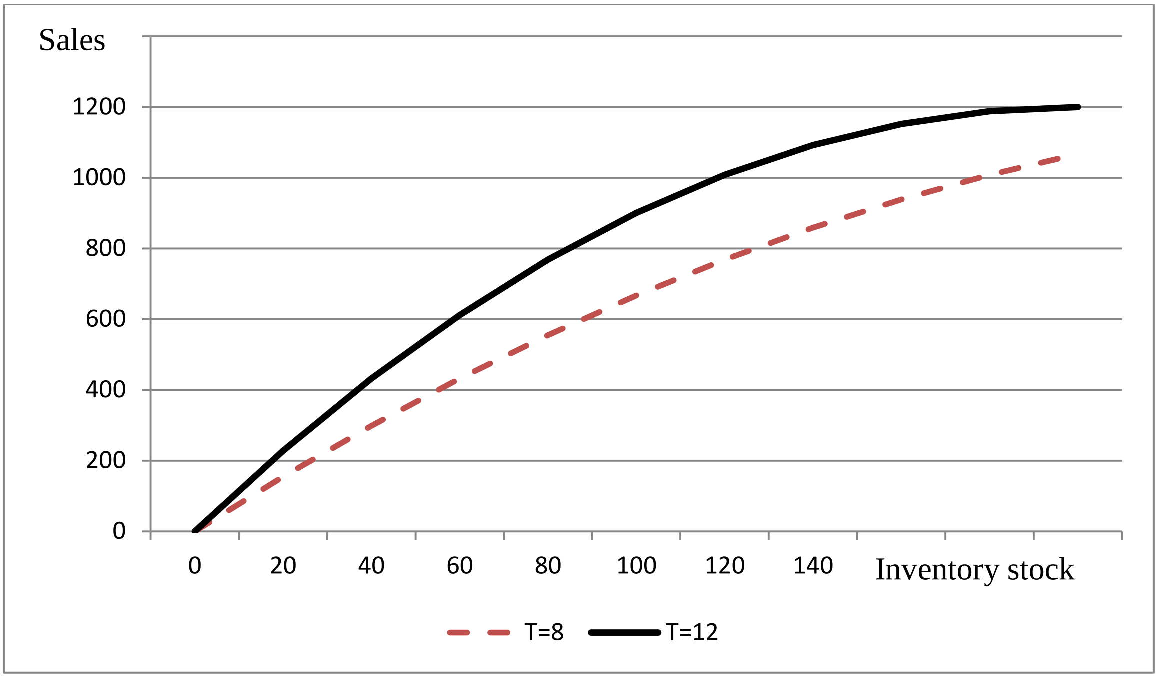

Note: *T* denotes the number of times the firm replenishes inventories per year.

To this we add yet another modification by assuming that there are some goods that are sold without holding stock. These may be perishable goods or goods produced on order. Sales of these goods are simply assumed to be equal to $\hat{D}$. Letting the fraction of storable finished goods be Λ, we obtain total sales as

$$S = \Lambda TMs + \left(1-\Lambda\right)\hat{D} = \kappa_1 \hat{Z} + \kappa_2 D^{\Sigma} P^{-\eta} - \kappa_3 \hat{Z}^2 D^{-\Sigma} P^{\eta} \tag{24}$$

where $\kappa_1 = \Lambda\left(1+\Psi\right)\mathrm{T}/\left(2\Psi\right)$, $\kappa_2 = 1-\Lambda-\Lambda\left(1-\Psi\right)^2 \Phi/\left(4\Psi\right)$ and $\kappa_3 = \Lambda T^2/\left(4\Psi\Phi\right)$.

The accumulation of finished goods inventories is governed by the function

$$\hat{Z}_t = \left(1-\delta_z\right)\hat{Z}_{t-1} + Y_t - S_t \tag{25}$$

where $\hat{Z}_t$ is the finished goods inventory stock at the end of the year, and δ_z is the rate at which inventories depreciate during the year. We also include a cost $c_z \cdot \hat{Z}_t$ that reflects other costs of holding inventories, such as the cost of providing storage space and managing the inventories.

Finally, we note that inventories consist not only of finished goods but also of inputs and goods in process, but our data do not allow us to distinguish between different types of inventories. To take this into account, we simply assume that the firm holds a stock of intermediate inputs that is proportional to current production: $h_z Y_t$.[27] We assume that these inputs can be bought without delay; hence, the total observed inventory stock Z_t is given by $Z_t = \hat{Z}_t + h_z Y_t$.

4.7 Profit maximization

The firm's profit-maximization problem is to choose $S_t, Y_t, P_t, \hat{K}_t, I_t, \hat{N}_t, H_t, u_t, x_t, \Omega_t, \hat{Z}_t$ to maximize

$$E_t \left\{ \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left[\begin{array}{l} S_\tau P_\tau - mY - W \hat{N}_t - \Phi_u \left(u_t - 1 + \frac{c_u}{2} \left(u_t - 1 \right)^2 \right) \hat{N}_\tau \\ -P^K \left(\hat{K}_\tau - (1-\delta_k) \hat{K}_{\tau-1} \right) - c_z \hat{Z}_t - \frac{\theta}{2} \left(\frac{P_\tau}{P_{\tau-1}} - 1 \right)^2 \end{array} \right] \right\} \tag{26}$$

subject to the following constraints (with associated shadow prices) that hold for all t:

$$Y_t = A u_t \left(a - \frac{a-1}{\Omega_{t-1}^{\xi}} \right) \left(\alpha \hat{K}_{t-1}^{\frac{\sigma-1}{\sigma}} + (1-\alpha) \left(x_t \hat{N}_t \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - \frac{c_N}{2} \left(H_t - \delta_n \bar{N} \right)^2 - \frac{c_K}{2} \left(I_t - \delta_k \bar{K} \right)^2$$

$$\nu_t \tag{27}$$

$$S_t = \kappa_1 \hat{Z}_t + \kappa_2 D_t^{\Sigma} P_t^{-\eta} - \kappa_3 \hat{Z}_t^2 D_t^{-\Sigma} P_t^{\eta} \qquad \mu_t \tag{28}$$

$$\hat{Z}_t = (1-\delta_z) \hat{Z}_{t-1} + Y_t - S_t \qquad g_t \tag{29}$$

$$F_k + \hat{K}_t = (1-\delta_k) \left(F_k + \hat{K}_{t-1} \right) + \lambda_{k1} I_t + \lambda_{k2} I_{t-1} + (1-\lambda_{k1} - \lambda_{k2}) I_{t-2} \qquad q_t \tag{30}$$

$$F_n + \hat{N}_t = (1-\delta_n) \left(F_n + \hat{N}_{t-1} \right) + \lambda_n H_t + (1-\lambda_n) H_{t-1} \qquad \gamma_t \tag{31}$$

$$\Omega_t = (1-\delta_\omega) \Omega_{t-1} + \chi (1 - x_t). \qquad \phi_t \tag{32}$$

The shadow price ν_t is the marginal value (and cost) of a unit of value added, g_t is the marginal value (and cost) of a unit of the final good, μ_t is the value of an additional unit of sales (the markup), and q_t, γ_t and ϕ_t are the shadow price prices associated with capital, labor and organizational capital. Defining $\hat{Y}_t = \left(\alpha \hat{K}_{t-1}^{\frac{\sigma-1}{\sigma}} + (1-\alpha) \left(x_t \hat{N}_t \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$ we can write the first-order conditions:

---

[27] We disregard the stock-flow aspect of inventories of inputs; see Humphreys, Maccini and Schuh (2001) for a more elaborate model.

$$S_t: \quad \mu_t = P_t - g_t \tag{33}$$

$$Y_t: \quad g_t = v_t + m \tag{34}$$

$$P_t: \quad S_t - \mu_t \eta \left(\frac{\kappa_2 D_t^{\Sigma}}{P_t^{1+\eta}} + \frac{\kappa_3 \hat{Z}_t^2}{D_t^{\Sigma} P_t^{1-\eta}} \right) - \theta \left(\frac{P_t}{P_{t-1}} - 1 \right) \frac{1}{P_{t-1}} + \beta \theta E_t \left\{ \left(\frac{P_{t+1}}{P_t} - 1 \right) \frac{P_{t+1}}{P_t^2} \right\} = 0 \tag{35}$$

$$K_t: \quad q_t = -P^K + \beta E_t \left\{ v_{t+1} A u_{t+1} \left(a - \frac{a-1}{\Omega_t^{\xi}} \right) \alpha \left(\frac{\hat{Y}_{t+1}}{\hat{K}_t} \right)^{\frac{1}{\sigma}} + (1-\delta_K)\left(P^K + q_{t+1} \right) \right\} \tag{36}$$

$$I_t: \quad v_t c_k \left(I_t - \delta_k \bar{K} \right) = \lambda_{k1} q_t + \beta \lambda_{k2} E_t \left(q_{t+1} \right) + \beta^2 \left(1 - \lambda_{k1} - \lambda_{k2} \right) E_t \left(q_{t+2} \right) \tag{37}$$

$$\hat{N}_t: \quad \gamma_t = v_t A u_t \left(a - \frac{a-1}{\Omega_{t-1}^{\xi}} \right) (1-\alpha) x_t^{\frac{\sigma-1}{\sigma}} \left(\frac{\hat{Y}_t}{\hat{N}_t} \right)^{\frac{1}{\sigma}} - W_t - \Phi_u \left(u_t - 1 + \frac{c_u}{2} \left(u_t - 1 \right)^2 \right) + \beta \left(1 - \delta_n \right) E_t \left(\gamma_{t+1} \right) \tag{38}$$

$$H_t: \quad v_t c_n \left(H_t - \delta_n \bar{N} \right) = \lambda_n \gamma_t + \beta \left(1 - \lambda_n \right) E_t \left(\gamma_{t+1} \right) \tag{39}$$

$$u_t: \quad \Phi_u \left(1 - c_u + c_u u_t \right) \hat{N}_t = v_t A \left(a - \frac{a-1}{\Omega_{t-1}^{\xi}} \right) \hat{Y}_t \tag{40}$$

$$x_t: \quad \phi_t \chi = v_t A u_t \left(a - \frac{a-1}{\Omega_{t-1}^{\xi}} \right) (1-\alpha) \hat{N}_t^{\frac{\sigma-1}{\sigma}} \left(\frac{\hat{Y}_t}{\mathrm{x}_t} \right)^{\frac{1}{\sigma}} \tag{41}$$

$$Z_t: \quad g_t = \mu_t \left(\kappa_1 - 2\kappa_3 \frac{\hat{Z}_t}{D_t^{\Sigma} P_t^{-\eta}} \right) - c_z + \beta \left(1 - \delta_z \right) E_t \left(g_{t+1} \right) \tag{42}$$

$$\Omega_t: \quad \phi_t = \beta E_t \left\{ v_{t+1} A u_{t+1} \frac{a-1}{\Omega_t^{\xi+1}} \xi \hat{Y}_{t+1} + \left(1 - \delta_{\omega} \right) \phi_{t+1} \right\}. \tag{43}$$

The total amounts of capital, labor and inventory stock are $K_t = F_k + \hat{K}_t$, $N_t = F_n + \hat{N}_t$, and $Z_t = h_z Y_t + \hat{Z}_t$. Finally, we also need to specify the stochastic process for the demand shock. In line with our empirical model, we assume that the demand shock follows an AR(2) process:

$$D_t = 1 + \rho_1 \left(D_{t-1} - 1 \right) + \rho_2 \left(D_{t-2} - 1 \right) + \varepsilon_t \tag{44}$$

To sum up, adjustment costs and implementation lags may help to explain the sluggish adjustment of labor and capital, while increasing returns to scale in production, variable utilization and variations in time spent building organizational capital could potentially explain the observed increase in factor productivity (as it is normally measured) in response to a demand shock. A positive inventory response may arise because of the stock-out motive and because a large fraction of inventories consist of inputs and goods in process. We now turn to the estimation, which will help us discriminate between these alternative explanations.

5. Estimation method

We follow Christiano, Eichenbaum and Evans (2005) and estimate the structural parameters in the theoretical model by finding the set of parameter values that make the impulse responses in the theoretical model match the impulse responses of its empirical counterpart.

5.1 Matching impulse-response functions

The target function is constructed as in Christiano, Eichenbaum and Evans (2005):

$$J = \min\left[\hat{\Psi} - \Psi\left(\vec{\gamma}\right)\right] V^{-1} \left[\hat{\Psi} - \Psi\left(\vec{\gamma}\right)\right]. \tag{45}$$

$\Psi\left(\vec{\gamma}\right)$ contains the impulse responses calculated with the theoretical model for different horizons as a function of the model parameter vector $\vec{\gamma}$, and $\hat{\Psi}$ is the empirical counterpart. V is a diagonal matrix with the variances from the empirical estimation. These variances are related to the 95% confidence intervals, which are shown in *Figure 5* below. We include 20 years of IRFs in the estimation.

5.2 Prior constraints

We constrain the parameters to be in an economically meaningful range; the prior intervals for the estimated parameters are shown in the columns denoted min and max in *Table 2* below. Additionally, we impose some restrictions on the steady-state levels to ensure that they are roughly consistent with what we know about the levels from accounting data.[28] In the baseline sample, the median cost of personnel relative to value added is 78 percent, and we constrain this ratio to be between 73 and 83 percent in the steady state. The median real capital (calculated as described above) relative to value added is 83 percent, and we constrain this ratio to be between 78 and 88 percent in the steady state.

The median ratio of inventories to production is 15 percent, but we should note that $\hat{Z}_t$ in the theoretical model is the stock of finished goods when the firm has just replenished inventories. This means that the steady-state level of inventories should be higher than the number observed in the data. We therefore constrain Z_t / Y_t to be between 15 and 25 percent

[28] A set of parameter estimates that are grossly inconsistent with what we know about the levels would be uninteresting.

in the steady state. We cannot distinguish between different types of inventories in our data, but we know the proportions of finished goods, inputs and goods in process for manufacturing as a whole. If we count goods in process as half finished goods and half inputs, then roughly half the inventory stock consists of finished goods. Therefore, we constrain the ratio of finished goods to stored inputs to be between 2/3 and 3/2. Finally, we check that profits are positive in the steady state.

5.3 Search algorithm

We use the search algorithm from Mickelsson (2016), which is based on the local algorithm of Nelder and Mead (1965). The basic idea behind this algorithm is to start with a large number of starting vectors that are spread out across the parameter space and then combine these vectors in a smart manner to approach the global maximum without getting stuck at local maxima or iterating too long on flat surfaces. Mickelsson (2016) shows that this algorithm does better than most commonly used search algorithms when the objective function has many local minima and flat surfaces.

5.4 Confidence intervals

To obtain confidence intervals for the parameters, we generate distributions of the estimates in the following way:

1. First, we create a new sample of firms of the same size as the original sample by drawing firms randomly from the original sample with replacement.
2. This sample is used to obtain a new estimate of the empirical IRFs.
3. The impulse responses from the empirical model are then used to estimate the parameters of the theoretical model, as described above.
4. The vector of parameter estimates is saved, and steps 1-3 are repeated 1000 times to obtain a distribution of estimates.[29]

6. Results

6.1 Some parameters are poorly identified or have corner solutions

There are many parameters in our theoretical model, and it is not surprising that some are poorly identified with these data. The ratio of inputs to total output (m) was set to 0.6 based

[29] We tested bootstrapping with the restriction that the number of firms from each industry should be the same in each sample, but this made very little difference to the results.

on aggregate data for manufacturing. Attempts to estimate the discount factor, the depreciation rates, the elasticity of substitution and price rigidity indicate that these parameters are poorly identified. For this reason, we set the discount rate (β) to 0.96 and all depreciation rates $(\delta_k, \delta_n, \delta_z)$ to 0.11, roughly in line with numbers in other studies.[30] Attempts to estimate the elasticity of substitution (σ) lead to estimates close to 1 but this estimate was very uncertain and we chose to set it to unity (Cobb-Douglas). We assume flexible prices $(\theta = 0)$. Later, we show that the dynamic responses are not very sensitive to changes in these parameters, which explains why they are poorly identified.

Our preliminary estimates suggested that variations in organizational capital play a small role, and we therefore decided to simplify the model by omitting this aspect. Thus, we set $\Omega = x = 1$ and omitted the equations that relate to Ω and *x*. We discuss this result in Section 7.

6.2 Replication of empirical responses

Figure 5 shows that the model replicates the empirical impulse-response functions almost perfectly. Also, the figure shows confidence intervals for the empirical IRFs, which have been calculated by bootstrapping (resampling the firms 1000 times with replacement).

In the first year, the increase in production is approximately 1.1 percent and *Figure 6* shows that most of this increase is achieved by increasing utilization by 0.8 percent, while the rest is achieved by increasing employment by approximately 0.3 percent. There is no response of the capital stock in the first year and this means that capital does not contribute to the increase in production in the first two years.

Adjustment costs and implementation lags in investment and hiring explain the slow adjustment of labor and capital. According to the estimates, the cost of utilization is not very convex, which can be seen from the fact that the marginal cost of production does not increase very much, although employment and capital respond very sluggishly. Note also that the price increases even less than the marginal cost, so the markup declines somewhat in response to an increase in demand.

[30] The depreciation rate for capital is a weighted average of the depreciation rates for machines and buildings used by Statistics Sweden. Bartelsman, Haltiwanger, and Scarpetta (2013) and Justiniano, Primiceri, and Tambalotti (2010) set the depreciation of capital to 0.10. Kryvtsov and Midrigan (2013) use monthly depreciation rates that imply yearly depreciation rates of 0.12 for capital and 0.13 for inventories. According to Statistics Sweden (AM 63 SM 1201), 12 percent of the permanently employed workers in Sweden left their jobs each year 1990-2011.

