Demand Shocks and Prices—Micro Evidence and Macro Implications

Christian Philip Hoeck
(University of Copenhagen and Danmarks Nationalbank)

Tobias Renkin
(Danmarks Nationalbank)

First draft: November 8, 2024
This version: March 20, 2025

(PRELIMINARY DRAFT)

Abstract

We estimate the response of domestic prices and total output of Danish manufacturing firms to persistent firm-level demand shocks that result from heterogeneity in firms' exposure to different export destinations. Our results suggest supply curves at the firm level are steep—a demand shock that increases output by 1% raises prices by 0.5%. We then augment the supply side of a New Keynesian model with firm-level demand shocks, and identify key parameters from matching the response of firms to those shocks in the model to our estimates. We show that in a model that fits firm behavior in the cross-section, the slope of firm-level supply curves contributes meaningfully to the slope of the Phillips curve. Nevertheless, with realistic cyclicality of real wages the Phillips curve is rather flat. Our preferred estimate of the overall Phillips curve slope is 0.025, which almost entirely reflects the aggregation of firms' supply curves.

Hoeck: Danmarks Nationalbank, Langelinie Allé 47, 2100 København Ø, Denmark, christian.hoeck@econ.ku.dk;
Renkin: Danmarks Nationalbank, Langelinie Allé 47, 2100 København Ø, Denmark, tobias.renkin@gmail.com;

The views expressed in this paper are those of the authors, and do not necessarily reflect the official views of Danmarks Nationalbank.

1 Introduction

The Phillips curve is an aggregate supply curve that determines the relationship between aggregate prices and the output gap. Its slope is crucial for the ability of macroeconomic policy to trade-off inflation and output in the short run. The low correlation of inflation with measures of the output gap over the two decades between 2000–2020 has sparked a lively debate over the flattening of the Phillips curve. This debate has highlighted the considerable uncertainty about the slope of the Phillips curve due to problems of weak identification in the traditional approach to estimation (Mavroeidis et al., 2014) and unaddressed simultaneity of aggregate supply and demand (McLeay and Tenreyro, 2019).

In this paper, we take the role of the Phillips curve as an aggregate supply curve seriously. Our key contribution is to explicitly connect the shape of the aggregate Phillips curve to estimated firm-level supply curves. We estimate the response of Danish manufacturing firms' prices and output to plausibly exogenous cross-sectional demand shocks and back out the average slope of supply curves. We then augment the production side of a New Keynesian model with similar firm-level demand shocks that mirror our empirical setting. We show that our estimates can identify a part of the Phillips curve slope that is directly related to the slope of firms' supply curves—the "capacity pressure" channel. We fit the response of firms' relative prices and output to firm-level shocks in the model to our empirical estimates. This identifies key parameters of the model, and allows us to make quantitative predictions about the aggregate Phillips curve.

Our analysis yields several new takeaways. First, supply curves at the firm level are somewhat steeper than implied by common calibrations of New Keynesian models. Second, this carries over to the aggregate—in a model that matches firm-level behavior, the capacity pressure channel contributes 0.032pp to the slope of the Phillips curve, and taking into account effects on inflation expectations 0.1pp to the Phillips multiplier after a moderately persistent monetary policy shock. These values are about 1.5 times the contribution of the capacity pressure channel in a textbook calibration of the New Keynesian model. However, with a realistic wage Phillips curve that matches the slightly negative co-movement of real wages and output after monetary policy shocks in Denmark, the overall Phillips curve is still rather flat.

Identifying a supply curve requires a demand shock. Our identification strategy relies on shift-share demand shocks to exporting firms. We combine heterogeneity in export destinations in the cross-section of Danish firms with fluctuations in aggregate import demand in those destinations over time.[1] We use local projections to estimate

[1]Similar export demand shocks have been used in Hummels et al. (2014) to estimate the effects of offshoring on Danish wages, and in Garin and Silvério (2024) to estimate the response of Portugese wages to firm-level demand shocks.

the dynamic response of prices and output to this demand shock, and IV local projections to directly estimate the slope of firms' supply curve. Since the demand shock we construct varies in the cross-section of firms, we are able to control for aggregate supply shocks and inflation expectations using time-sector fixed effects. Our remaining identification assumption is that the variation in exposure to different export destinations is orthogonal to firm-level supply shocks. By estimating the response to a cross-sectional demand shock rather than an aggregate shock, our approach thus sidesteps the most important identification issues that plague the estimation of the Phillips curve slope from aggregate data.

In our empirical work, all aggregate variation is absorbed in time-sector fixed effects. This includes movements in prices that arise from equilibrium interaction with other firms, changes in inflation expectations, and the dynamics of real wages. To make predictions about macroeconomic relationships, we combine our firm-level estimates with a model that fills in the aggregate dynamics absorbed in our empirical approach. We augment the production side of a New Keynesian model with persistent firm-level demand shocks that mirror our empirical setting. We show that our empirical estimates can be used to identify the importance of the capacity pressure channel—the part of the Phillips curve slope that is driven by the slope of firms' supply curve. In contrast, our estimates, and others that use cross-sectional variation in demand, are not informative about what we call the "wage pressure" channel—the part of the price Phillips curve slope that is driven by pass-through of wage dynamics that result from the wage Phillips curve.