Figure 5. Effects of a firm-specific demand shock in theoretical and empirical model

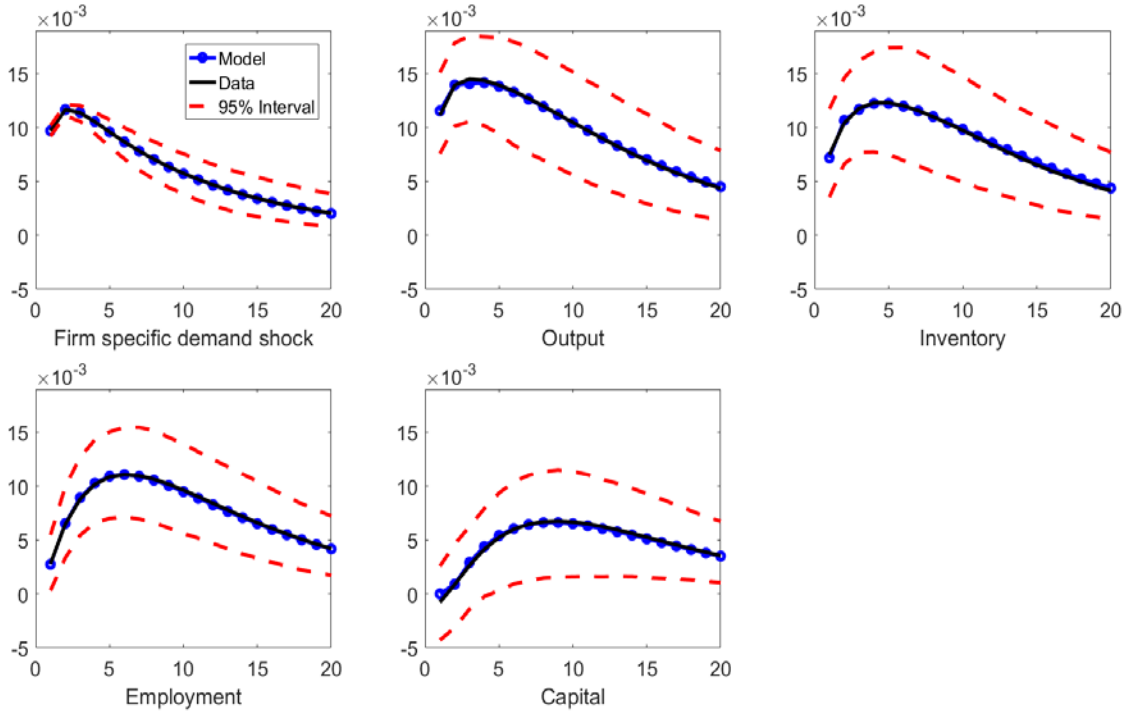

Figure 6. Effects of a firm-specific demand shock on variables in the theoretical model

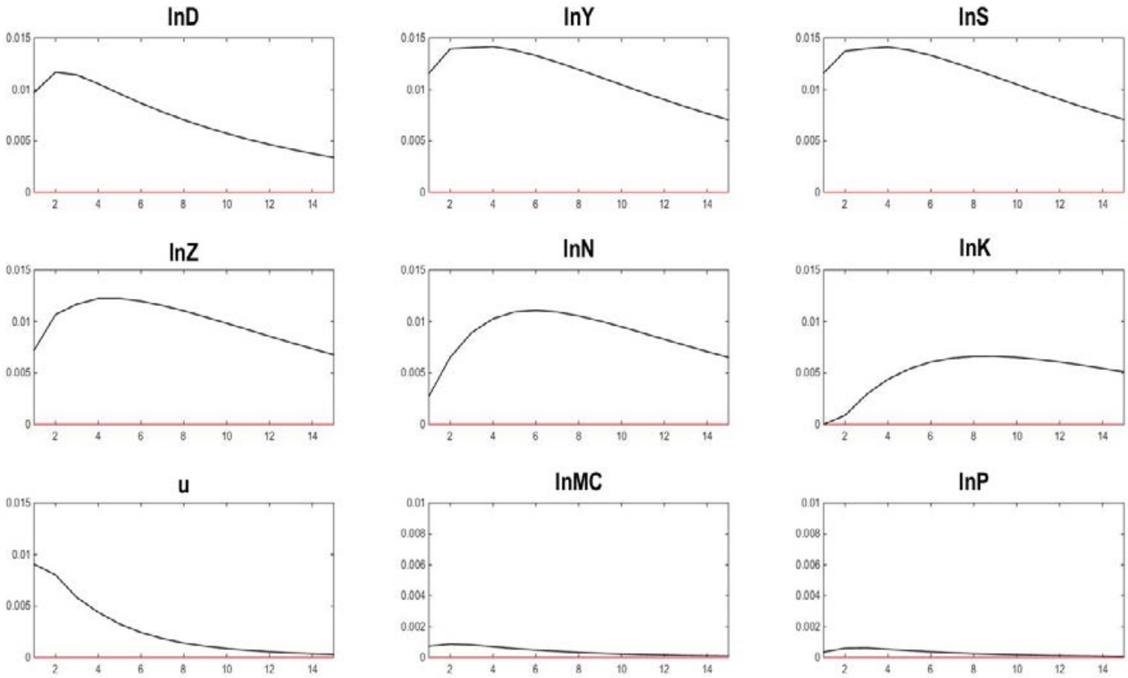

6.3 Parameter estimates and confidence intervals

The distributions obtained by bootstrapping are shown in *Figure 7* and the parameter estimates and confidence intervals are shown in *Table 2*. Convergence evaluation of the distributions shows that the distributions have converged. [31]

Figure 7. Distributions for deep parameters

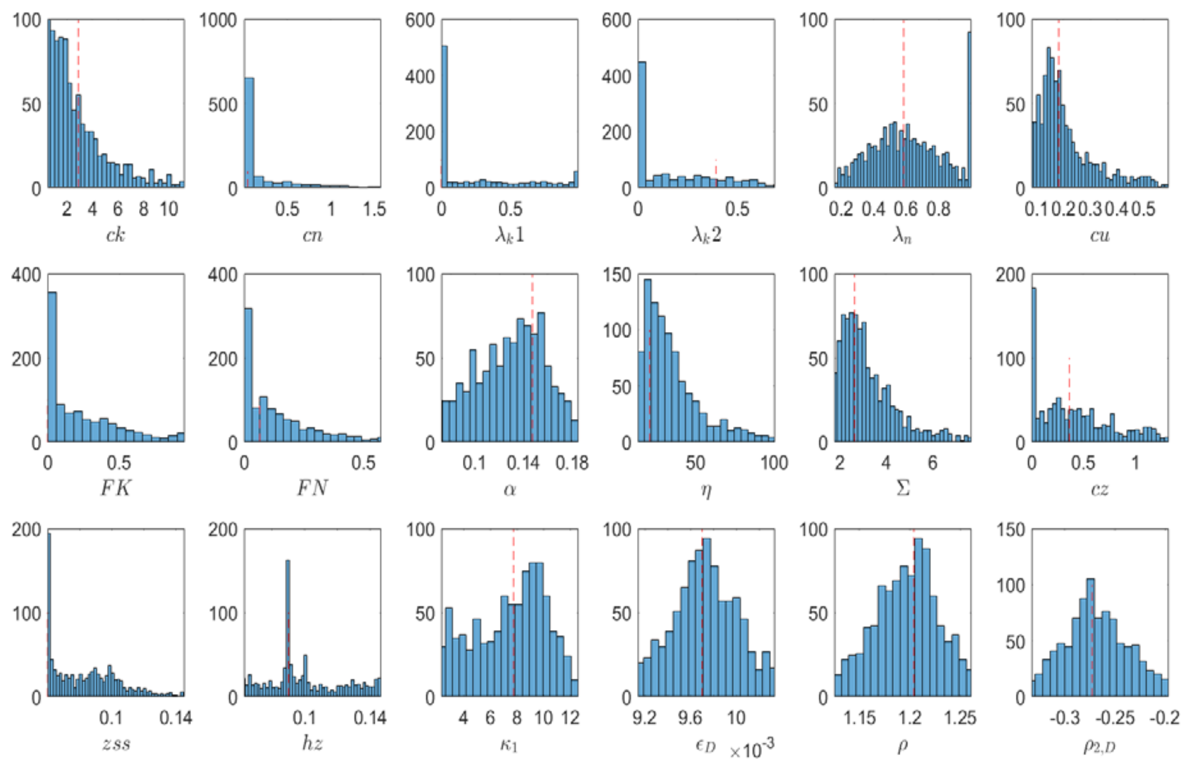

Note: 50 of 1000 estimates are outside the intervals shown.

6.4 Discussion of parameter values and their effects on the dynamics

Looking at *Figure 7* we see that many parameter estimates are uncertain. Still, the estimates indicate what type of model is needed to match the impulse-response functions. In this section, we discuss the parameter estimates and the impact that the different parameters have on the dynamics. *Figure A5* in the Appendix shows how the dynamics change when we change one parameter at the time, keeping the other parameters constant.

[31] *Figure A4* in the appendix shows the effect on the target function when we change one parameter while keeping all other parameters constant. These plots allow us to confirm that the estimate is at a minimum.

Adjustment costs $(c_k = 2.900, c_n = 0.041)$

Adjustment costs for capital play a central role. If we set the adjustment costs to zero, it has major effects on the dynamics: much more of the adjustment is made by changing the production factors and there is a much smaller change in utilization in response to a demand shock.

Implementation lags $(\lambda_{k1} = 0, \lambda_{k2} = 0.391, \lambda_n = 0.584)$

None of the investment in capital that is made in response to a demand shock is implemented in the same year that the shock occurs, 39 percent is implemented the year after the shock, and 61 percent is carried out 2 years after the shock. As one would expect, the share of hiring that occurs in the first year is higher (58 percent). If we omit implementation lags by setting $\lambda_{k1} = \lambda_n = 1$, the capital response speeds up and the utilization increases less, but otherwise, the responses are similar to the baseline.

Convexity of the utilization cost $(c_u = 0.182)$

Our estimates suggest that, with substantial adjustment costs and a relatively flat cost of utilization, firms meet an increase in demand by telling their workers to work more. If we set $c_u = 2$, so the cost of utilization becomes more convex, we obtain a very small increase in utilization. The output and inventory responses are reduced while labor increases more than production in the first years.

Fixed costs $(F_k = 0, F_n = 0.066)$

Increasing returns to scale were proposed by Hall (1988) as an explanation for why production varies more (in percentage terms) than employment in response to demand-side shocks. We find that increasing returns play little role: if we set $F_k = F_n = 0$, it has negligible effects on the dynamic responses.

Furthermore, increasing returns are not sufficient to match the impulse-response functions without variations in utilization. *Figure 8* shows estimates where we have eliminated variations in utilization by setting $\Omega = u = 1$. In this case, we find evidence of strongly increasing returns with respect to labor $(F_k = 0, F_n = 0.37)$ but the impulse-response functions are poorly matched. To understand why, note that the empirically estimated increase in employment is only 23 percent of the increase in output in the first year. We would need enormously increasing returns to scale to match this response. At the same time, employment almost catches up with production after a few years, and this observation is inconsistent with very strongly increasing returns to scale. Put differently, an increase in demand is associated with an *immediate* increase in measured factor productivity, but most of this effect is *temporary*, disappearing after some time. This observation is in line with Sbordone (1996, 1997).

Figure 8. Effects of a firm-specific demand shock in a model without variations in utilization or organizational capital and empirical model

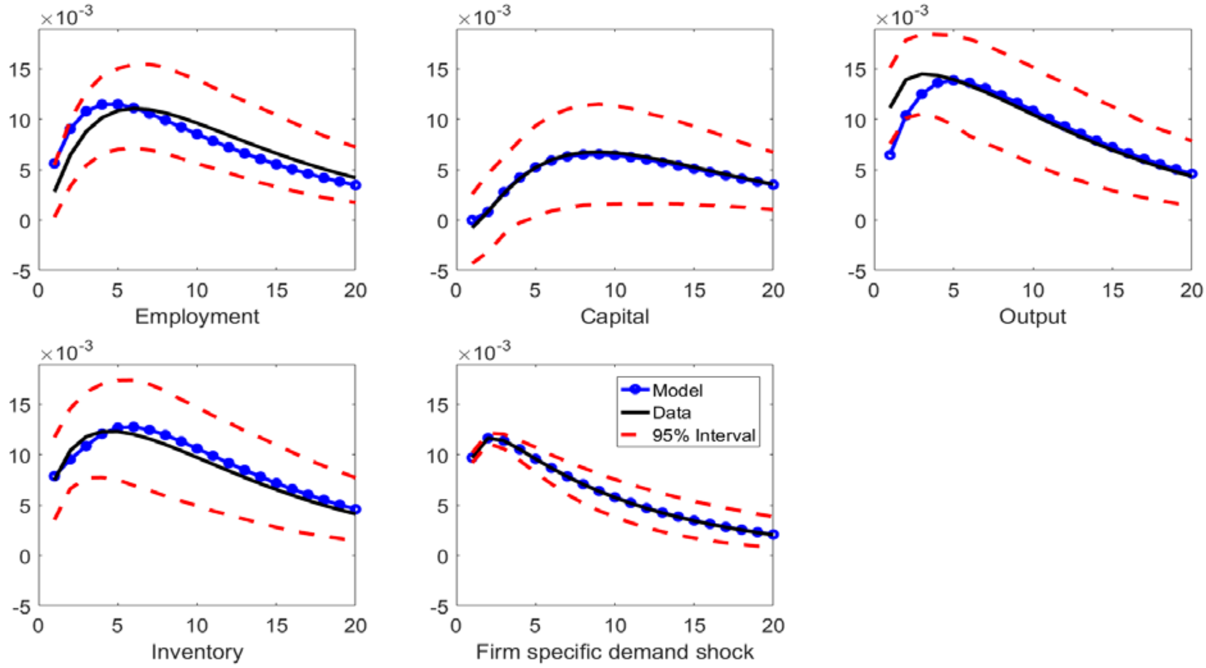

Note: The model is constrained by setting $u = \Omega = 1$.

Organizational capital

We did not find an important role for organizational capital. Instead, we could match the impulse-response functions very well with flexible utilization. Note that variations in effective work hours (u_t) and variations in time spent investing in organizational capital (x_t) have similar effects on production today. The difference is that variations in utilization are associated with a direct cost today, while disinvestment in organizational capital shows up as lower productivity in the future.

Table 2. Parameter estimates and confidence intervals

Description	Param.	Estimate	90 % min	90 % max	min	max
Capital adjustment cost	ck	2.900	0.478	11.260	0.00	300
Labor adjustment cost	cn	0.041	0.000	1.583	0.00	300
Investment implemented in t	λk1	0.000	0.000	0.990	0.00	1.00
Investment implemented in t+1	λk2	0.391	0.000	0.686	0.00	1.00
Hiring implemented in t	λn	0.584	0.159	1.000	0.00	1.00
Slope of utilization cost	cu	0.182	0.081	0.593	0.01	10000
Fixed capital	FK	0.000	0.000	0.952	0.00	3.00
Fixed employment	FN	0.066	0.000	0.574	0.00	3.00
Importance of capital	α	0.147	0.072	0.185	0.01	0.99
Price elasticity	η	20.582	13.057	100.542	2.00	500
Sensitivity to demand shock	Σ	2.653	1.779	7.669	0.10	10
Inventory holding cost	cz	0.365	0.000	1.324	0.00	2.00
Stock finished goods/production	zss	0.060	0.060	0.145	0.00	1.00
Stock inputs/production	hz	0.090	0.063	0.147	0.00	0.20
Kappa 1 (inventory model)	κ_1	7.748	2.320	12.616	0.20	700
Std. of the shock	ε	0.010	0.009	0.010	0.00	0.02
First term in demand process	ρ	1.204	1.125	1.261	-2.00	2.00
Second term in demand process	ρ2	-0.273	-0.333	-0.197	-2.00	2.00
Fixed parameters		**Value**				
Subjective discount factor	β	0.96				
Depreciation of capital	δk	0.11				
Depreciation of employment	δn	0.11				
Depreciation of inventory	δz	0.11				
Price stickiness	θ	0				
Restriction		**Estimate**	**min**	**max**		
Wage share of value added		0.73	0.73	0.83		
Capital/value added		0.78	0.78	0.88		
Inventories/production		0.15	0.15	0.25		
Stocks of inputs/finished goods		1.50	0.67	1.50		
Implied values						
Steady state markup		0.08				
Steady state profits/value added		0.11				
Kappa 2 (inventory model)		0.57				
Kappa 3 (inventory model)		11.77				
Target function		0.71				

To investigate the differences between the two models, we re-estimated the full model, allowing organizational capital to vary but imposing a very steep marginal cost of utilization $\left(c_u = 10000\right)$. The result was that we failed to match the strong first-period increase in production and the resulting impulse-response functions look very similar to *Figure 8*. Our interpretation is that this has to do with the hump-shape of the demand shock. The initial unexpected increase in demand is followed by an expected further increase in demand in the following year, so the need for organizational capital is even higher in the year after the initial shock, and this makes it unprofitable to de-cumulate organizational capital in the first period. For this reason, a model without variations in effort is unable to match the strong increase in production in the first period. As it turns out, we can match the impulse-response functions almost perfectly without variations in organizational capital once we allow for flexible factor utilization.

Distribution of factor returns $\left(\alpha = 0.147\right)$

The parameter α and the markup determine factor returns. The cost of labor is 73% of value added, which is the lower bound that we set for the labor share.

Price elasticity and markup $\left(\eta = 20.582,\ \mu_{ss} = 0.079\right)$

In a model without inventories, a price elasticity of 21 would imply a markup equal to 5 percent, but the possibility of stock-out makes the effective demand curve less elastic for a given inventory stock. A steady-state markup of 7.9 percent is in the reasonable range. Carlsson and Smedsaas (2007) estimate the markup for Swedish manufacturing firms to be 17%.

Demand sensitivity $\left(\Sigma = 2.653\right)$

According to the estimate, a demand shock of one percent shifts the demand curve of the typical firm by 2.7 percent, but the estimate is very uncertain as can be seen in *Figure 7*.[32]

---

[32] The high point estimate may reflect the fact that many manufacturing firms produce investment goods and durable goods, the demand for such goods being relatively sensitive to shocks. Aggregate consumption and investment, which are used to construct the demand variable, contain large portions of services, and we know that the demand for services is much more stable than the demand for manufactured goods. Furthermore, investment in stocks of intermediate inputs by *other* firms in the same industry will respond to the demand shock and contribute to the volatility of demand for the goods produced by an individual firm. Note also that we capture the combined effect of the demand shock and the resulting industry price response and both shift the demand for goods produced by an individual firm (see discussion in Section 3.3).

That this parameter is poorly identified can be understood by noting that different combinations of demand shifts and price responses can lead to the same increase in sales. There may be a large shift of the demand curve $(\Sigma \text{ large})$ but the effect on sales may be mitigated by a substantial price response combined with a high price elasticity $(\eta \text{ large})$. Alternatively, there may be a small shift of the demand curve $(\Sigma \text{ modest})$ and the effect on sales may be only partially mitigated because of a small price response and/or a low price elasticity $(\eta \text{ small})$. We would need price data and other types of shocks, e.g. cost shocks, to better pin down these parameters.

Inventory model $\left(c_z = 0.365, \hat{Z}_{ss} = 0.060, h = 0.090, \kappa_1 = 7.748\right)$

In our model, there are three costs of holding finished goods inventories: the financing cost, depreciation of the inventory stock $(\delta_z \cdot Z)$ and a storage cost $(c_z \cdot Z)$. Having set $\beta = 0.96$ and $\delta_z = 0.11$, we find that the yearly storage cost is large (37 percent of the value), steady-state inventories end up at the lower bound (15 percent of production) and 2/3 of the inventory stock consists of inputs. It appears that, by setting finished-goods inventories at the lower bound, we are able to match the fact that the inventories track production with a modest lag. Since there are four structural parameters in the inventory model $(\Lambda, \Psi, \mathrm{T}, \Phi)$ but only three estimated parameters $(\kappa_1, \kappa_2, \kappa_3)$ the latter do not have clear economic interpretations.

Elasticity of substitution $(\sigma = 1 \text{ imposed})$

Preliminary estimates indicted that the elasticity of substitution is poorly identified from these data. If we reduce the elasticity of substitution to 0.4 based on estimates by Chirinko, Fazzari and Meyer (2011), there is a smaller increase in employment, because the marginal product of labor falls more rapidly for a given capital stock, and a larger increase in utilization, but otherwise, the dynamics are similar.

Discount factor and depreciation rate for capital $(\beta = 0.96, \ \delta_k = 0.11 \text{ imposed})$

We had difficulty estimating these parameters, so we set them at reasonable values. If we change the discount factor to 0.92 or the depreciation rate for capital to 7 percent, it has little effect on the dynamic responses. This explains why these parameters are poorly identified.

Separation rate for labor $(\delta_n = 0.11 \text{ imposed})$

The parameter δ_n can be interpreted as an exogenous separation rate, and we would then expect a value of approximately 0.10. There is, however, an alternative interpretation, which is that a firm incurs adjustment costs when it changes the *number* of workers instead of when it hires; this would imply that $\delta_n = 0$. If we set $\delta_n = 0$, this has very small effects on the responses, so this parameter is not well identified with the data that we have.

Depreciation rate for inventories $(\delta_z = 0.11 \text{ imposed})$

The depreciation rate for inventories was also set to 11 percent. If we increase it to 20 percent, it has small effects on the dynamics. If we decrease it to 0.05, it affects the dynamics, but we view such a depreciation rate as implausibly low. Technical changes and changes in fashion and design may make goods unsellable, so the depreciation rate of finished goods should be relatively high.

Price stickiness $(\theta = 0 \text{ imposed})$

Although we assume that prices are completely flexible, there is a very weak price response to the demand shock, which explains why we could not estimate the degree of price rigidity with any precision. The weak price response is due to utilization being very flexible, so the marginal cost increases only slightly in response to the demand shock. Note, however, that we are not using price data in the estimation, so our inference about prices is indirect.

7. Relation to previous research

Adjustment costs

Like Galeotti, Maccini and Schiantarelli (2005) and many others, we find substantial adjustment costs and these play a key role in the dynamics. We are fully aware that quadratic adjustment costs are a crude approximation and that many authors have found evidence of asymmetric, linear or lumpy adjustment costs.[33] We are not questioning these results, but we see quadratic adjustment costs as a useful stand-in for various types of adjustment costs. Our

[33] For reviews and recent evidence, see Hamermesh and Pfann (1996), Adda and Cooper (2003), Cooper and Haltiwanger (2006), and Bond and Van Reenen (2007). Adjustment costs associated with employment could also represent search frictions, but the results in Carlsson, Eriksson and Gottfries (2013) and Stadin (2015) contradict this interpretation.

model is meant to capture the *average* reaction of firms to demand shocks, and we manage to fit the average response very well with this simple specification of adjustment costs.

Increasing returns vs. factor utilization

Hall (1988) shows that shocks that should be uncorrelated with technology are associated with variations in output that are more than proportional to the corresponding variation in inputs, which he interprets as evidence of increasing returns to scale. With increasing returns, firms make losses if the price is equal to the marginal cost, and since firms typically do not make losses, not even in periods of low demand, Hall concludes that firms must have very substantial market power. He estimates markups of more than 100 percent for most industries.

Increasing returns may take the form of increasing returns in the long-run production function, or there may be some form of *short-run increasing returns* because some factors of production are *quasi-fixed*. Hall (1988) discusses a case where some predetermined amount of overhead labor determines the firm's maximum production capacity, and a fixed amount of production labor is needed per unit of output actually produced. This means that overhead labor is a fixed cost in the short run, and when the firm operates below capacity, the marginal cost is the cost of the required production work; again, the firm would make a loss if the price would be equal to the marginal cost. Importantly, Hall assumes that there is no cost of increasing the utilization of overhead labor.

An alternative explanation is that there are variations in utilization (e.g., effort), so the services of labor and capital vary more than the measured inputs, and variations in utilization are *costly*. Burnside, Eichenbaum and Rebelo (1993) find that a model with constant returns to scale, perfect competition, implementation lags in employment and variations in effort fits the data well and that it can account for a positive correlation between the growth rates of the Solow residual and government expenditures. Note that if the cost of increasing utilization is not properly accounted for, the marginal cost will be underestimated.

Our identification strategy is conceptually similar to that used by Hall (1988) and we also find that firms have market power, but increasing returns are not found to be important for medium-term dynamics. Instead, costly variations in utilization play an important role. Note, however, that we do not have data for hours worked – only the reported number of "full-time equivalent employees" which likely does not fully reflect the variations in hours. This means that the variation in utilization that we find may represent variations in unregistered work hours or effort but also registered overtime. In this respect, our results are not directly comparable to Hall's results.

Like Burnside, Eichenbaum and Rebelo (1993), we find that adjustment lags and variations in utilization are important, but our dynamic specification is quite different. Their model has no adjustment costs, but employment is determined one period in advance, leading to variations in utilization with very low persistence. According to our estimates, adjustment costs lead to very sluggish adjustments and large and persistent variations in utilization. In this respect, our results are similar to those of Fairise and Langot (1994) and Braun and Evans (1998).

Imbs (1999) adjusts Solow residuals for variations in utilization of capital and labor, and he finds that the adjusted residuals are substantially less pro-cyclical than standard series. In his model, utilization can be backed out due to specific functional forms for the utilization costs.[34] Our utilization cost function is more general, so our results are more data-driven and less dependent on theoretical assumptions, but the conclusions are similar.

Straight time and overtime

Lucas (1970), Sargent (1978) and Hansen and Sargent (1988) argue that imperfect substitution between straight time and overtime can help to explain the pro-cyclical pattern of the standard Solow residual. Hansen and Sargent (1988) write the production function $Y_t = h_1 A_t K_t^{\alpha} n_{1t}^{1-\alpha} + h_2 A_t K_t^{\alpha} n_{2t}^{1-\alpha}$, where n_{1t} is straight time and n_{2t} is overtime. This production function is equivalent to our production function when the elasticity of substitution is set to unity. To see this, rewrite it as $Y_t = A_t K_t^{\alpha} n_{1t}^{1-\alpha}\left(h_1 + h_2\left(n_{2t}/n_{1t}\right)^{1-\alpha}\right)$. Changing notation so that $n_{1t} = N_t$ and $u_t = h_1 + h_2\left(n_{2t}/n_{1t}\right)^{\alpha}$, we obtain our production function.[35]

External economies of scale

Caballero and Lyons (1990, 1992) show that when aggregate output is included in industry-level production-function regressions, the estimated coefficient for this variable becomes positive and significant for many industries, and there is little evidence of internal economies of scale. Their interpretation is that there are external economies of scale, and such "Marshallian" externalities have been considered by many authors; see Cooper and Haltiwanger (1996) and Braun and Evans (1998).

[34] In Imbs' (1999) model, the cost functions for utilization of capital and labor have only one free parameter, so that parameter can be backed out from the steady-state conditions. This is not the case in our model.

[35] Hall (1996) estimates a model with straight time and overtime on macro data and finds that such a model provides greater magnification and propagation of shocks than the model by Burnside, Eichenbaum and Rebelo (1993).

In line with Caballero and Lyons (1990, 1992), we find that internal economies of scale play a small role, but aggregate externalities cannot explain our results because any such effects are picked up by the time dummies. There could also be externalities within an industry, but the estimated impulse-response functions are qualitatively similar when we include interactions between time dummies and industry dummies, suggesting that externalities within industries are not driving our results (see *Figure A3* in the Appendix).

Sbordone (1996, 1997) argues that labor hoarding could explain the results of Caballero and Lyons (1990, 1992). The point is that aggregate output may act as a proxy for factor utilization when there are demand-side shocks and it takes time to adjust labor and capital. Sbordone shows that aggregate output has a *persistent* effect on industry-level output but only a *transitory* effect on industry-level productivity, which supports the labor hoarding interpretation. Our results point strongly in the same direction.

Time spent on maintenance, cleaning, organizing and training

There is evidence that workers spend a substantial fraction of their time on tasks that increase future rather than current production; see, e.g., Fay and Medoff (1985) and Kim and Lee (2007). Kim and Lee show theoretically that even without adjustment costs, skill accumulation will be countercyclical in a real business cycle model because the opportunity cost of skill accumulation is higher when productivity is high, and similar ideas have emerged in the growth literature.[36] Bean (1990) constructs a real business cycle model in which production factors are used either in production or to accumulate human capital. He shows that in periods of high government expenditures, there will be less human capital accumulation, so measured productivity will be high. Bean argues that this is consistent with the observed high measured productivity during wars in data for the UK, and he finds that a shock to government expenditures has a negative long-run effect on growth.

We incorporated these ideas in our model by assuming that workers spend some of their time building "organizational capital" that increases future productivity. However, our estimates did not support the idea that firms invest less in organizational capital when there is high demand. The model with a relatively flat cost of utilization does a better job matching the empirical impulse-response functions.

---

[36] See Aghion and Saint-Paul (1998) and DeJong and Ingram (2001). Cooper and Johri (2002) assume instead that there is learning-by-doing, so the accumulation of "organizational capital" is positively related to the level of production.

Price rigidity

Rotemberg and Summers (1990) argue that price rigidity can explain pro-cyclical productivity under perfect competition. They assume that firms must fix prices before demand is known. With free entry, the price must be equal to the *average* cost, so the price will be above the marginal cost in a recession. In booms, firms produce at the point where the marginal cost equals the predetermined price, so there is rationing in periods of high demand. Thus, the price will be higher than the marginal cost on average, which could explain the results found by Hall (1988).

In our model, the markup is sufficiently large and utilization is sufficiently flexible so that price exceeds the marginal cost throughout the adjustment to a one standard deviation shock and this means that firms always want to satisfy demand.[37] However, a positive demand shock makes demand high relative to the stock of finished goods, so there is increased rationing in the sense that a larger fraction of customers do not find their desired variety.

Inventory dynamics

As noted by Bils and Kahn (2000), inventory behavior provides clues to the nature of business cycles. Some researchers have viewed pro-cyclical inventory investment as evidence that the costs of producing must be low in boom periods.[38] Our study has nothing to say about the role of productivity shocks, but we find a very strong positive response of inventory holdings to demand shocks, which is well explained by a model with a stock-out motive, as suggested by Kahn (1987, 1992). Other studies finding support for inventory models with a stock-out motive are Bils and Kahn (2000), Wen (2005), Galeotti, Maccini and Schiantarelli (2005) and Kryvtsov and Midrigan (2013).[39]

While inventory investments respond positively to demand shocks, they fail to keep pace with shipments. In the year that the shock occurs, there is a large increase in production and sales but only a small increase in the stock of finished goods, so the *ratio* of finished-goods inventories to production (and sales) decreases when demand increases. To understand this, note that the first-order condition for finished-goods inventory holdings is

[37] The finding of a flat marginal cost curve is consistent with Carlsson and Nordström-Skans (2012), but not with Galeotti, Maccini and Schiantarelli (2005).
[38] See e. g. Khan and Thomas (2007).
[39] In a closely related approach, Kydland and Prescott (1982), Christiano (1988) and Ramey (1989) introduce inventories as a factor of production.

$$MC_t - \beta\left(1-\delta_z\right)E_t\left(MC_{t+1}\right) + c_z = \left(\kappa_1 - 2\kappa_3\frac{\hat{Z}_t}{D_t^{\Sigma}P_t^{-\eta}}\right)\left(P_t - MC_t\right). \tag{46}$$

where $MC_t = v_t + m$. The left-hand side is the cost of holding an additional unit of finished-goods inventories: the marginal cost today plus the storage cost minus the expected discounted marginal cost next year. The right-hand side is the effect of inventories on sales (due to reduced stock-outs) times the markup (the value of selling one more unit). In a model without stock-outs, inventories have no effect on sales, so the right-hand side is zero. This implies pure production-smoothing: inventories are adjusted until the sum of the marginal cost today and the storage cost is equal to the discounted marginal cost next year.

When inventories contribute to sales, the desired ratio of inventories to demand depends on prices. As noted by Bils and Kahn (2000), at least one of the following must occur for the *ratio* of inventories to sales to *decrease* when there is an increase in demand:

i) marginal costs increase relative to discounted future marginal costs, making it more expensive to hold inventories, or

ii) the markup declines so the gain from holding inventories is reduced.

Bils and Kahn argue that there is little evidence of predictable changes in marginal costs but that the markup is lower in booms. Looking closely at *Figure 6* we can see that marginal cost increases when the shock occurs and then it increases further between year 1 and 2 because of the hump-shape of demand. Thus, the decline in the inventory/sales ratio in the first year must be due to a decline in the markup, as argued by Bils and Kahn (2000).

But why does the markup decline? This is harder to understand. But note that, without price rigidity, the first order condition for the price can be written

$$P_t S_t = \eta\mu_t\left(\frac{\kappa_2}{\Phi} + \Phi\kappa_3\left(\frac{\hat{Z}_t}{\Phi D_t^{\Sigma}P_t^{-\eta}}\right)^2\right)\Phi D_t^{\Sigma}P_t^{-\eta} \tag{47}$$

where $\Phi D_t^{\Sigma}P_t^{-\eta}$ is potential sales (what the firm would sell if it would never stock out). When considering an increase in the price, the firm balances the direct effect on revenue from a higher price against the resulting decline in sales. As demand increases, there is a decline in the ratio of finished goods to potential sales and this reduces the price sensitivity of demand, which should *increase* the markup. But there is also another effect, which comes from the fact that the direct increase in revenue from a price increase is proportional to *actual* sales while the negative effect on sales is proportional to *potential* sales. Because of stock-outs, actual sales increase much less (1.2 percent) than potential sales (1.8 percent) in the first year and this explains the decline in the markup.

We should note some caveats here, however. First, we do not use price data in the estimation. Second, a more realistic model of the demand side and financial conditions may lead to a more countercyclical markup and this might change our conclusions; see Rotemberg and Woodford (1999), Lundin et al. (2009), and Gilchrist et.al (2017). Third, we identify responses to shocks that are more persistent than normal business cycle fluctuations.

Permanent and transitory shocks

Our autoregressive model allows for a hump-shaped and persistent demand shock, but we do not allow for unit roots and we do not distinguish permanent and transitory shocks. An alternative would be to allow for permanent and transitory shocks, as in Franco and Philippon (2007) and Carlsson, Messina and Nordström Skans (2017), but we leave this for future research. Carlsson, Messina and Nordström Skans (2017) use long-run restrictions to identify demand shocks, and they also find that demand shocks have substantial effects on production, employment and production per worker.

8. Conclusion

Investment, hours worked, labor productivity and inventory holdings are all pro-cyclical, but there is no consensus on how to interpret these correlations. A positive correlation between production and output per worker may arise because productivity shocks drive both variables or because demand-side shocks lead to variations in factor utilization. A positive correlation between production and inventory holdings may arise because firms build up inventories when it is easy to produce or because firms need more inventory holdings in order to not stock out when demand is high. Since all variables are endogenous on the macroeconomic level, it is difficult to establish causality without additional assumptions about functional forms and the stochastic nature of the shocks, as in estimation of DSGE models.

The same problem occurs if we use panel data for individual firms: a positive correlation between production and output per worker can be interpreted in different ways, and without some exogenous source of variation, it is difficult to establish causality.

In this paper, we tried to study the *causal effects* of demand-side shocks on firms' dynamic adjustments. We combined macro data with input-output tables, foreign industry data, and micro-level export shares to identify demand-side shocks that can be taken as exogenous for individual firms. We used firm-level panel data to study how firms in general react to such shocks. We found that production and inventory holdings respond quickly to an increase in demand while the registered inputs of labor and capital respond slowly. Then, we used this information to estimate the structural parameters of a theoretical model. We found that the responses can be well explained by a theoretical model with adjustment costs, implementation lags, variable utilization and a stock-out avoidance motive for inventories. Although most individual parameter estimates are uncertain, we think that our estimates indicate what kind of model is needed to match these responses.

We study only the effects of specific demand-side shocks, so we cannot draw any conclusions about the relative importance of supply and demand shocks for business cycle fluctuations. Still, our estimates may be useful as references for researchers who are estimating or calibrating macro models of the business cycle. Specifically, we have found evidence of substantial variations in factor utilization. Any study that ignores this aspect will, most likely, produce incorrect conclusions. Variations in utilization may be interpreted as technology shocks and the markup will be vastly overestimated if the marginal cost of utilization is ignored.

The analysis could be extended in many ways. Financial constraints could be included in the theoretical model. Other sectors than manufacturing could be examined and heterogeneity between firms could be analyzed. Effects of other shocks could be analyzed and other outcome variables could be included in the empirical analysis.

An outcome variable of particular interest is the price that the firm sets. Although price setting is an important aspect of the theoretical model, we did not include prices in the empirical model. The main reason is that the sample would be drastically reduced if we would include only those firms for which we have firm-specific price data. A second reason is that, to properly model price setting, we would need to take account of customer relations in the product market, which slow down the effect of a price change on sales. We omitted customer relations in the present study because the theoretical model is already quite complicated.[40]

It is worth emphasizing that our methodology is fundamentally different from that used in many studies that estimate general equilibrium models. In such studies, the modeling of labor supply plays a central role. In our estimation, the time dummies pick up any general equilibrium effects. We assume that each individual firm can hire as much labor and buy as much capital as it desires at given prices, which may vary over time. We ask how an

[40] Lundin et al. (2009) and Gilchrist et.al (2017) study the interaction between customer relations and financial constraints.

individual firm reacts to a firm-specific shock. This makes our analysis more limited and partial but also more focused compared to estimation of a general equilibrium model. An improved understanding of firm behavior can help us to design key building blocks of our macroeconomic models.

References

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, 2015, Identification Properties of Recent Production Function Estimators, *Econometrica*, 83, 2411-2451.

Adda, Jérome, and Russel Cooper, 2003, *Dynamic Economics – Quantitative Methods and Applications*, MIT Press, Cambridge, Massachusetts.

Aghion, Philippe, and Gilles Saint-Paul, 1998, Virtues of Bad Times, *Macroeconomic Dynamics*, 2, 322-344.

Arellano, Manuel, and Stephen Bond, 1991, Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations, *The Review of Economic Studies*, 58, 277–297.

Baily, Martin N., Eric J. Bartelsman and John Haltiwanger, 2001, Labor Productivity: Structural Change and Cyclical Dynamics, *Review of Economics and Statistics* 83, 420-433.

Bartelsman, Erik, John Haltiwanger, and Stefano Scarpetta, 2013, Cross-Country Differences in Productivity: The Role of Allocation and Selection, *American Economic Review*, 103, 205-334.

Basu, Susanto, 1996, Procyclical Productivity: Increasing Returns or Cyclical Utilization?, *Quarterly Journal of Economics*, 111, 719-751.

Basu, Susanto, and John G. Fernald, 1997, Returns to Scale in U.S. Production: Estimates and Implications, *Journal of Political Economy*, 105, 249-283.

Bean, Charles R., 1990, Endogenous Growth and the Procyclical Behavior of Productivity, *European Economic Review*, 34, 355-363.

Biddle, Jeff E., 2014, The Cyclical Behavior of Labor Productivity and the Emergence of the Labor Hoarding Concept, *Journal of Economic Perspectives*, 28, 197-212.

Bils, Mark, and James A. Kahn, 2000, What Inventory Behavior Tells Us about Business Cycles, *American Economic Review*, 90, 458-481.

Blinder, Alan S., and Louis J. Maccini, 1991, Taking Stock: A Critical Assessment of Recent Research on Inventories, *Journal of Economic Perspectives*, 5, 73-96.

Bond, Stephen, and John Van Reenen, 2007, Microeconometric Models of Investment and Employment, in *Handbook of Econometrics*, 6A, Elsevier.