We fit the response of firms' relative prices and output to firm-level demand shocks in the model to the response we estimate from the data. The model does very well at fitting firm behavior. In the fitted model, the slope of firms' supply curve is steeper than implied by common calibrations of the New Keynesian model. In the aggregate, this results in an important contribution of capacity pressure to the slope of the Phillips curve. In particular, the capacity pressure channel adds about 0.032pp to the slope of the Phillips curve and taking into account effects on inflation expectations 0.1pp to the Phillips multiplier after a moderately persistent monetary policy shock. With additional estimates of the response of wages to monetary policy shocks, we can make predictions about the overall slope of the Phillips curve as well. Since wages in Denmark are slightly counter-cyclical, we arrive at a rather flat overall Phillips curve with a slope of 0.025. The positive slope is entirely the result of the capacity pressure channel, while the contribution of the wage pressure channel is slightly negative.

Our work is related to the large literature estimating the slope of the Phillips curve in different settings. This estimation is subject to two important identification concerns. First, any shock to output might affect inflation directly through the Phillips curve slope, and indirectly through inflation expectations. Second, aggregate supply and de-

mand shocks move prices in different directions, and unless the output gap is observed without error, it is necessary to use aggregate demand shocks for identification. Most of the classical literature on Phillips curve estimation focuses on identifying the slope of the Phillips curve separately from the effect of expectations, using aggregate data and rational expectation assumptions that motivate (typically internal) instruments to deal with the forward-looking component in the error term. This literature is surveyed in detail in Mavroeidis et al. (2014). They conclude that the traditional approach is subject to severe identification issues—in particular, aggregate instruments are too weak to reliably estimate the slope of the Phillips curve from time series data.

A second identification issue of simultaneity between aggregate supply and aggregate demand has received more attention recently and is presented succinctly in McLeay and Tenreyro (2019). If monetary policy is conducted systematically to limit variation in output after aggregate demand shocks, then co-movements between output and inflation will result mostly from supply shocks and will not be informative about the slope of the Phillips curve. One solution to this identification issue is to use deviations from monetary policy rules to identify the trade-off between output and inflation.[2] Barnichon and Mesters (2020) and Barnichon and Mesters (2021) follow this approach and obtain point estimates that suggest the Phillips curve has flattened after 1990, but still has a positive slope of 0.12 to 0.18. However, deviations of monetary policy from policy rules are infrequent and small, and the resulting shocks are weak instruments that yield imprecise estimates without additional structural assumptions.

Given the substantial unresolved identification issues in the estimation of the Phillips curve from aggregate data, a nascent literature has started to develop alternative approaches using panel data of regional aggregates or at the firm level. We add to this literature. Our primary contribution is to use credible and strong firm-level demand shocks to estimate the slope of firms' supply curves as a building block that we use together with a structural model to make predictions about the aggregate Phillips curve. This approach has two advantages. First, the identification assumptions required to estimate firms' supply curves are weaker than those required for aggregate or regional Phillips curve estimation because we can absorb all aggregate supply shocks in sector-time fixed effects. Second, the shocks we use for identification are strong instruments compared to those used in direct estimation of aggregate Phillips curves. A bottom-up approach such as ours also has disadvantages, and our second main contribution is to clarify these. While aggregate approaches ideally identify the overall slope of the Phillips curve, we show that estimates using cross-sectional demand shocks are only able to credibly identify the capacity pressure channel and are not informative about the slope of the wage Phillips curve. However, given the substantial identification issues of aggregate approaches, we view our results as an important step forward, even if we ultimately identify a narrower

[2]There is a larger literature that estimates reduced form effects of monetary policy on inflation, output, and other variables using deviations from monetary policy rules. This is very similar in spirit, but usually doesn't explicitly back out the slope of the Phillips curve. Notable recent examples include Gertler and Karadi (2015), Nakamura and Steinsson (2018), Jarocinski and Karadi (2020).

parameter.

The paper most closely related to ours is Gagliardone et al. (2024), who use microdata on Belgian manufacturing firms to estimate the marginal cost formulation of the Phillips curve, i.e. pass-through of current and expected future marginal cost into prices. This pins down important Phillips curve parameters, but is not informative about the relationship between prices and *output*. To estimate this relationship, Gagliardone et al. construct sectoral demand shocks from aggregate monetary policy shocks interacted with estimated sectoral sensitivities to those shocks. Since they build on aggregate monetary policy shocks, their estimates are subject to the same concerns about weak identification highlighted in the literature estimating the Phillips curve using aggregate data. Indeed, we replicate their approach using Danish data in Appendix F, and show that once we fully account for the estimation uncertainty associated with their multi-step estimation approach—sectoral sensitivities are estimated in a first-step—their instruments appear critically weak, and weak-instrument robust methods yield imprecisely estimated, insignificant coefficients. Our main contribution is therefore to identify firms' supply curves based on strong cross-sectional demand shocks.

Our paper is also closely related to the literature estimating regional Phillips curves, most notably McLeay and Tenreyro (2019) and Hazell et al. (2022). The main difference between their approach and ours is that they estimate the slope of the overall (regional) Phillips curve, but require stronger identification assumptions than our firm-level estimation. McLeay and Tenreyro (2019) use city-level CPI and unemployment data and use fixed effects to control for variation in inflation expectations and aggregate supply shocks at the national level. They estimate a relatively steep unemployment-based hybrid Phillips curve with a Phillips multiplier of -0.379. However, their estimate might still be biased if local unemployment is partially driven by local supply shocks, and might not be informative about the shape of the national Phillips curve if prices of tradeable goods are less responsive to local aggregate demand than to national aggregate demand.

Hazell et al. (2022) address these identification issues. They use unemployment and prices of non-tradeables (i.e., mostly services) at the state level to identify the slope of regional Phillips curves. To address the possibility that local unemployment is partially driven by local supply shocks to non-tradeable production, they construct a shift-share instrument using variation in the local exposure to national shocks to tradeable sectors. They estimate an unemployment-based Phillips curve that is flat, with a slope coefficient of -0.006. Hazell et. al. require several relatively strong identification assumptions. For example, they assume that national tradeable shocks do not spill over into local tradeable supply, which could be violated, for example, due to inflows of workers or capital. Moreover, their estimates are "scaled" correctly only if aggregate state unemployment that results from shocks to tradeable

industries is a good measure of slack for non-tradeable industries. This would be violated if labor mobility between the tradeable and non-tradeable sectors is limited. Our firm-level approach requires only weak assumptions, but only identifies the capacity-pressure channel of the slope of the Phillips curve, while the estimates in Hazell et. al. also identify the part of the slope coming from the wage Phillips curve.

The paper proceeds as follows. In section 2, we introduce the datasets we use throughout the paper. Section 3 explains our identification strategy using a simple static model and discusses the construction of demand shocks and the equations we estimate. Section 4 discusses the firm-level results. In section 5 maps our firm-level estimates to the aggregate Phillips curve. We conclude in section 6.

2 Data

Our work is based on register data covering production, sales, and prices of Danish manufacturing firms at the product and destination level. We combine these firm-level datasets with macroeconomic data on countries' product-level imports and exports to construct firm-level shift-share demand shocks. While much of the firm-level microdata we use is available at a quarterly or monthly frequency, trade data covering a large enough sample of countries over a sufficiently long period of time is only available at the annual level. Consequently, most of our empirical analysis is at the annual frequency. Our analysis covers the 2001–2021 period, since some firm-level datasets are not available before that.

Production and sales microdata. We use data on sales, production and exports of Danish manufacturing firms collected from various administrative sources. Data on total global sales and production at the product level comes from large scale administrative survey (VARS) that is used to produce the Danish contribution to the Eurostat PRODCOM database. The survey covers all manufacturing firms with more than 10 employees and provides quarterly sales and production quantities at the level of 8-digit Combined Nomenclature (CN) product categories. Data on goods export and import values and quantities is based on administrative survey and customs data (UHDM). This data is collected for all exporters above a small yearly minimum export cutoff and provides monthly export sales and quantities at the level of 8-digit CN product categories. We complement this data with basic firm information from annual balance sheets available in the Danish business register (FIRM) and the Danish accounting statistics (FIRE). The variables we use from these data sets are available for the universe of Danish firms. Finally, we use survey data on self-reported capacity utilization from the Danish Business Sentiment survey (Konjunkturbarometer). This

dataset covers roughly 450 manufacturing firms.

Our main measure of firm output is an output index constructed from VARS micro data. In VARS, firms report total quarterly sales and produced quantities at the level of the 8-digit combined nomenclature. We aggregate this to yearly data at the level of 6-digit Harmonized System codes[3]. We then construct a firm-level output index as:

$$Q_{i,t} = \sum_{j \in J} \eta_j Q_{i,j,t} \tag{1}$$

, where the weights η_j correspond to the average unit value of product category j over the sample period. We keep the unit value weights fixed over the sample period instead of constructing a chained index, since lagged weights are frequently missing due to fluctuations in firms' product portfolio.

Producer price index microdata Our price data comes from the Danish Producer Price Index (PPI) survey. The PPI is based on a monthly survey in which firms report prices for a persistent selection of their product portfolio. In an average month, the data covers about 3,500 price quotes from 500 firms. Products are classified using 8-digit CN codes. Firms mainly report domestic prices. Some firms also provide export prices, but the survey does not contain information on the export destination. The reported prices are transaction prices in Danish kroner, including temporary sales and discounts. The survey is designed to allow adjustments for quality changes and product substitutions. When quality changes or product substitutions occur, firms report both a lagged and current price for the new product, based on which a quality-adjusted price change can be computed. The dataset is very balanced, with few gaps in price series. We perform quality adjustments and winsorize price changes at ± 1 log points in the monthly data. We then transform the dataset to annual frequency by keeping the price in the first month of each year. The Danish PPI survey has been previously used in Dedola et al. (2019), who provide important price-setting moments and show that the data is comparable to other European producer price datasets.

Macro data on imports and exports Finally, we use macroeconomic data on imports and exports during the 2000–2021 period from the UN Comtrade database. Comtrade covers trade flows between a source and a destination country at the product level. Our baseline analysis uses flows at the 4-digit Harmonized System (HS) code level. We construct country-product level import growth rates that exclude imports from Denmark. These imports will serve as shocks in our shift-share demand instrument, and we leave out imports from Denmark in the construction to rule out a source of possible reverse causality.

[3]We use the 1996 definition of the Harmonized System, and convert the concurrent HS codes to their 1996 counterpart using conversion tables provided by United Nations Statistics Division (2022)

Sample description Our baseline estimation sample covers manufacturing firms that report domestic prices in the PPI and output in VARS. Since all manufacturing firms with more than 10 employees are covered in VARS, the binding constraint is usually participation in the PPI survey. We impose two additional constraints on the sample. First, we impose a balance requirement that firms have more than 20 employees for more than 5 consecutive years during the sample period. This excludes small firms that would otherwise go in and out of the sample as they cross the VARS coverage threshold of 10 employees. Second, we require firms to have an export share of at least 5% of their total sales in the period they are hit by a given shock.