Braun, R. Anton, and Charles Evans, 1988, Seasonal Solow Residuals and Christmas: A Case for Labor Hoarding and Increasing Returns, *Journal of Money, Credit and Banking*, 30, 306-330.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 1993, Labor Hoarding and the Business Cycle, *Journal of Political Economy*, 101, 245–273.

Carlsson, Mikael, Stefan Eriksson, and Nils Gottfries, 2013, Product Market Imperfections and Employment Dynamics, *Oxford Economic Papers*, 65, 447–470.

Carlsson, Mikael, Julián Messina, and Oskar Nordström-Skans, 2017, Firm-Level Shocks and Labor Flows, mimeo, Uppsala University.

Carlsson, Mikael, and Oskar 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, Mikael, and Jon Smedsaas, 2007, Technology Shocks and the Labor-Input Response: Evidence from Firm-Level Data, *Journal of Money, Credit and Banking*, 39, 1509–1520.

Chirinko, Robert S., Steven M. Fazzari, and Andrew P. Meyer, 2011, A New Approach to Estimating Production Function Parameters: The Elusive Capital–Labor Substitution Elasticity, *Journal of Business & Economic Statistics*, 29, 587–594.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 2005, Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy, *Journal of Political Economy*, 113, 1–45.

Christiano, Lawrence J., 1988, Why Does Inventory Investment Fluctuate so Much? *Journal of Monetary Economics*, 21, 247-280.

Cooper, Russel W., and John C. Haltiwanger, 2006, On the Nature of Capital Adjustment Costs, *Review of Economic Studies*, 73, 611-633.

Cooper, Russel W., and Alok Johri, 2002, Learning-by-doing and Aggregate Fluctuations, *Journal of Monetary Economics*, 49, 1539-1566.

DeJong, David, and Beth F. Ingram, 2001, The Cyclical Behavior of Skill Acquisition, *Review of Economic Dynamics*, 4, 536-561.

Eriksson, Stefan, and Karolina Stadin, 2017, What are the Determinants of Hiring? The Importance of Product Market Demand and Search Frictions, *Applied Economics*, 49, 5144-5165.

Fairise, Xavier, and Francois Langot, 1994, Labor Productivity and the Business Cycle: Can R.B.C Models be Saved?, *European Economic Review*, 38, 1581-1594.

Fay, Jon A., and James L. Medoff, 1985, Labor and Output over the Business Cycle: Some Direct Evidence, *American Economic Review*, 75, 638-655.

Field, Alexander J., 2010, The Procyclical Behavior of Total Factor Productivity in the United States, 1890-2004, *Journal of Economic History*, 70, 326-350.

Franco, Francesco, and Thomas Philippon, 2007, Firms and Aggregate Dynamics, *Review of Economics and Statistics*, 89, 587–600.

Gilchrist, Simon, Raphael Schoenle, Jae Sim and Egon Zakrajšek, 2017, Inflation Dynamics during the Financial Crisis, *American Economic Review*, 107, 785-823.

Galeotti, Marzio, Louis J. Maccini and Fabio Schiantarelli, 2005, Inventories, Employment and Hours, *Journal of Monetary Economics*, 52, 575-600.

Gertler, Mark, and Peter Karadi, 2015, Monetary Policy Surprises, Credit Costs, and Economic Activity, *American Economic Journal: Macroeconomics*, 7, 44-76.

Greenwood, Jeremy, Zvi Hercowitz and Per Krusell, 2000, The Role of Investment-Specific Technological Change in the Business Cycle, *European Economic Review*, 44, 91-115.

Hall, Robert E., 1988, The Relation between Price and Marginal Cost in U.S. Industry, *Journal of Political Economy*, 96, 921-947.

Hansen, Gary D., and Thomas J. Sargent, 1988, Straight Time and Overtime in Equilibrium, *Journal of Monetary Economics*, 21, 281-308.

Hamermesh, Daniel S., and Gerard Pfann, 1996, Adjustment Costs in Factor Demand, *Journal of Economic Literature*, 34, 1264-1292.

Hart, Robert A., and James R. Malley, 1999, Procyclical Labor Productivity: A Closer Look at a Stylized Fact, *Economica*, 66, 533-550.

Humphreys, Brad R., Louis J. Maccini and Scott Schuh, 2001, Input and Output Inventories, *Journal of Monetary Economics*, 47, 347-375.

Imbs, Jean M., 1999, Technology, Growth, and the Business Cycle, *Journal of Monetary Economics*, 44, 65-80.

Justiniano, Alejandro, Giorgio Primiceri, and Andrea Tambalotti, 2010, Investment Shocks and Business Cycles, *Journal of Monetary Economics*, 57, 132–145.

Kahn, James A., 1987, Inventories and the Volatility of Production, *American Economic Review*, 77, 667-679.

Kahn, James A., 1992, Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding, *Quarterly Journal of Economics*, 107, 481-510.

Khan, Aubhik, and Julia K. Thomas, 2007, Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies, *American Economic Review*, 97, 1165-1188.

Kim, Daehaeng, and Chul-In Lee, 2007, On-the-job Human Capital Accumulation in a Real Business Cycle Model: Implications for Intertemporal Substitution Elasticity and Labor Hoarding, *Review of Economic Dynamics*, 10, 494-518.

Kryvtsov, Oleksiy, and Virgiliu Midrigan, 2013, Inventories, Markups, and Real Rigidities in Menu Cost Models, *Review of Economic Studies*, 80, 249-276.

Kydland, Finn, and Edward Prescott, 1982, Time to Build and Aggregate Fluctuations, *Econometrica*, 50, 1345-1370.

Levinsohn, James, and Amil Petrin, 2003, Estimating Production Functions Using Inputs to Control for Unobservables, *Review of Economic Studies*, 70, 317-341.

Lucas, Robert E. Jr., 1970, Capacity, Overtime, and Empirical Production Functions, *American Economic Review*, 60, Papers and Proceedings 23-27.

Lundin, Magnus, Nils Gottfries, Charlotte Bucht and Tomas Lindström, 2009, Price and Investment Dynamics: Theory and Plant-Level Data, *Journal of Money, Credit, and Banking*, 41, 907-934.

Mickelsson, Glenn, 2016, Estimation of DSGE Models with Uninformative Priors, in Mickelsson, Glenn, *DSGE Model Estimation and Labor Market Dynamics*, PhD Thesis, Economic Studies 163, Department of Economics, Uppsala University.

Mickelsson, Glenn, and Karolina Stadin, 2016, Demand Shocks and Labor Hoarding: Matching Micro Data, in Mickelsson, Glenn, *DSGE Model Estimation and Labor Market Dynamics*, PhD Thesis, Economic Studies 163, Department of Economics, Uppsala University.

Nelder, J. A., and R. Mead, 1965, A Simplex Method for Function Minimization, *The Computer Journal*, 7, 308–313.

Olley, Steven, and Ariel Pakes, 1996, The Dynamics of Productivity in the Telecommunications Equipment Industry, *Econometrica*, 64, 1263-1295.

Prescott, Edward C., 1986, Theory Ahead of Business Cycle Measurement, *Carnegie-Rochester Conference Series on public Policy*, 25, 11-44.

Ramey, Valerie A., 1989, Inventories as Factors of Production and Economic Fluctuations, *The American Economic Review*, 79, 338-354.

Ramey, Valerie A., 2016, Macroeconomic Shocks and Their Propagation, *Handbook of Macroeconomics*, 2, 71-162.

Rotemberg, Julio J., and Lawrence H. Summers, 1990, Inflexible Prices and Procyclical Productivity, *Quarterly Journal of Economics*, 105, 851-873.

Rotemberg, Julio J., and Michael Woodford, 1999, The Cyclical Behavior of Prices and Costs, *Handbook of Macroeconomics*, edited by John Taylor and Michael Woodford, Amsterdam: North-Holland Elsevier.

Sargent, Thomas J:, 1978, Estimation of Dynamic Labor Demand Schedules under Rational Expectations, *Journal of Political Economy*, 86, 1009-1044.

Sbordone, Argia M., 1996, Cyclical Productivity in a Model of Labor Hoarding, *Journal of Monetary Economics*, 38, 331-361.

Sbordone, Argia M., 1997, Interpreting the Procyclical Productivity of Manufacturing Sectors: External Effects or Labor Hoarding? *Journal of Money, Credit, and Banking*, 29, 26-45.

Stadin, Karolina, 2014, The Dynamics of Firms’ Factor Demand, in Stadin, Karolina, *Employment Dynamics*, PhD Thesis, Economic Studies 145, Department of Economics, Uppsala University.

Stadin, Karolina, 2015, Firms’ Employment Dynamics and the State of the Labor Market, Working Paper Series 2015:20, IFAU.

Urban, Timothy L., 2005, Inventory models with inventory-level-dependent demand: A comprehensive review and unifying theory, *European journal of Operations Research*, 162, 792-804.

Wen, Yi, 2005, Understanding the Inventory Cycle, *Journal of Monetary Economics*, 52, 1533-1555.

APPENDIX

Table A1. Industries used in the estimations and the number of firms in each industry in the baseline panel (all in the manufacturing sector)

Industry (SNI 92)	Number of firms
15 Food products and beverages	31
17 Textiles	13
18 Wearing apparel; furs	2
19 Leather and leather products	4
20 Wood and products of wood and cork (except furniture)	54
21 Pulp, paper, and paper products	30
22 Printed matter and recorded media	16
23 Coke, refined petroleum products and nuclear fuels	1
24 Chemicals, chemical products, and man-made fibers	30
25 Rubber and plastic products	65
26 Other non-metallic mineral products	16
27 Basic metals	18
28 Fabricated metal products, except machinery and equipment	189
29 Machinery and equipment n.e.c.	168
30 Office machinery and computers	2
31 Electrical machinery and apparatus n.e.c.	44
33 Medical, precision and optical instruments, watches, and clocks	24
34 Motor vehicles, trailers, and semi-trailers	54
35 Other transport equipment	8
36 Furniture; other manufactured goods n.e.c.	49
Total number of firms	818

Note: Industry definitions in SNI92 and SNI2002 are almost the same at the 2-digit level.

Calculation of real capital stock

The real capital stock (*Kr*) consists of machines and buildings. In the firm-level panel data, we have firms' book values of buildings and machinery, but generous depreciation allowances imply that the book values of these stocks are much lower than their economic values. With a too low value of the stock of capital, we would exaggerate the volatility of the capital stock measured as log changes. For this reason, we tried to construct a better measure of the real capital stock. We did this in three steps:

First, we obtained industry-level estimates of capital stocks and book values from Statistics Sweden. Using these data for the years 2000-2005, we calculated an *average* ratio of book value to economic value at the industry level (2-digit SNI92) for buildings and machines separately. This ratio was then used to scale up the book values of buildings and machines for the first year that a firm appears in the sample.

Adding buildings and machines together and dividing by a price index for investments, we express the first-year capital stock in prices in the year 2000.

Finally, we calculated capital stocks for subsequent years by subtracting depreciation and adding investments in machines and buildings deflated by the investment price index. This was repeated for each year that the firm appears in the sample. We set the depreciation rate of capital to 11 percent based on a weighted average of the depreciation rates for buildings and machines used by the Statistics Sweden.

Note in *Table 1* that the median ratio of (owned) real capital to total real production (output) – 0.33 – is more than twice as large as the median book value of capital relative to production – 0.15. This is due to depreciation being much higher in the accounting than the estimated economic depreciation rate (11 percent). With larger capital stocks, the implied log changes are correspondingly smaller.

Value added is approximately 39 percent of total production (output) in this sample, and the median ratio of real capital to value added is 87 percent. This may strike readers as a low value, but a substantial fraction of the capital that firms use is rented. Firms often lease cars, trucks and other types of equipment, and they may also rent the buildings in which they operate. We do not have data for rented capital and we treat it as a flexible input in our theoretical model.

Figure A1. Growth rate of the number employed and hours worked in manufacturing, extraction, energy and environment (percentage)

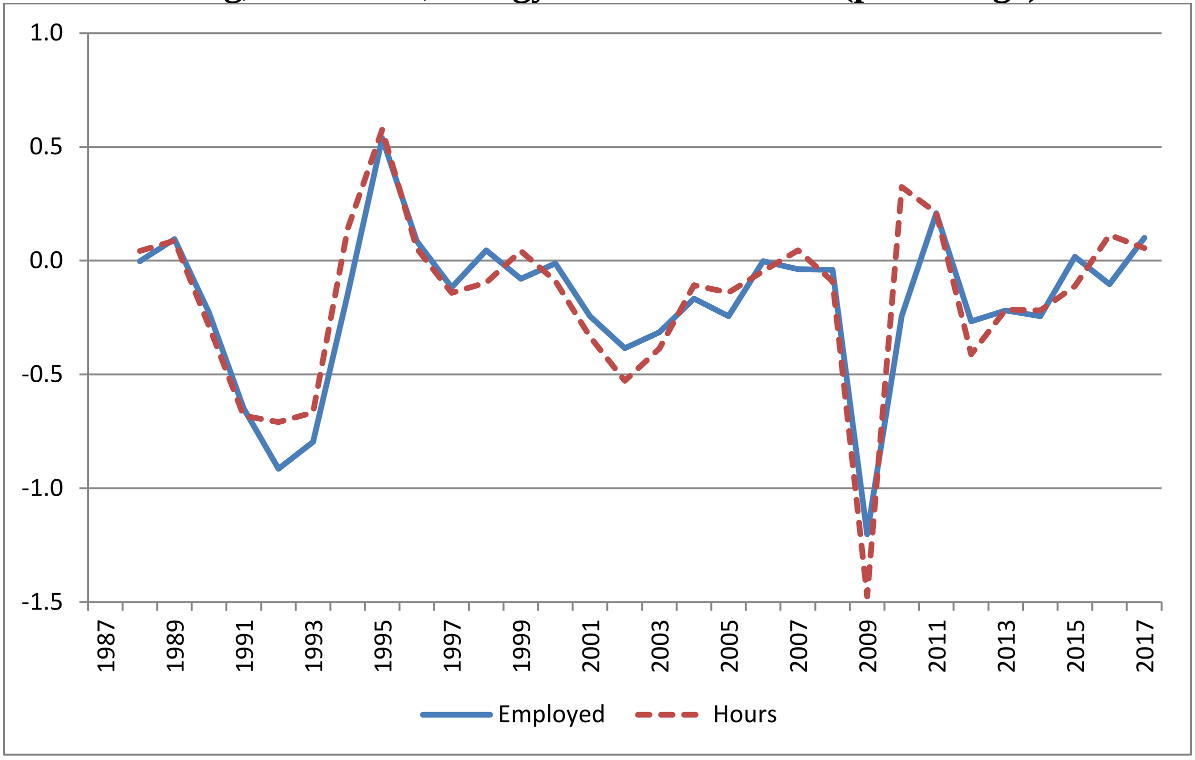

Note: The series includes extraction, energy and environment because of data availability. The correlation between the two variables is 0.91. Data from Statistics Sweden homepage.

Figure A2. The relative price of investment goods

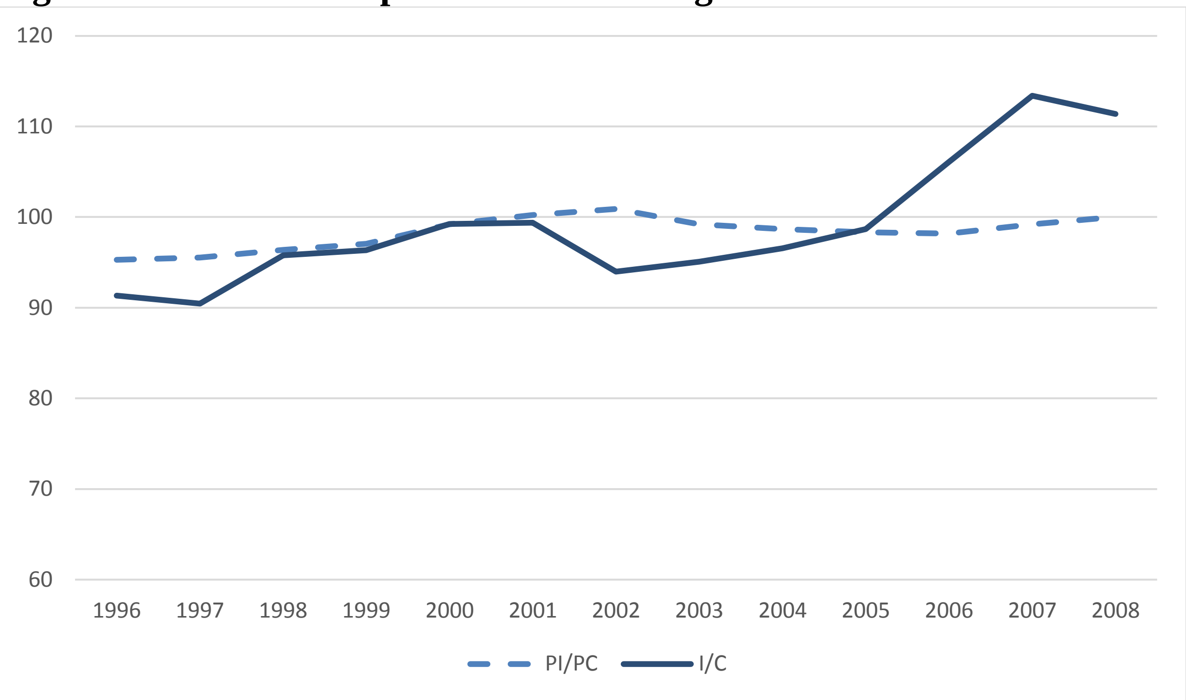

Note: I/C is the ratio of investment to consumption (volume indexes), and PI/PC is the ratio of the corresponding deflators. Aggregate national account data from Statistics Sweden.

Table A2. Estimated empirical model

a) Equations for firm-level production, employment, capital stock, and inventories

	(1) lnYr	(2) lnN	(3) lnKr	(4) lnZr
L.lnYr	0.551***	0.134***	0.101***	0.188***
	(0.020)	(0.013)	(0.014)	(0.021)
L2.lnYr	0.024	-0.033***	-0.024	0.009
	(0.017)	(0.011)	(0.015)	(0.019)
L.lnN	0.119***	0.657***	0.059***	0.115***
	(0.021)	(0.029)	(0.020)	(0.028)
L2.lnN	-0.019	0.019	0.000	-0.020
	(0.020)	(0.020)	(0.021)	(0.024)
L.lnKr	0.037***	0.059***	0.830***	0.015
	(0.013)	(0.011)	(0.021)	(0.019)
L2.lnKr	-0.029**	-0.047***	-0.075***	-0.016
	(0.013)	(0.009)	(0.016)	(0.017)
L.lnZr	0.148***	0.072***	-0.003	0.492***
	(0.013)	(0.009)	(0.010)	(0.022)
L2.lnZr	-0.055***	-0.034***	-0.014	0.007
	(0.012)	(0.008)	(0.010)	(0.016)
lnD	1.148***	0.289**	-0.080	0.770***
	(0.206)	(0.139)	(0.178)	(0.216)
L.lnD	-0.721***	-0.076	0.120	-0.479**
	(0.204)	(0.133)	(0.184)	(0.226)
Observations	8,180	8,180	8,180	8,180
Number of FAD_F_Id	818	818	818	818
R-squared	0.574	0.652	0.680	0.507

b) AR(2) process for product demand

	lnD
L.lnD	1.201***
	(0.027)
L2.lnD	-0.270***
	(0.027)
Observations	8,180
Number of FAD_F_Id	818
R-squared	0.993
St.D. of residual	0.0097

Note: Robust standard errors in parentheses, clustered at the firm level; *** p<0.01, ** p<0.05, * p<0.1. Time dummies and firm fixed effects are included in all regressions. Baseline balanced panel of private firms in manufacturing with at least 10 employees existing 1997-2008. Firms with extreme values or missing values are removed. The estimation method is OLS. The demand process was specified as AR(2) because the coefficient on the third lag was not significantly different from zero.