This results in a sample of 855 firms over the 2001–2021 period. The average firm in our sample has 234 employees and sales of 96 million euros. This is small by global standards, as Danish manufacturing is dominated by small and medium-sized enterprises. However, the sample does include some very large firms, and firm size measures are very skewed, with the median substantially below the mean. Most firms export a large share of their production, and average goods exports are about half of average sales. On average, firms export 17 different HS products categories to 27 different countries. Finally, most firms report several prices in the PPI survey, and the average firm reports 5 prices.

	Mean	Median	10th percentile	90th percentile
Sales (Mio EUR)	96.83	26.97	6.73	144.44
Employment (FTE)	234.01	97.53	29.16	428.00
Assets (Mio EUR)	106.98	19.04	4.28	132.70
Goods exports (Mio EUR)	45.25	13.46	2.26	80.41
Exports (Mio EUR)	58.45	14.87	2.51	93.53
Imports (Mio EUR)	26.03	5.89	0.55	43.08
Export destinations	26.70	23.00	7.00	52.00
Exported products (4-digit HS)	18.61	12.00	3.00	41.00
Prices reported in the PPI	4.87	3.00	1.00	9.00
Firms				855
Observations				11,898

Table 1: Descriptive statistics for our main estimation sample

3 Estimating Firm-level Supply Curves Using Export Demand

The aim of this paper is to explicitly link the aggregate Phillips curve to the slope of supply curves at the firm level. This section discusses our estimation of supply curves in the microdata. Since firms' output and prices are determined simultaneously by supply and demand, we require a demand shifter to identify supply curves. We use demand shifters that arise from heterogeneity in firms' export exposure to aggregate fluctuations in different destination countries. Firms persistently export their products to different destination countries, and aggregate fluctuations in these countries are not perfectly synchronized. This leads to variation in firm-level demand that we exploit. We introduce a simple static framework describing the price-setting problem of a single firm to motivate the construction of the demand shifter we use in our empirical work and rigorously illustrate our identification assumptions. We keep this framework as simple as possible for illustrative purposes and show that our strategy is robust in more complicated settings in extensions in Appendix C. In section 5, we map our empirical estimates to a New Keynesian model with sticky prices.

3.1 A Simple Analytical Framework

The firm we consider produces output in a home market h and sells it in K markets (including the home market) indexed by k. The firm uses labor in the home market h as its only freely adjustable input and produces using a Cobb-Douglas production technology. We assume that the firm's capital stock is fixed in the short run and normalize it to 1:

$$Q_i = (A_h V_i L_i)^{1-\alpha}.$$

Since the capital stock is fixed and $\alpha \in (0,1)$, the firm operates with decreasing returns to scale and its short-run supply curve is upward sloping. $(A_h V_i)^{1-\alpha}$ is the firm's total factor productivity, which we decompose into an idiosyncratic component V_i and a component A_h that is shared across all producers in market h.

The firm produces a perfectly tradeable product, and sets one price in all markets. We assume prices are fully flexible for now. The firm faces constant elasticity market demand curves with the elasticity of demand σ in each market k, i.e., $Q_{i,k} = (P_i/\overline{P}_k)^{-\sigma}\Lambda_{i,k}Z_k$. $\overline{P}_k$ is the price index of the firms' competitors in market k. Z_k and $\Lambda_{i,k}$ are two demand shifters. $\Lambda_{i,k}$ is a taste shifter that we think of as permanent and that gives rise to differential exposure of firms to different markets. Z_k is an aggregate shock that affects demand for all firms in market k. We can think of Z_k as aggregate income in market k. Profit maximization results in a price-setting policy of constant markups over

marginal cost. Taking logs and differentiating, we get the following supply-demand system that determines changes in prices and quantities (we denote logarithms with lowercase letters):

$$\Delta p_i = \delta \Delta q_i + \Delta w_h - \Delta a_h - \Delta v_i \qquad \text{Inverse supply} \qquad (2)$$

$$\Delta q_i = -\sigma \Delta p_i + \sigma \sum_{k=1}^{K} \gamma_{i,k} \Delta \overline{p}_k + \sum_{k=1}^{K} \gamma_{i,k} \Delta z_k \qquad \text{Demand} \qquad (3)$$

We denote the slope of the supply curve with $\delta = \alpha/(1-\alpha)$. The variable $\gamma_{i,k} = Q_{i,k}/Q_i$ measures the share of the firm's output sold in market k. If initial price levels and aggregate shocks are identical, then $\gamma_{i,k} = \Lambda_{i,k}/\sum_k \Lambda_{i,k}$. We denote output-weighted averages with a tilde, e.g. $\sum_{k=1}^{K} \gamma_{i,k} \Delta z_k = \Delta \tilde{z}_i$. The slope of the flexible price inverse supply curve (2) in this setting is determined by the production function exponent α and only depends on the total output of a firm, not on the part sold in any specific market, because marginal cost is shared across production for all markets. Combining the inverse supply and demand functions, the reduced form relationship between the firm's price and the exogenous variables is given by:

$$\Delta p_i = \frac{\delta}{1+\delta\sigma} \Delta \tilde{z}_i + \frac{\delta\sigma}{1+\delta\sigma} \Delta \tilde{\overline{p}}_i + \frac{1}{1+\delta\sigma} (\Delta w_h - \Delta a_h - \Delta v_i).$$

The price of a firm depends on total output and thus on demand shifters and competitor prices in all markets weighted by their share in the firm's output. It also depends on the home country supply shifters Δw_h and Δa_h. An upward shift in total demand leads to an increase in the firm's output and marginal cost and consequently its price—this effect is reflected by the numerator of the coefficients. The higher price lowers the firms' output which partially offsets the cost increase through a movement along the demand curve—this is represented by the denominator of the coefficients.

In our empirical analysis we estimate how domestic prices respond to demand shifters that resemble $\Delta \tilde{z}_i$. We will absorb the home market supply shifters Δw_h and Δa_h and other factors that don't vary between firms in sector-time fixed effects. Our critical identification assumption is thus that the weighted demand shifter $\Delta \tilde{z}_i$ is independent of the remaining idiosyncratic supply shock Δv_i. Our demand shifter $\tilde{z}_i$ is a shift-share or "Bartik" instrument, and it is well understood that such instruments are valid if the shares—i.e. export exposure to different markets $\gamma_{i,k}$—are orthogonal to idiosyncratic supply shocks (see e.g. Goldsmith-Pinkham et al. (2020)). Under this condition, a reduced

form regression of price changes Δp_i on the demand shifter $\Delta \tilde{z}_i$ identifies:

$$\beta_{\tilde{z}} = \frac{\delta}{1+\delta\sigma}\left(1+\sigma\frac{\text{cov}(\Delta\tilde{z}_i, \Delta\tilde{\overline{p}}_i)}{\text{var}(\Delta\tilde{z}_i)}\right).$$

The coefficient is a mix of the slope of the inverse supply curve δ and the slope of the demand curve σ. We explicitly allow for the fact that the demand shifter might affect demand directly and indirectly through a correlation of aggregate conditions with competitor prices in destination markets. We can identify the slope of the inverse supply curve from an IV regression of prices on output (again, absorbing the home market supply shifters in fixed effects), using the weighted supply shifter as an instrument for output:

$$\beta_{\tilde{y}}^{IV} = \frac{\text{cov}(\Delta p_i, \Delta\tilde{z}_i)}{\text{cov}(\Delta q_i, \Delta\tilde{z}_i)} = \delta$$

This framework provides some key takeaways. First, our identification strategy relies on absorbing shared home market supply shocks such as local factor cost and productivity in fixed effects, and orthogonality of export shares with firm-specific supply shocks. Second, when estimating the response of domestic prices to demand shocks, it doesn't matter for identification of the supply curve where the shock originates—variation in export destinations provides us with a setting in which we can control for home market supply shocks while still retaining variation in demand, but the firm's response to a foreign demand shock is identical to its response to a domestic shock as long as both are weighted correctly by the firms' exposure $\gamma_{i,k}$. This result comes from the fact that marginal cost is the same for sales in all destination markets. Third, while we like to think of z_k as an aggregate demand shifter in the destination market, it does not matter in practice if fluctuations in foreign demand arise from aggregate demand or supply shocks in the destination. Theory suggests that the covariance with destination prices $\text{cov}(\Delta p_i, \Delta\tilde{z}_i)$ would be positive if z_k mostly captures aggregate demand shocks, and negative if it captures mostly aggregate supply shocks. In both cases, $\tilde{z}_i$ is a valid instrument, and an IV estimate recovers the slope of the firms' supply curve.

In Appendix C, we discuss extensions of this framework. We show that our strategy still identifies firms' supply curve in a setting with strategic complementarity and pricing-to-market. Moreover, in section 5, we consider firms' pricing decision with sticky prices to tie our results directly to the Phillips curve of a New Keynesian model.

3.2 Estimation

The key takeaway from the previous section is that we can use variation in firms' exposure to variation in demand in different export destinations to estimate the slope of firms' supply curve, if we construct a proper demand shifter and control for supply shocks in the home market using fixed effects.

Construction of demand shifters. We construct shift-share demand shifters in line with section 3.1—we use annual import growth as a proxy for demand fluctuations in destination countries and firm-level export shares in sales to measure the exposure of Danish firms to different destinations. We use import growth as a measure of fluctuations for two reasons. First, it is more directly related to the demand for Danish products than, for example, GDP growth or measures of the output gap. Second, both country-level import data and firm-level Danish export data are available at the level of disaggregated product categories. This allows us to construct shift-share instruments using variation in both firms' export destinations and the product composition of their production. Our baseline demand shifters are constructed using imports and export shares at the level of 4-digit Harmonized System product categories.[4]

For each country k and product j, we calculate annual import growth rates $\Delta im_{k,j,t}$. We exclude imports from Denmark from this calculation to rule out a possible source of reverse causality. We then calculate the share of exports of product j to country k in the total sales of each firm (including domestic sales) in the previous year. Our shift-share demand shifter is then calculated as:

$$\Delta \tilde{z}_{i,t} = \sum_{k \in K} \sum_{j \in J} \omega_{i,k,j,t-1} \Delta im_{k,j,t} \tag{4}$$

Note that the shares $\omega_{i,k,j,t-1}$ will in general not add up to one due to the presence of domestic sales. We calculate the total coverage of the demand shifter as $\Omega_{i,t} = \sum_{k \in K} \sum_{j \in J} \omega_{i,k,j,t-1}$ and recenter the shift-share instrument following Borusyak et al. (2022) by controlling for total coverage. In our baseline analysis, we winsorize $\Delta \tilde{z}_{i,t}$ at the 5th and 95th percentile each year to remove outliers. We also exclude firm observations with a total export exposure $\Omega_{i,t}$ below 0.05 from our baseline estimation sample.