Robustness of impulse-response functions

Figure A3 shows the impulse-response functions obtained for alternative samples and specifications. *Panel a* repeats the results from the baseline estimation. *Panel b* shows the results for an unbalanced panel, where we include firms that existed for only part of the sample period, and *Panel c* shows the results for a sample where we exclude only extreme observations and not the entire time series for firms with an extreme observation. In both cases, the impulse-response functions are very similar to our baseline results. *Panel d* shows estimates for the whole manufacturing sample including extreme observations. The main difference is that production now responds immediately while the lag is similar for inventories and employment. In *Panel e*, we add firms in the construction industry to the sample, and this has virtually no effect on the estimates. *Panel f* shows estimates for one-plant firms only. The responses are more sluggish but qualitatively similar.

Panel g shows the impulse-response functions for a model where we include 3 lags in the estimated empirical model. The responses of production and inventories become more hump-shaped, peaking 2 years after the peak in demand. *Panel h* shows the IRFs for a model where we include industry-specific linear trends in the model. The overall picture is similar, but the shock becomes less persistent when some of the industry variation is mopped up by the trend, the half-time of the demand shock falls from 9 to 7 years, and the inventory response becomes weaker.

In *Panel i*, we have mopped up the cross-industry variation by adding interactions between industry dummies and time dummies. This means that the only reason why the demand variable varies across firms is because firms have different export shares. Now, the magnitude of the demand shock is only 57 percent of the demand shock in the baseline estimation, which is not surprising because the shock is defined differently and reflects differences between different markets rather than differences between different industries. Conditional on the size of the typical shock, the qualitative response is similar to the baseline, with production and inventories responding first, then labor, and finally capital responding very sluggishly.

As discussed in *Sections 2* and *3*, one may be worried that the demand variable is not exogenous after all. For this reason, it is interesting to estimate the model using only the domestic or only the foreign component of the demand variable. In *Panel j*, we use only the domestic component of demand, $(1-\delta_i)\left[\phi_j^C \ln C_t + \phi_j^G \ln G_t + \phi_j^I \ln I_t + \left(1-\phi_j^C - \phi_j^G - \phi_j^I\right)\ln EX_t\right]$

and in *panel k*, we use only the foreign component of the demand variable, $\delta_i\left(\sum_m \omega_{j,m} \ln Y^F_{j,m,t}\right)$.

The responses become smaller conditional on the shock and the profiles change. Still, the timing of the responses is similar to the baseline, with production and inventories responding first, and labor and capital responding more slowly.

Panel l shows estimates where we measure production by *real value added* instead of total production (output). Production and the capital stock respond more strongly in this case, but otherwise, the results are similar. For the baseline estimation, we do not use real value added because we view the deflators of value added as more unreliable than the producer price indexes.

To sum up, all alternative estimations show that production, inventories, labor and the capital stock respond positively to the demand shock. *Table A3* summarizes the impulse-responses by showing the peak responses for the different variables and specifications. While the exact timing varies a bit, the qualitative results in specifications a) to l) are similar: production and inventories respond relatively quickly to the demand shock while labor responds with a longer lag and capital with an even longer lag.

We also divided the sample into larger and smaller firms, where larger firms are defined as having a mean employment (across years) of at least 50 employees. The responses of larger firms (*Panel m*) are similar to the baseline estimation, but larger firms appear to be more cyclical in the sense that they respond more to the demand shock. It is likely that large firms produce more investment goods and durable goods. Smaller firms (*Panel n*) respond much less to the demand shock, and the response of the capital stock is weaker. One reason may be that small firms often rent the capital that they use. Clearly, there is interesting heterogeneity among firms, and our baseline estimation should be viewed as a rough estimate of the average responses across firms.

Finally, *Panel o* shows the results when we include industry trends but exclude the time dummies. Now, the shock is much less persistent and reflects business cycle variation to a much greater extent. As explained above, we do not use this variation for the estimation because of the risk of spurious correlation due to unobserved aggregate shocks. If some unobserved exogenous factor causes a boom in investment that contributes to higher aggregate demand, we will have reverse causation if we estimate without time dummies.

Figure A3. Robustness of impulse-response functions

a) Baseline estimation

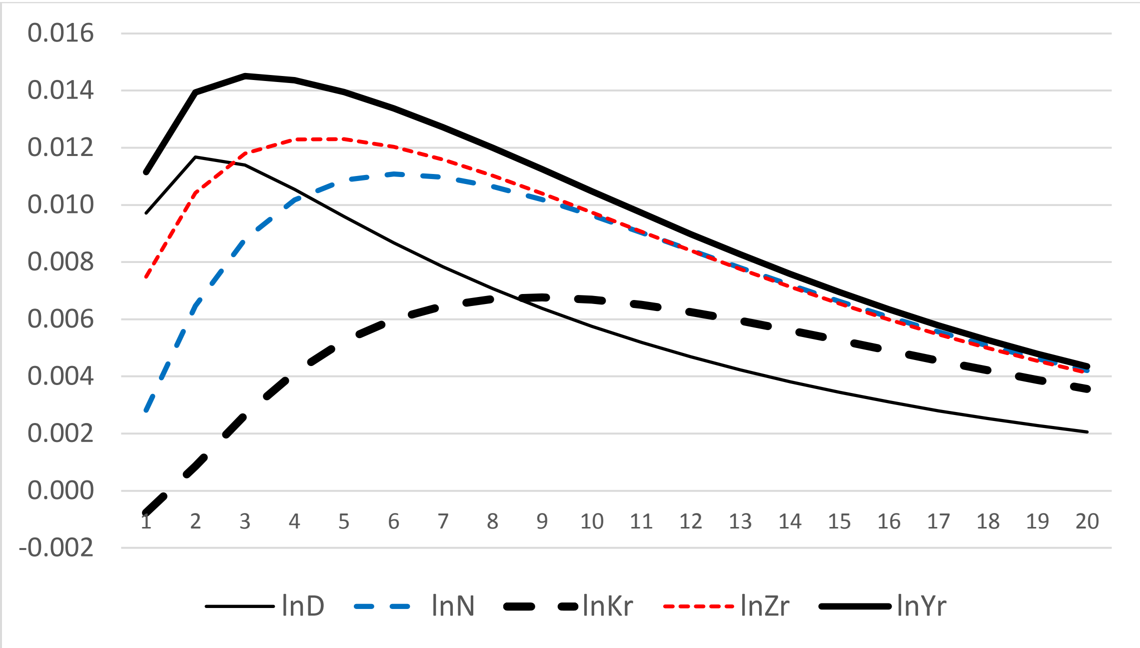

Note: Impulse-response functions from reduced-form model estimated on a balanced panel consisting of private manufacturing firms with at least ten employees, no extreme observations, and no missing values 1997-2008 (818 firms). Firm and time fixed effects are included. The variables are in logs and the time units on the horizontal axis are years. With two lags, the number of observations included in the estimations is 8180.

b) Unbalanced panel

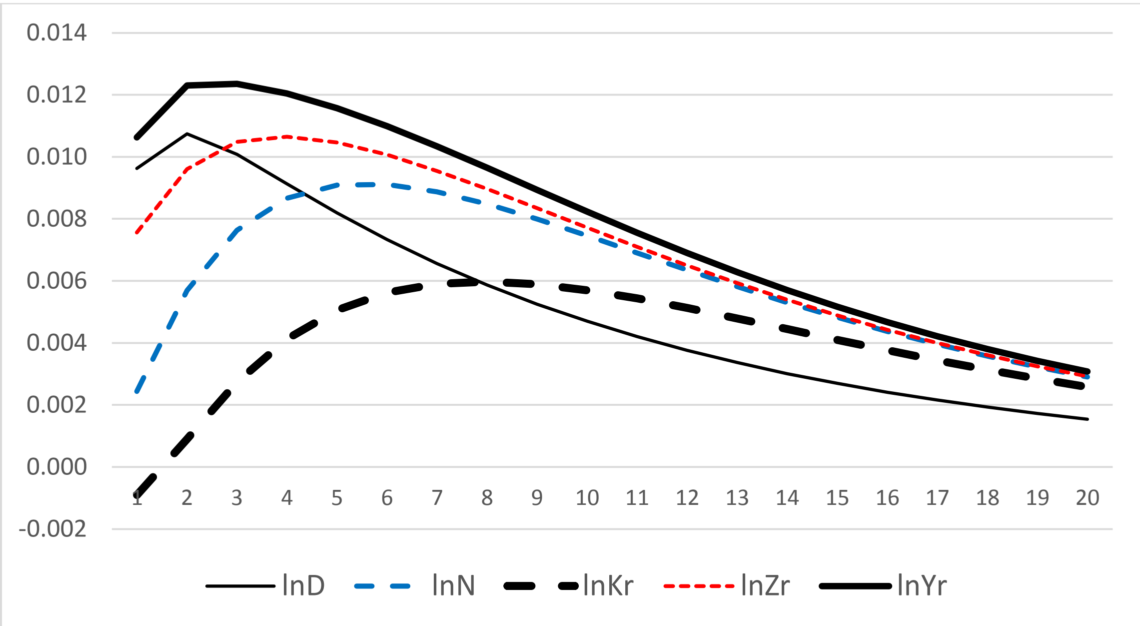

Note: Private manufacturing firms with at least ten employees, excluding firms with extreme observations. Number of observations included in the estimations: 11306.

c) Excluding extreme observations instead of firms with extreme observations

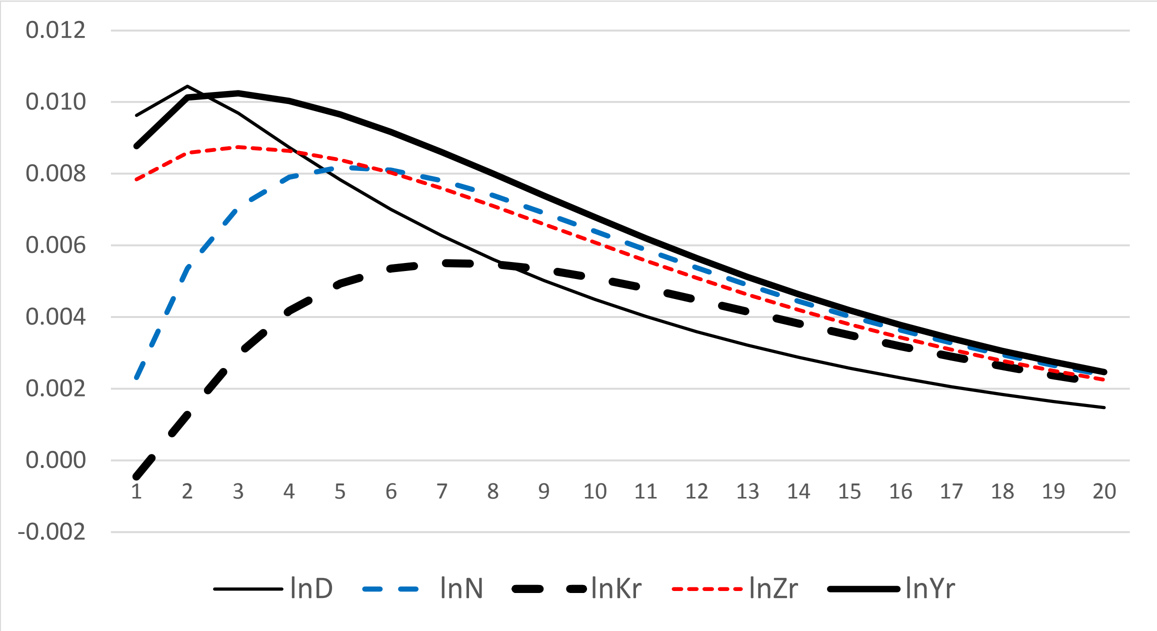

Note: Private manufacturing firms with at least ten employees. Extreme observations excluded, but we do use data for other years for firms with extreme observations. Number of observations used in the estimations: 17179.

d) Including extreme observations

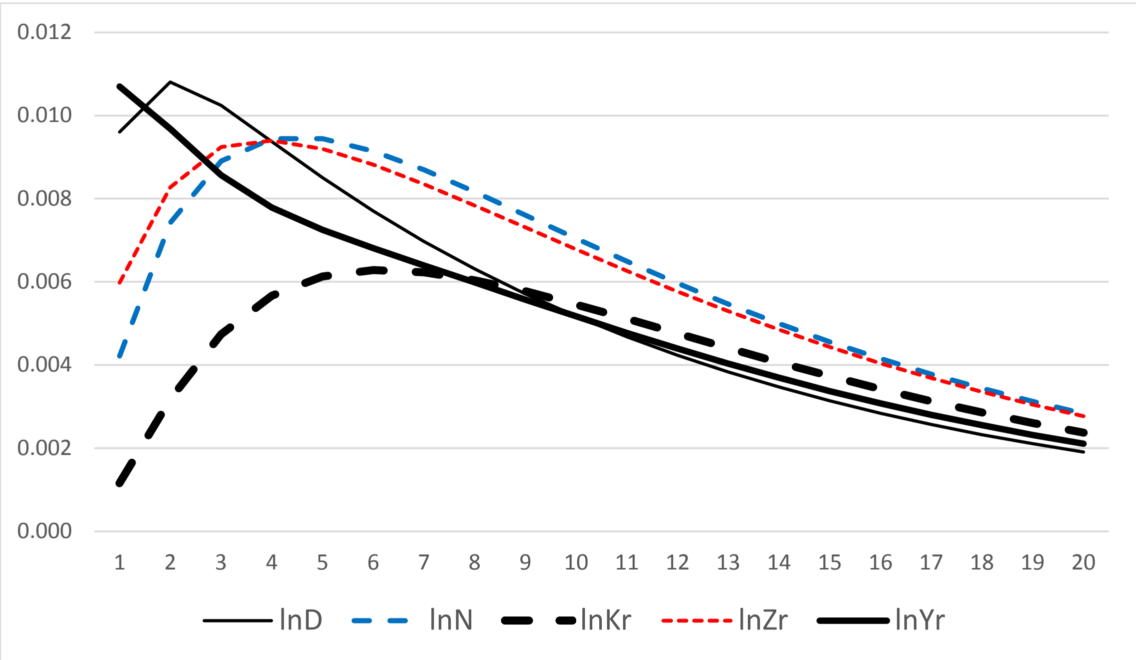

Note: Private manufacturing firms with at least ten employees. Unbalanced panel with extreme observations included. Number of observations used in the estimations: 31164.

e) Adding construction industry firms to the baseline sample

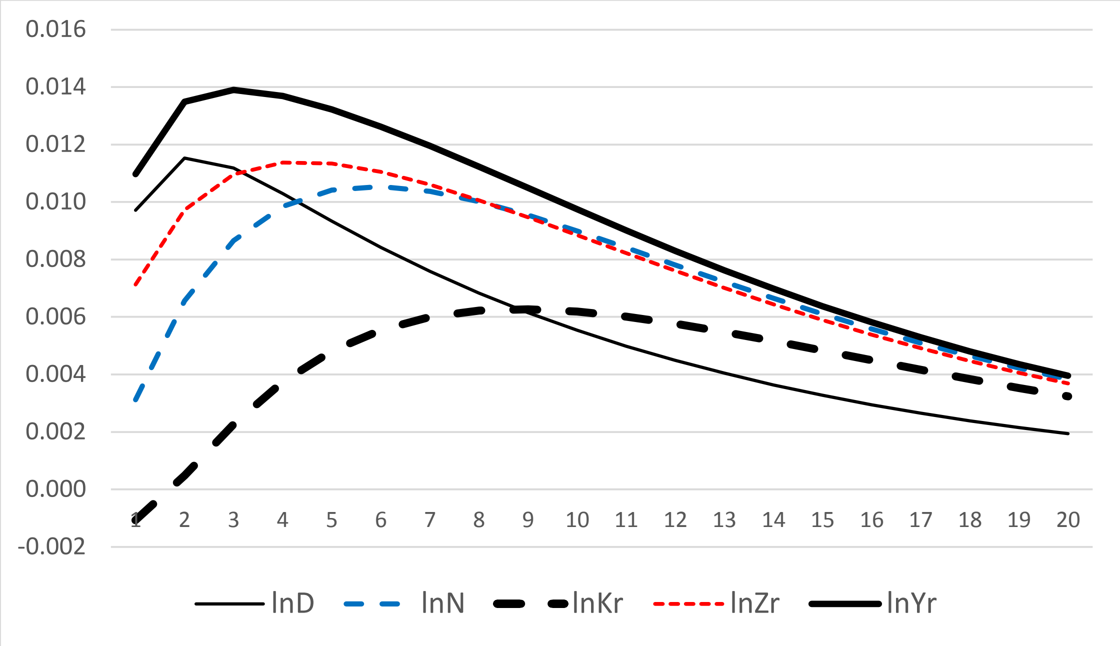

Note: Panel of private firms in manufacturing and construction with at least ten employees, excluding firms with extreme observations or missing values. Number of observations included in the estimations: 8570.

f) One-plant firms only

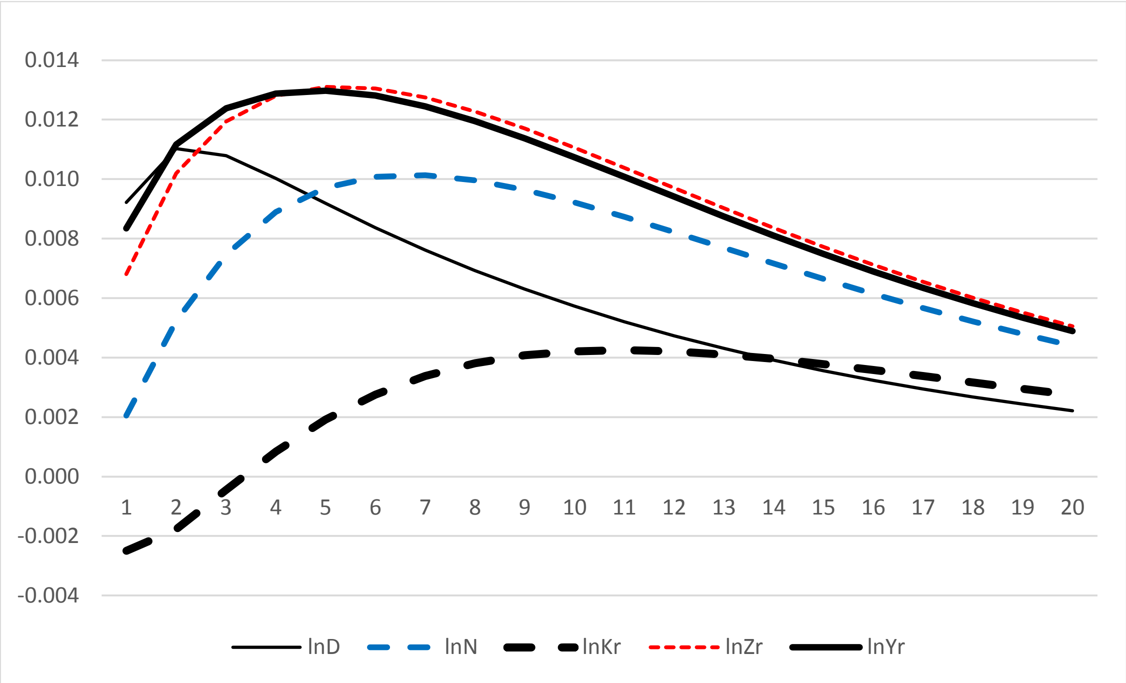

Note: Panel of private firms in manufacturing with at least ten employees and only one plant, excluding firms with extreme observations or missing values. Number of observations included in the estimations: 6124.