Properties of the demand shifters. Figure 1 shows some properties of the demand shifters we construct. Panel (a) shows the mean and cross-sectional standard deviation. Since imports grow with output, our demand shifter is

[4]We deal with changes in Harmonized System product classifications over time by converting both the firm-level export data and the aggregate import data to the 1996 version of the Harmonized System using conversion Tables provided by United Nations Statistics Division (2022)

on average positive and varies over international business cycles. Our estimation only uses cross-sectional within-sector variation, as the mean will be absorbed in fixed effects. The cross-sectional standard deviation amounts to 4.5% on average over the sample period and increases to 10% during the great recession (which consequently contributes a lot to identification). In panel (b), we show the local projection of the demand shifter on its cumulative sum. Changes in our demand shifters feature almost no autocorrelation—past demand shifts do not predict current shifts. Consequently, we can reasonably treat our demand shifter as a demand *shock*. The shocks slowly decay over the following 5 years. The dynamics are well-described by an AR(1) process with annual persistence 0.97, which we plot for comparison. In our empirical estimation described below, we will focus on two types of specifications: one that treats the demand shifter as a shock as it is, and one that controls for lagged values of the demand shifter as is commonly done in the literature using local projections.

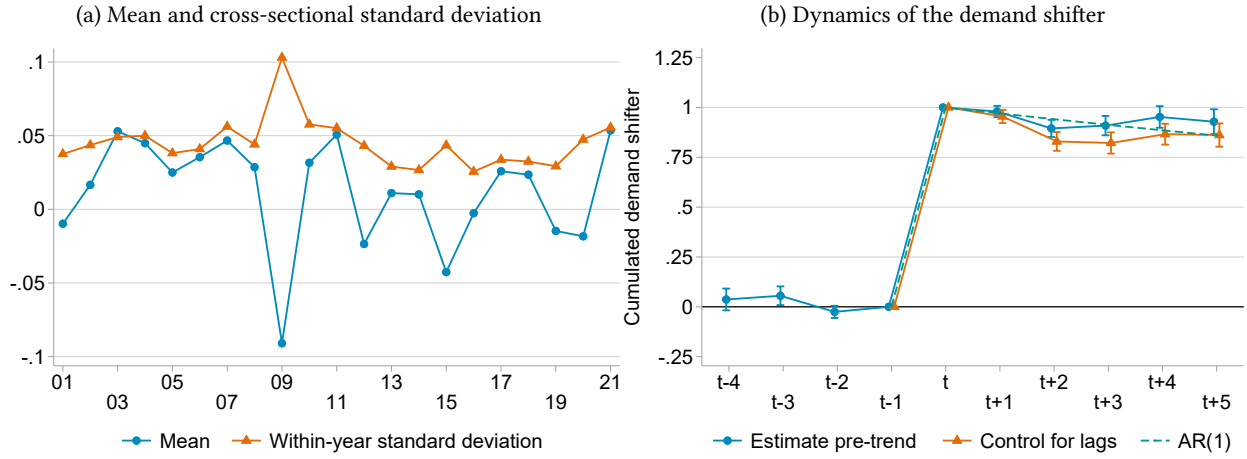


Figure 1: Properties of the demand shifter

Reduced form local projections. We estimate the relationship between prices, quantities and demand shocks using panel local projections following Jordà (2005). These local projections are estimated either at the firm level (for output and other firm-level outcomes) or firm-product level (for price outcomes). Our main reduced-form specification has changes in domestic prices for product j of firm i in Denmark on the left-hand side:

$$\Delta^h p_{i,j,t} = \beta^h \Delta \tilde{z}_{i,t} + \sum_{k=1}^{2} \alpha_k^h \Delta \tilde{z}_{i,t-k} + \sum_{k=1}^{2} \gamma_k^h \Delta p_{i,t-k} + \Omega_{i,t} \times T_{s(i),t} + v_{i,t} \tag{5}$$

All our baseline estimates include industry-time fixed effects $T_{s(i),t}$ that absorb aggregate supply shocks. Moreover, we always control for the sum of export shares interacted with the industry-time effects to recenter the shift-share

instrument and account for possibly incomplete coverage as suggested in Borusyak et al. (2022).

Local projections are often estimated including lags to isolate shocks in exogenous variables. Since there might be autocorrelation in the endogenous variables, including lagged exogenous variables necessitates including lagged endogenous variables as well. In our first baseline specification, we omit those lags, since our demand shifter features almost no autocorrelation, and can reasonably be treated as a shock as is. This has several advantages. First, it allows us to estimate placebo coefficients for negative h—we can show that there is no correlation between shocks and lagged outcomes, which serves as a test for parallel pre-trends. Second, in a short panel, including lagged endogenous variables might introduce Nickell bias (see Anderson and Hsiao, 1982, Arellano and Bond, 1991). In our second baseline specification, we include lags of $\Delta\tilde{z}$ and lagged first-differences of the dependent variable to make sure the small autocorrelation in $\Delta\tilde{z}$ doesn't affect our results. In robustness checks, we also implement the Anderson and Hsiao (1982) estimator using further lags as instruments for this case. Our main takeaway, illustrated in our results below, is that our results are not sensitive to the inclusion of lagged controls.

IV estimation. In addition to our reduced form estimates, we directly estimate the elasticity of firms' supply curve at different horizons using an IV-local projection. We regress price changes over horizon $t+h$ on firms output growth $\Delta q_{i,t}$ in period t, using $\Delta\tilde{z}_{i,t}$ as an instrument:

$$\Delta_h p_{i,j,t+h} = \beta^h \Delta q_{i,t} + \Omega_{i,t} \times T_{s(i),t} + v_{i,t} \tag{6}$$

This estimates the response of prices at horizon $t + h$ to a demand shock that increases log output by one in period t, and can therefore be directly interpreted as the slope of the supply curve at horizon h. The parallel to the alternative dynamic panel specification we estimate for our reduced form specification is to include lagged values of output growth and lagged shocks as controls to produce a first-stage regression that is identical to our reduced form regressions of demand shocks on output.

Our first-stage regression of quantities on the demand shifter involves firm-level variables, whereas our structural equation involves firm-product level prices. The standard approach would be to "stack" firm-level observations, so that each firm-level observation of quantities and the demand shifter is duplicated for each product in the first stage regression. This leads to a first stage that weights each firm by the number of products it sells, and is thus not the same regression as the firm-level reduced form regressions we also present below. To be consistent, we use the two-sample TSLS estimator of Inoue and Solon (2010). That allows us to estimate the first-stage regression at the

firm level, and the structural regression at the firm-product level. It also allows us to directly compare our first-stage results to samples in which we do not observe prices in robustness checks. We estimate standard errors that are clustered at the firm level for all our results, and follow Pacini and Windmeijer (2016) to estimate clustered standard errors for the two-sample TSLS estimator.

Concerns for identification. We anticipate two possible concerns about our identification strategy. First, firms that export to destinations with permanently higher growth rates could exhibit permanently higher growth in their prices as well—i.e., the parallel trend assumption could be violated. We address this concern by estimating placebo coefficients for negative horizons in our baseline local projections as a direct test for differences in pre-shock trends of endogenous variables that correlate with the demand shifter. We largely find no significant placebo coefficients. We also control for lagged values of the demand shifter and the dependent variable in our second baseline specification. Finally, we include firm effects that would absorb differential linear trends in our (differenced) local projection in a robustness check. We find no indication that differential trends would be a problem, and our results do not meaningfully differ between these different approaches.

Second, firm-level supply shocks could correlate with our demand shocks. A plausible scenario would be that firms import intermediates from a similar set of destinations as they export to, and variations in aggregate conditions in destination markets could then affect demand as well as input prices. To preclude this possibility, we construct a shift-share "supply shock" parallel to how we constructed the demand shock, but using firms import shares rather than export shares to weight aggregate import growth in source countries. We include this control in our baseline specification with controls. None of our results are meaningfully affected by it.

4 Empirical Results

In this section, we report our baseline empirical results—the first stage effect of demand shifters on measures of output, the reduced form effect of the demand shifter on prices, and the IV estimate of the slope of firms' supply curve. We also show that these results are robust to a large suite of robustness checks that vary our sample restrictions and estimation procedure.

4.1 Baseline results

First stage results on output. We present our main results, starting with the effect of the demand shock on measures of firms' output in Figure 2. These estimates are also the first stage of the IV estimates presented further below. They test whether the demand shock we construct is a *relevant* instrument, i.e. whether it actually shifts demand. presents our estimates. The figures include results for our two baseline specifications—first, our simple local projection of equation 5, and a second specification that controls for lagged shocks, lagged changes in the endogenous variable and our import-weighted control "supply shock". Panel (a) shows the effect on firm output. In both specifications, output increases by about 0.5 log points contemporaneously for every unit increase in export-weighted demand, increases slightly more the year after and then decays. The shock decays slightly faster in the specification with controls, but otherwise the two specifications yield very similar results. The placebo pre-shock coefficients are all close to zero and insignificant.

For a subset of firms, we observe self-reported capacity utilization (in %). In panel (b), we show that capacity utilization increases by about 20% for a unit increase in our demand shock. This validates an important assumption—some factors cannot be freely adjusted in the short run, leading to upward-sloping marginal cost and supply curves. In line with the idea that firms can adjust capacity in the longer run, the effect on capacity utilization appears to be more short-lived than the effect on output. Coefficient estimates, standard errors and summary statistics for all three output measures are reported in Tables 4 (output), and 8 (capacity utilization). Table 6 also reports results for sales, which—consistent with an effect on prices—respond slightly stronger to demand shocks than output.