g) Including three lags of the endogenous variables

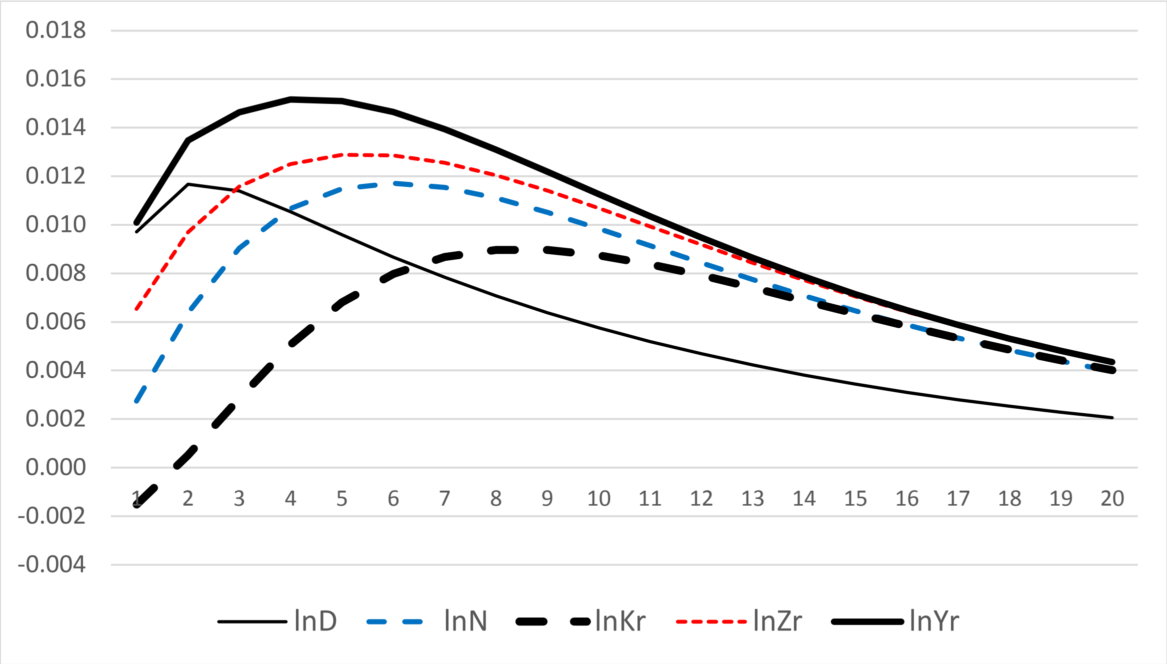

Note: Baseline panel of private manufacturing firms with at least ten employees, excluding firms with extreme observations or missing values. Number of observations included in the estimations: 7386.

h) Including linear industry trends

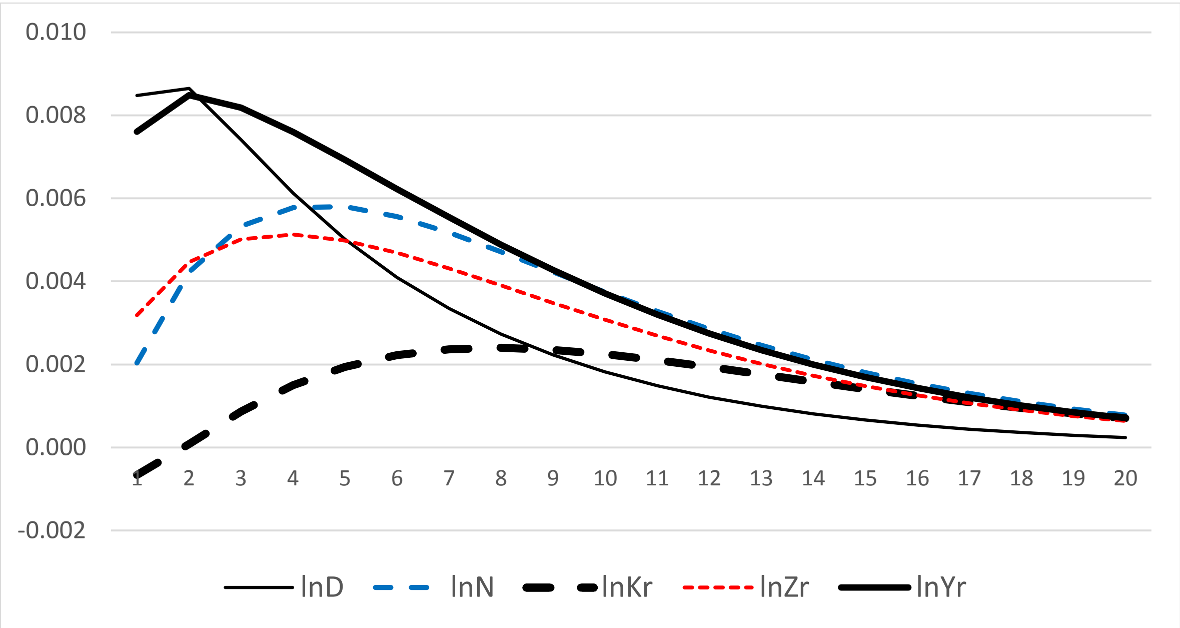

Note: Baseline panel of private manufacturing firms with at least ten employees, excluding firms with extreme observations or missing values. Number of observations used in the estimations: 8180.

i) Including industry dummies interacted with time dummies

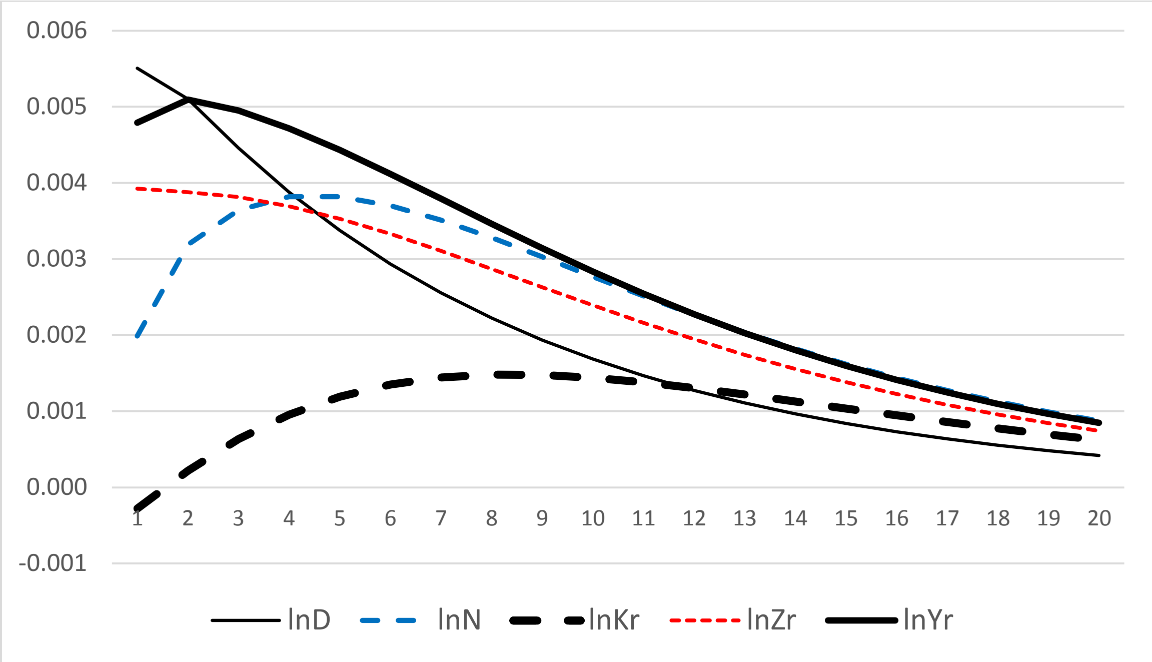

Note: Baseline panel of private manufacturing firms with at least ten employees, excluding firms with extreme observations or missing values. Number of observations used in the estimations: 8180.

j) Using only the domestic part of the demand variable

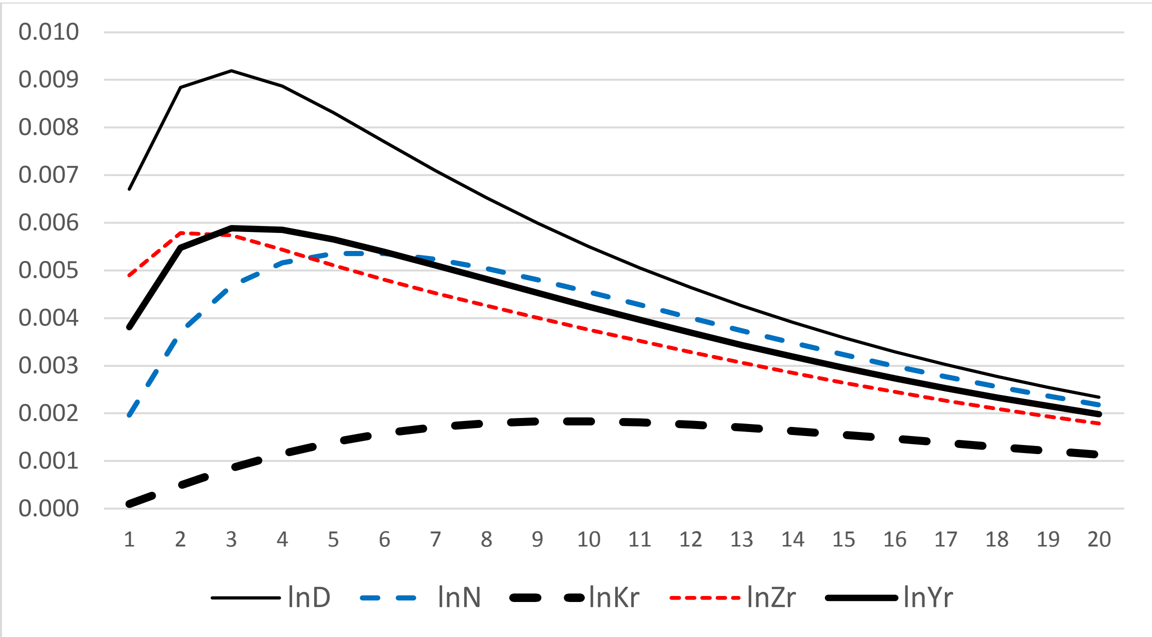

Note: Baseline panel of private manufacturing firms with at least ten employees, excluding firms with extreme observations or missing values. Number of observations used in the estimations: 8180.

k) Using only the foreign part of the demand variable

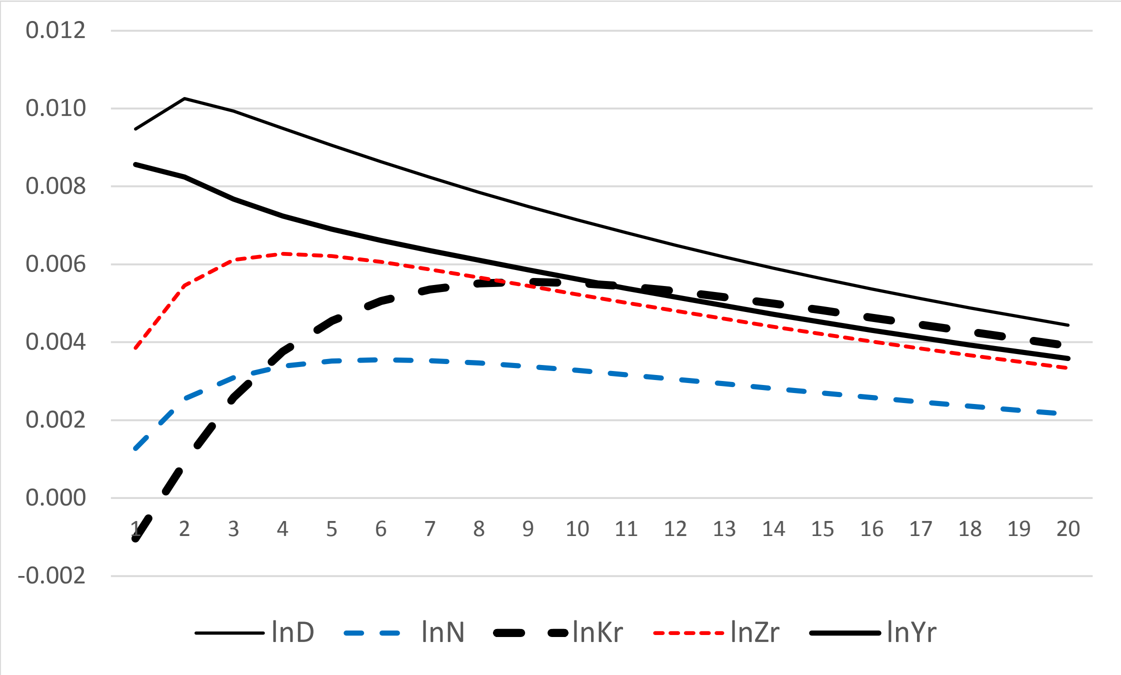

Note: Baseline panel of private manufacturing firms with at least ten employees, excluding firms with extreme observations or missing values. Number of observations used in the estimations: 8180.

l) Using value added as the measure of production

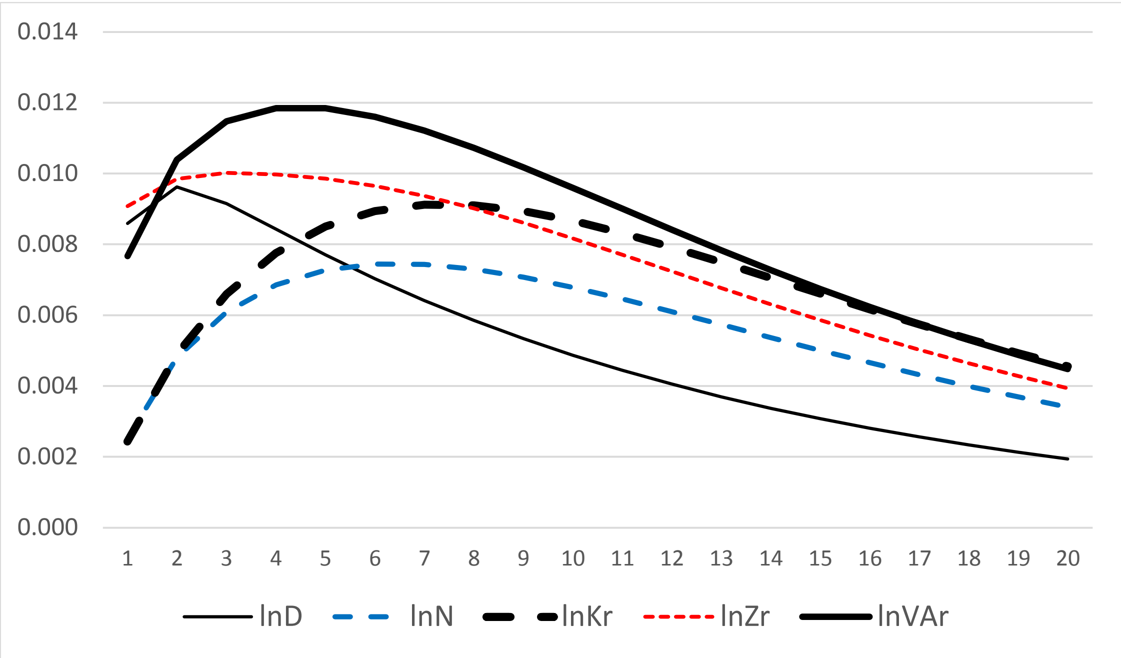

Note: Panel of private manufacturing firms with at least ten employees, excluding firms with extreme observations or missing values. Number of observations included in the estimations: 6149.

m) Larger firms (mean N ≥ 50)

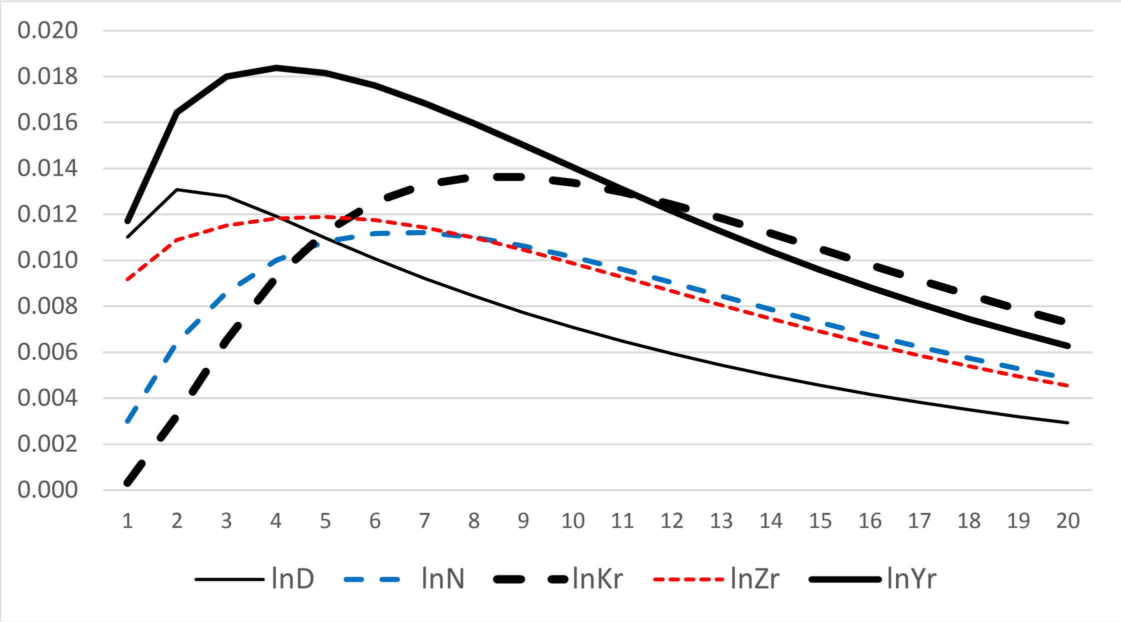

Note: Panel of private manufacturing firms with an average number of employees of at least 50, excluding firms with extreme observations or missing values. Number of observations included in the estimations: 4300.

n) Smaller firms (mean N < 50)

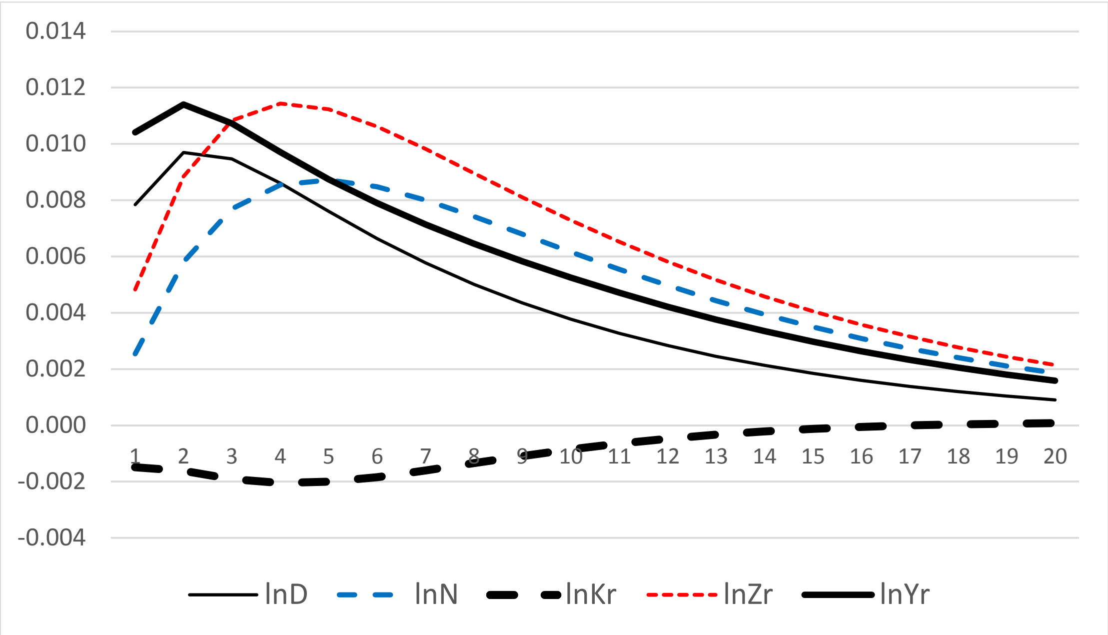

Note: Panel of private manufacturing firms with at least ten employees but with an average number of employees of less than 50, excluding firms with extreme observations or missing values. Number of observations included in the estimations: 3880.

o) Including linear industry trends but no time dummies

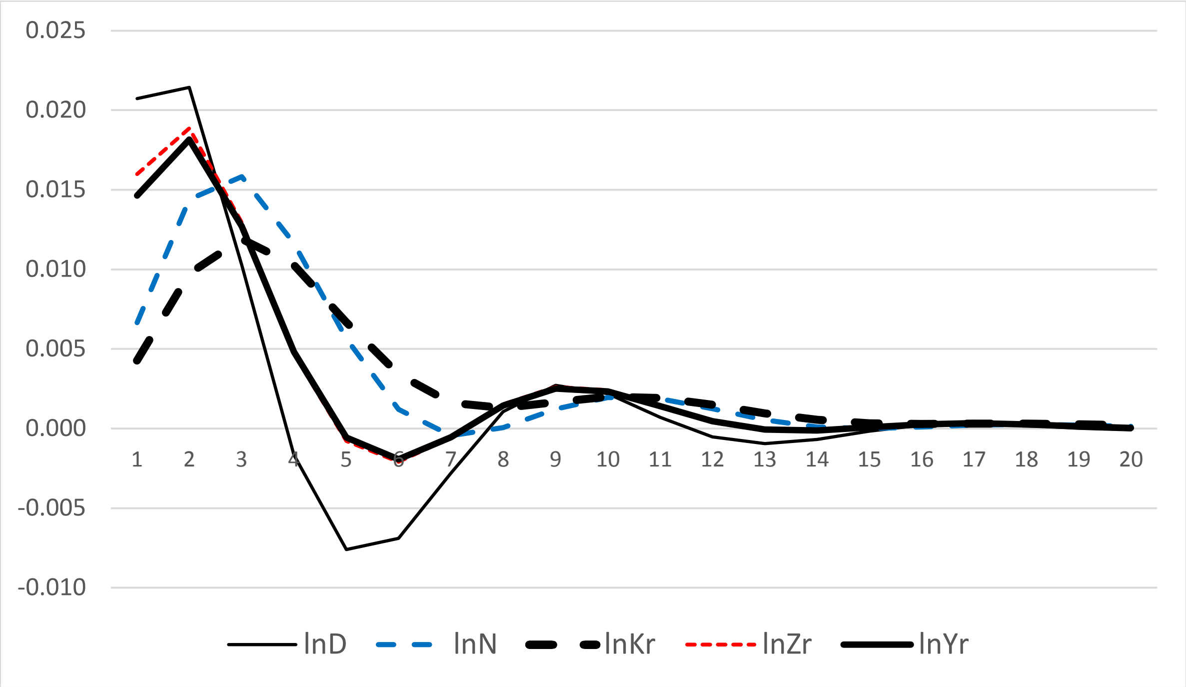

Note: Baseline panel of private manufacturing firms with at least ten employees, excluding firms with extreme observation or missing values s. Number of observations included in the estimations: 8180.

Table A3. Year of peak response for alternative specifications and samples

Specification		D	Y	Z	N	K
a	**Baseline estimation**	2	3	5	6	9
b	**Unbalanced panel**	2	3	4	6	8
c	**Excl. extreme obs. only**	2	3	3	5	7
d	**Including extreme obs.**	2	1	4	4	6
e	**Add construction**	2	3	4	6	9
f	**One-plant firms**	2	5	5	6	10
g	**3 lags**	2	4	5	6	8
h	**Industry trends**	2	2	4	5	8
i	**Industry x time**	1	2	1	4	8
j	**Domestic demand**	3	3	2	6	10
k	**Foreign demand**	2	1	4	6	9
l	**Value added**	2	4	3	6	7
m	**Larger firms**	2	4	5	7	9
n	**Smaller firms**	2	1	4	5	-
o	**No time dummies**	2	2	2	3	3

Construction of the steady state

The estimation involves a very large number of repeated simulations of the model. To save time in this process, we need to calculate the steady-state values analytically instead of searching for the steady state. Let variables with a bar denote steady-state values. Without loss of generality, we can choose units so that

$$\bar{Y} = \bar{K} = \bar{N} = \bar{\Omega} = \bar{P} = \bar{D} = \bar{u} = 1\,,\ \bar{I} = \delta_k\,,\ \bar{H} = \delta_n\,,\ \bar{q} = \bar{\gamma} = 0. \tag{48}$$

For given values of δ_ω and χ, we can calculate x in the steady state: $\bar{x} = 1 - \delta_\omega / \chi$, and then our normalizations imply a value for *A*:

$$A = \left(\alpha + (1-\alpha)\,\bar{x}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{1-\sigma}}. \tag{49}$$

It is convenient to view the steady-state inventory stock of finished goods as a parameter to be estimated. Denoting this value $\bar{Z}$, we obtain steady-state sales as $\bar{S} = 1 - \delta_z \bar{Z}$. In the steady state, we have

$$\kappa_1 \bar{Z} + \kappa_2 - \kappa_3 \bar{Z}^2 = \bar{S} \qquad\qquad \frac{\bar{S}/\eta}{1-\bar{v}-m} = \kappa_2 + \kappa_3 \bar{Z}^2$$

$$\frac{r_z\left(\bar{v}+m\right)+c_z\bar{\mu}}{1-\bar{v}-m} = \kappa_1 - 2\kappa_3\bar{Z} \qquad\qquad \text{where } r_z = 1+\beta\left(1-\delta_z\right).$$

Multiplying the last equation by $\bar{Z}$ and summing both sides of these equations, we can solve for the marginal cost of real value added in the steady state:

$$\bar{v} = \frac{1-m-1/\eta-\left(r_z m + c_z\right)\bar{Z}/\bar{S}}{1+r_z\bar{Z}/\bar{S}}.$$

For a given estimate of κ_1, we can then use the equations above to solve for κ_3 and κ_2.

We can also find the capital price and wage that are consistent with our normalizations:

$$P^K = \beta\bar{v}\alpha / \left(1-\beta\left(1-\delta_K\right)\right) \qquad \text{and} \qquad W = \bar{v}\left(1-\alpha\right)\bar{x}^{\frac{\sigma-1}{\sigma}}.$$

The first-order conditions for x and Ω yield

$$a = \frac{\left(1-\beta\left(1-\delta_\omega\right)\right)\left(1-\alpha\right)}{\beta\xi\chi\bar{x}^{1/\sigma}} + 1.$$

Furthermore, the normalizations imply that $\Phi_u = \bar{v}$ and $\bar{\phi} = \beta\bar{v}\left(a-1\right)\xi / \left(1-\beta\left(1-\delta_\omega\right)\right)$.

Figure A4. Effects on the target function of variations in one parameter at the time, keeping other parameters constant

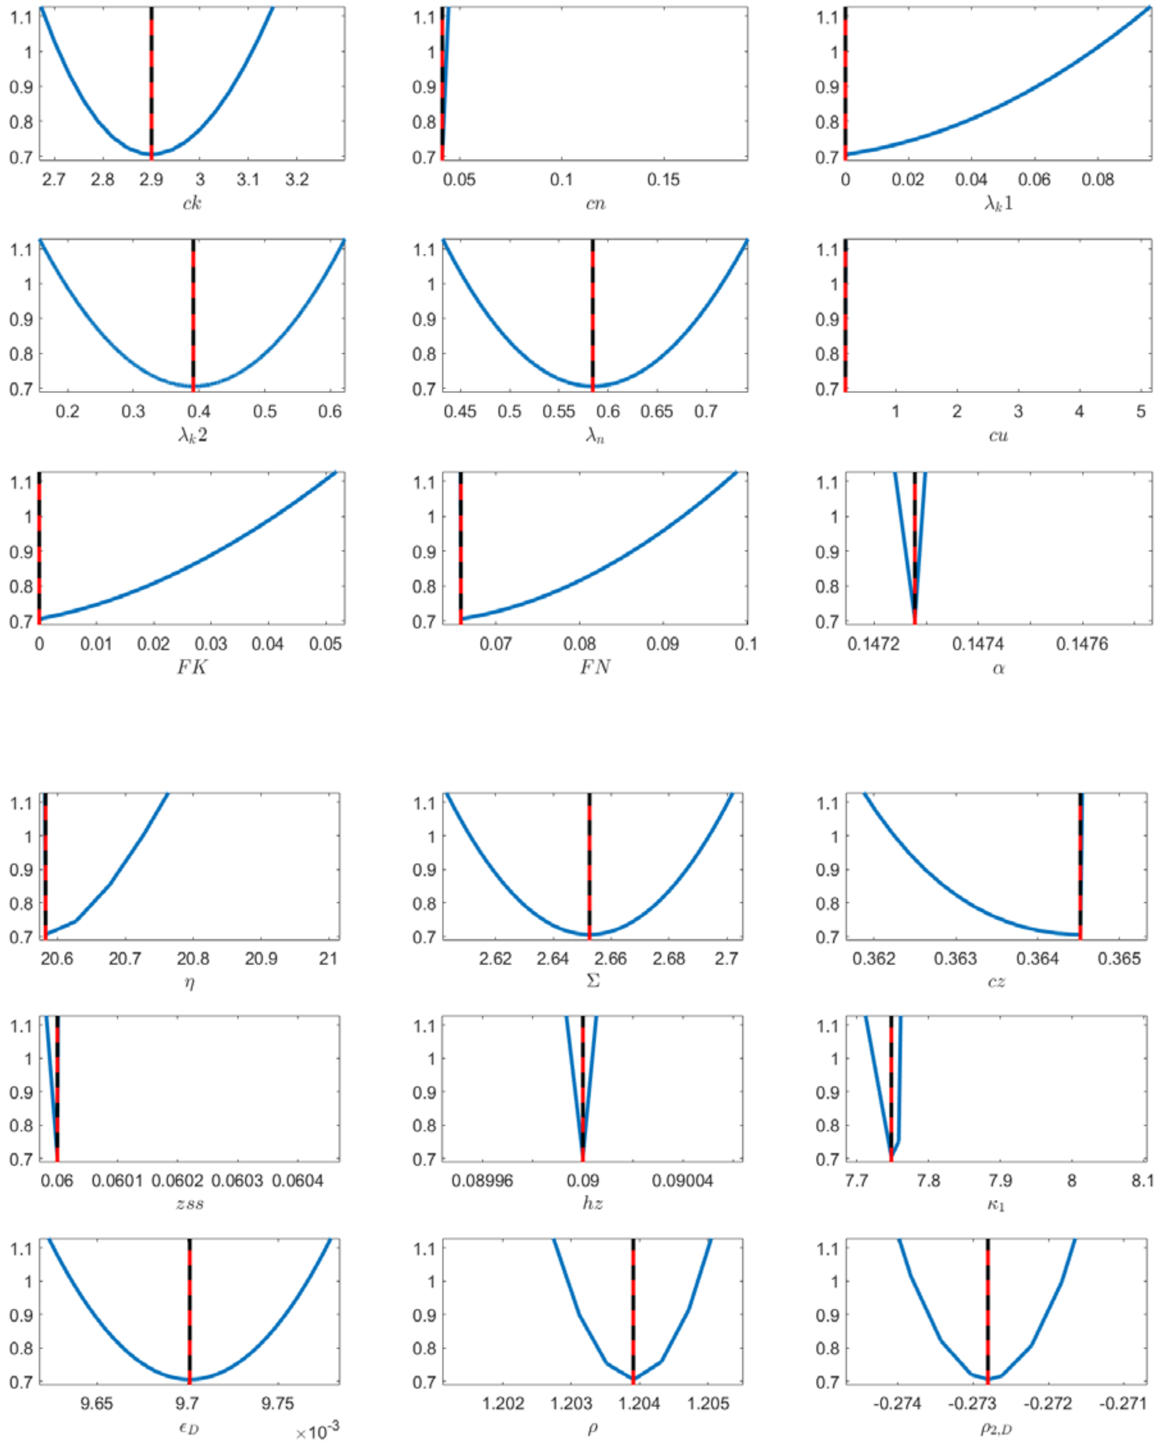

Note: In some cases, the parameters end up at corner solutions due to the prior constraints that we have imposed on the labor share, capital ratio, and inventory ratio (see Section 5.2). These plots do not tell us what happens to the target function when we change several parameters simultaneously.

Figure A5. Effects on the dynamics changing one parameter at the time

a) Baseline estimates of parameters

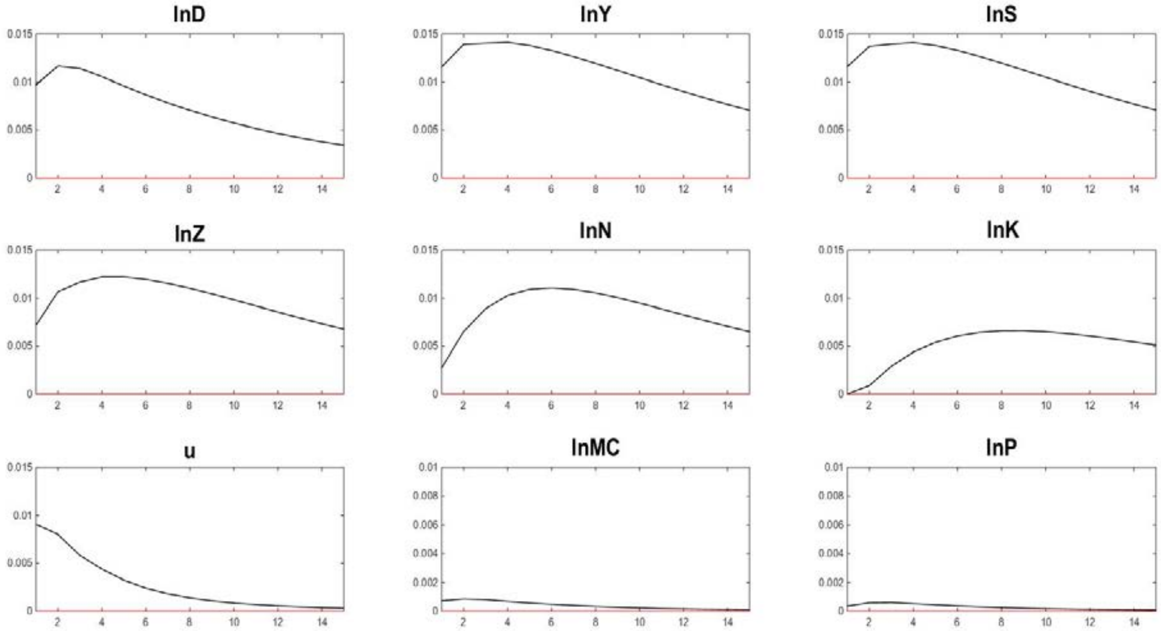

b) Setting $c_n = c_k = 0.1$

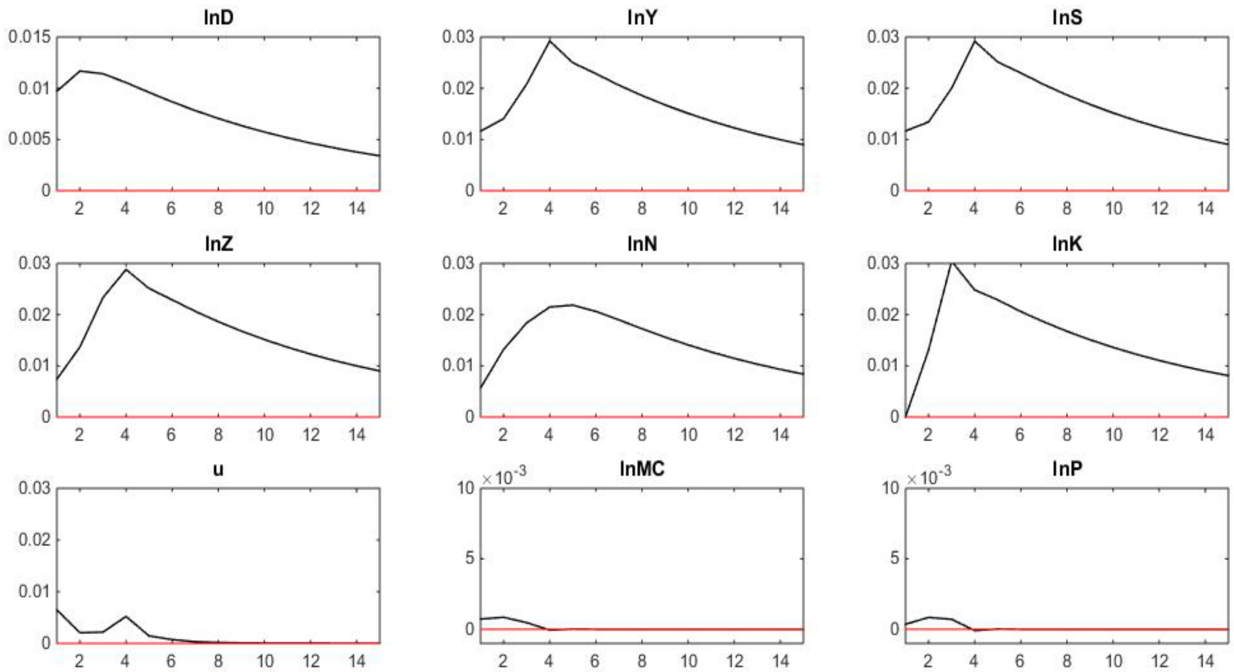

c) Setting $\lambda_k = \lambda_n = 1$

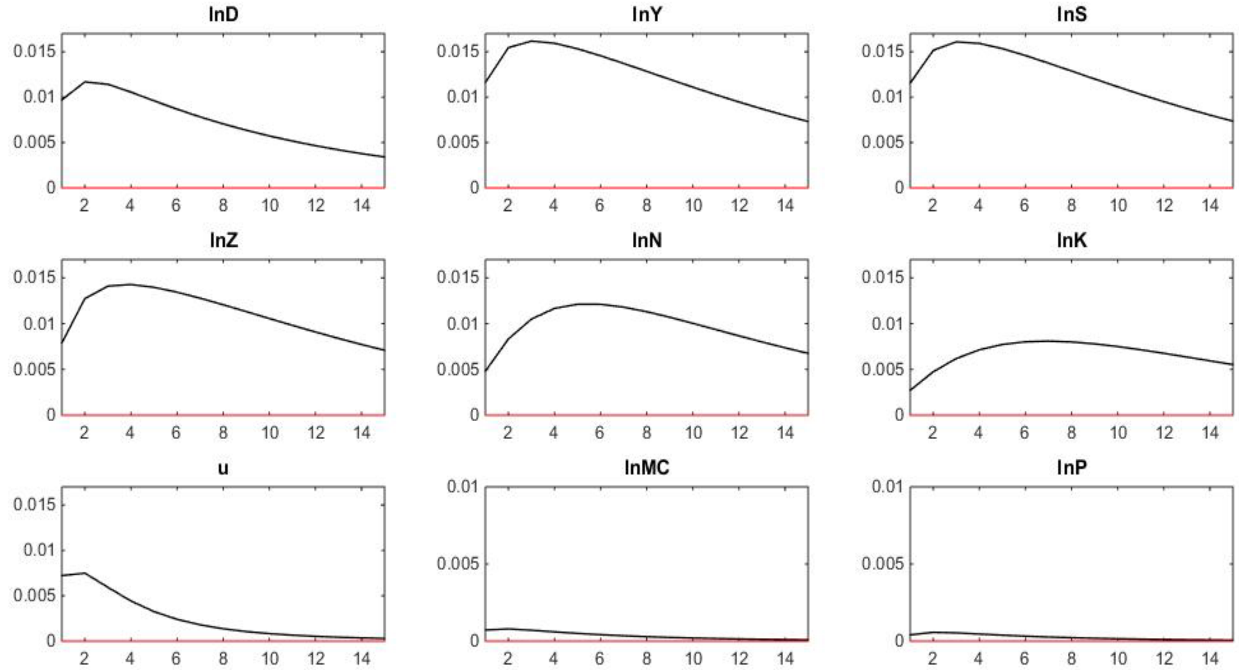

d) Setting $c_u = 2$

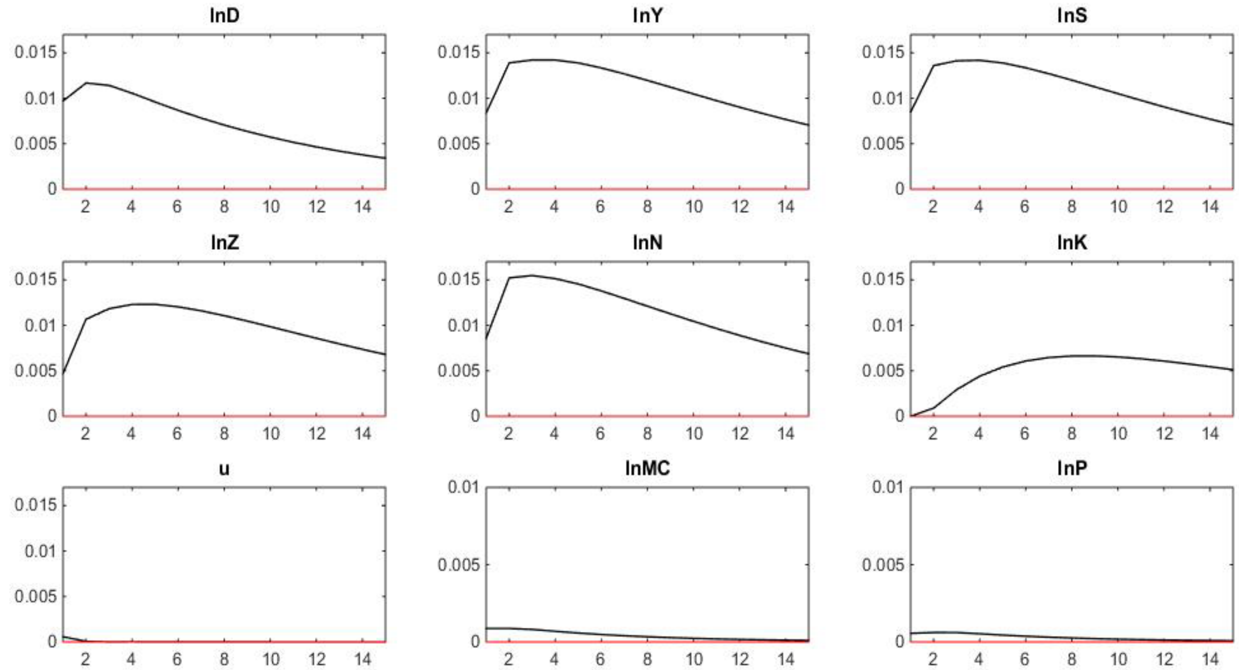

e) Setting $F_n = F_k = 0$

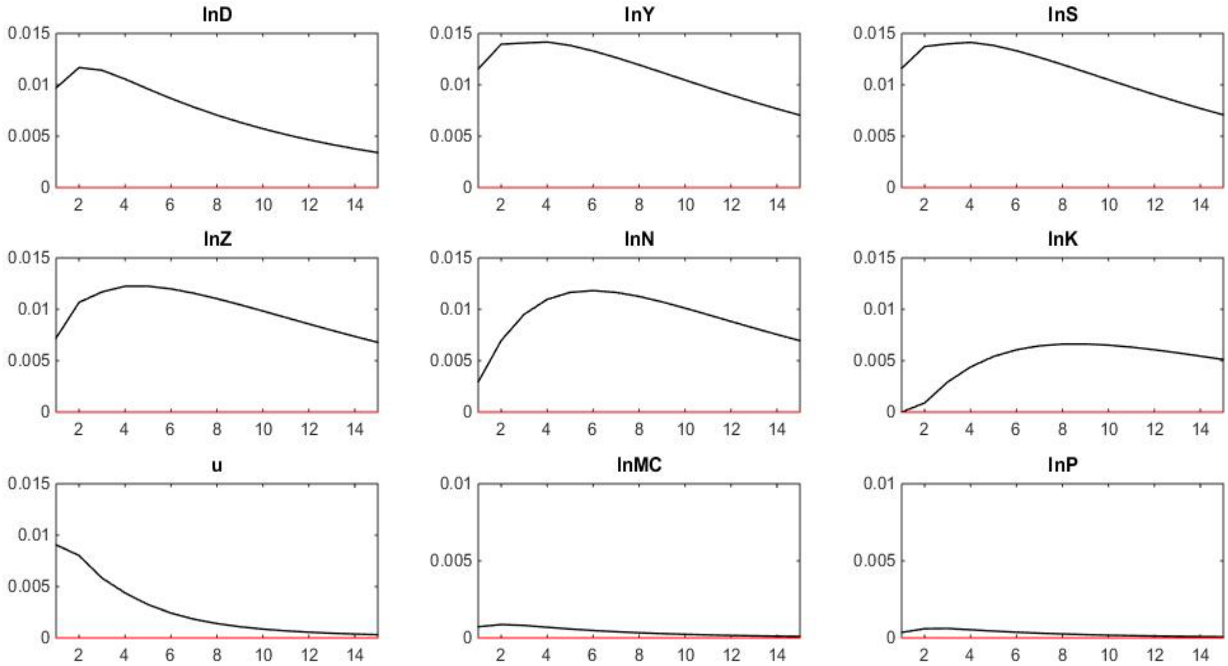

f) Setting $\sigma = 0.4$

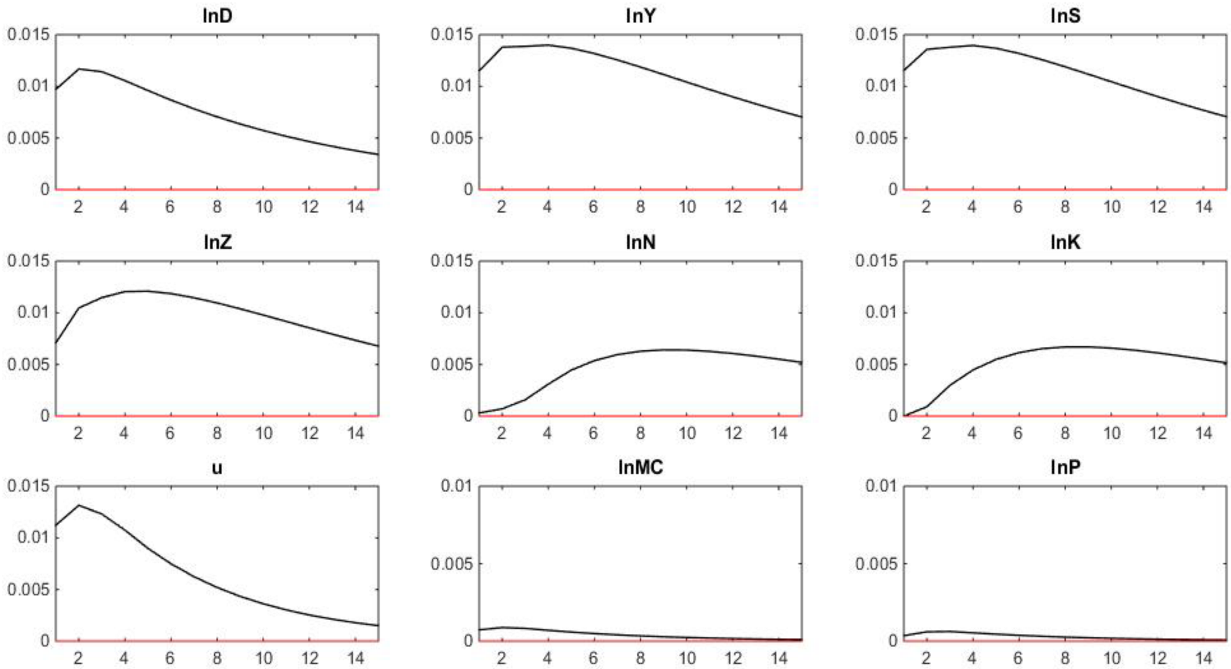

g) Setting $\delta_k = 0.07$

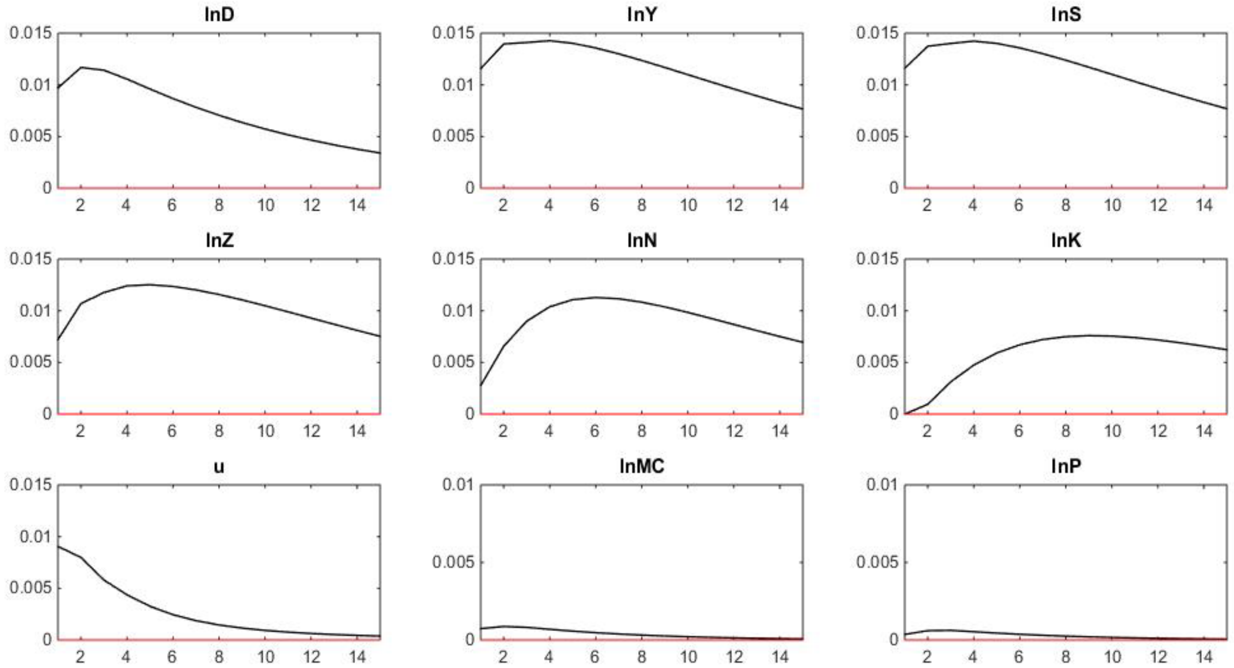

h) Setting $\delta_n = 0$

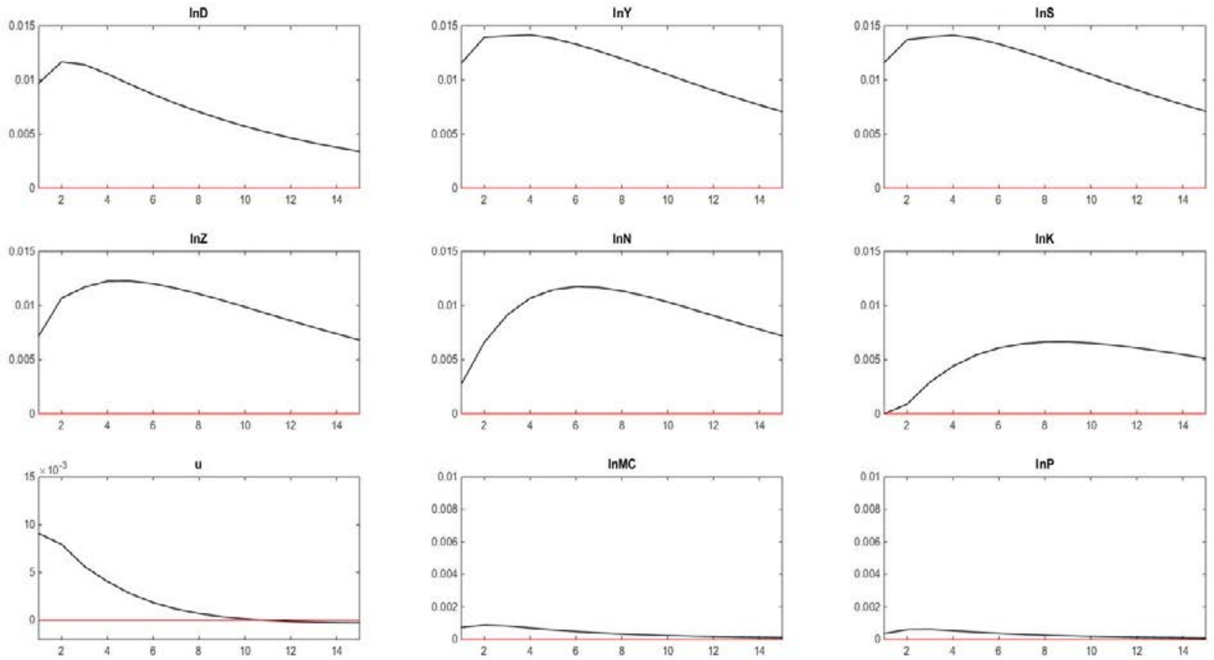

i) Setting $\delta_z = 0.20$

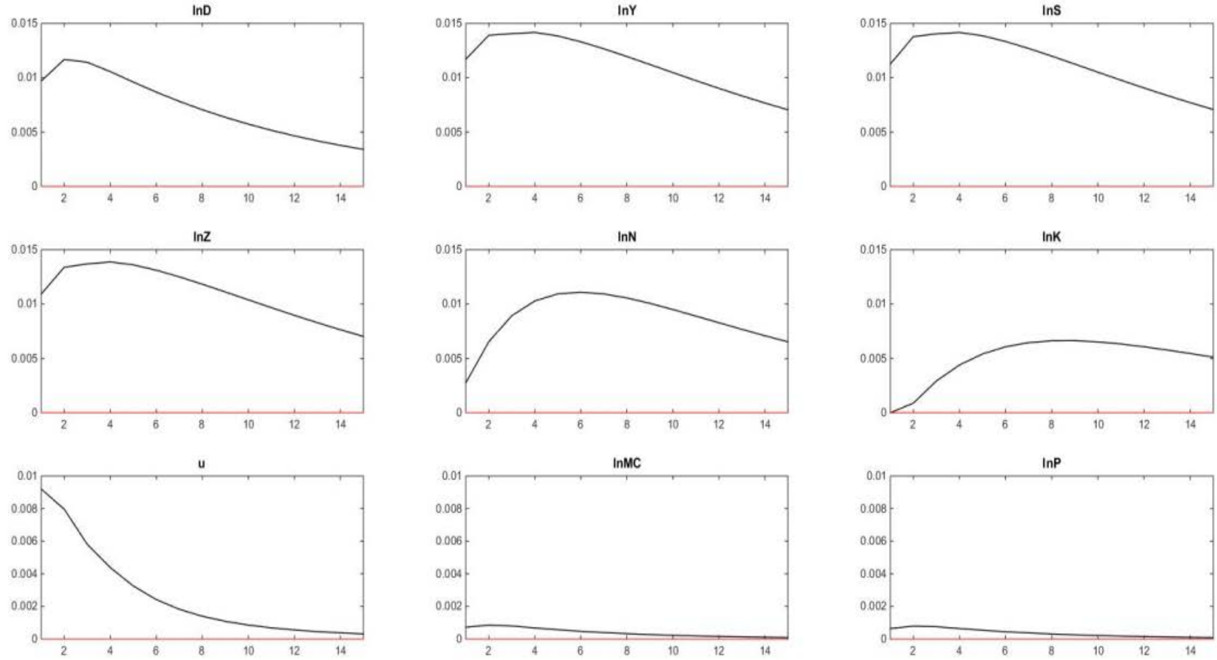

Note: We chose to hold the steady-state level of finished goods inventories constant in these experiments, so the parameters related to inventories $(\kappa_2$ and $\kappa_3)$ change when we change some parameter while keeping the other structural parameters constant. The reason for this is technical as we need to solve analytically for the steady state.