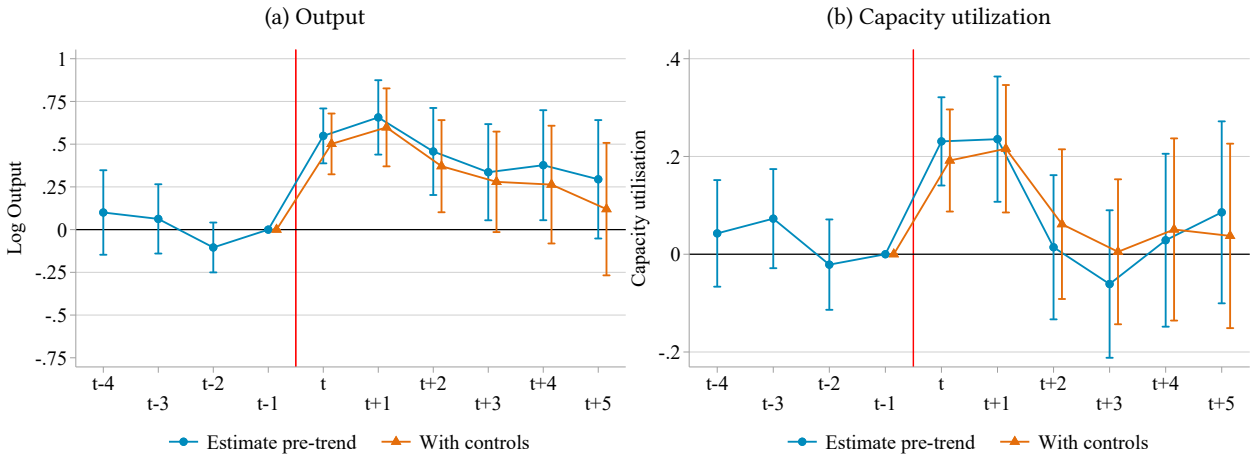


Figure 2: Effects of demand shocks on output and capacity utilization

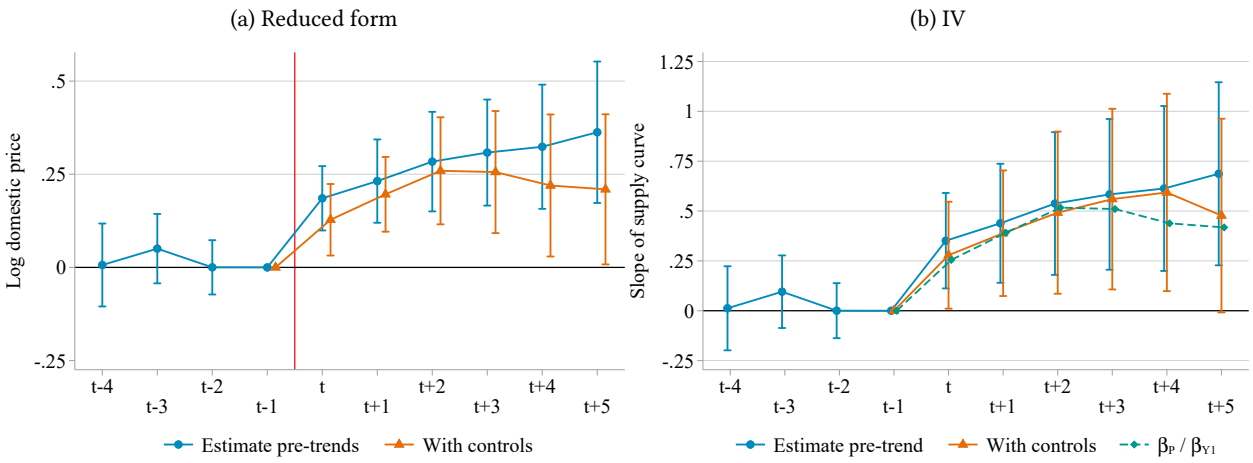


Figure 3: Effects on prices

Reduced form results on prices. In Figure 3, we show the effects of demand shocks on prices. Panel (a) shows the reduced form effect of our demand shocks on prices. Prices increase by about 0.15 log points for every unit increase in the demand shifter on impact and continue to rise slowly for 3 years up to a maximum increase of about 0.25 log points. Four years after the shock hits, prices start to slightly decline again. There is no significant pre-trend in price dynamics in the periods before a firm is hit by a shock. The results with and without controls are similar, but in our baseline with controls, prices decay from year four, while they continue to increase slightly in the specification without controls. In terms of magnitude, the results on prices line up nicely with the difference between the effects on output and sales, even though the three come from different data sources. Coefficient estimates, standard errors and other summary statistics for the reduced form effects on prices are reported in Table 10.

IV results. We present our baseline IV results in panel (b) of Figure 3. The coefficients from this estimation can be interpreted as the elasticity of prices over horizon $t + h$ to a shock demand shock that increases output in period t. Our IV estimates suggest that a 10% increase in demand leads to a peak price increase of 5%. Like in our reduced form estimates, we find no significant pre-trends. The elasticity increases over the first five years after a shock and then starts to decline. As with all our estimates, the coefficient estimates from the two baseline specifications with and without controls are very similar. In addition, the reduced form effect on prices at horizon h relative to the first-stage coefficient at a horizon of 1 year as a dashed line. Table 12 columns (1) and (2) present the coefficients for our two baseline specifications, as well as other summary statistics. The cluster-robust F-statistic for our baseline specification without controls amounts to 45, suggesting our demand shocks are a strong instrument. The baseline specification with additional controls gives a lower F-statistic of 8. Since both specifications are just-identified and

our approach to inference based on standard errors clustered at the firm level is rather conservative, we are not concerned about bias arising from this weaker first-stage (Angrist and Pischke, 2008). Moreover, in Table 13, we report results for the same specification using a larger first-stage sample that also includes manufacturing firms that are not included in the PPI, which yields a more precise first stage with an F-statistic of 14, and almost identical estimates for the response of prices.

Effects by initial capacity utilization. For a subset of about 450 firms, we observe self-reported capacity utilization. For these firms, we test how the effect of a demand shock depends on initial capacity utilization x_{t-1}. We divide a given demand shock into a below capacity component $\Delta z^{-}_{i,t} = \min(\Delta z_{it}, 1 - x_{t-1})$ that falls within the free capacity of a firm, and an above capacity component $\Delta z^{+}_{i,t} = \Delta z_{i,t} - \Delta z^{-}_{i,t}$ that exceeds the capacity limit. By construction, negative demand shocks are entirely contained in the below capacity component of a shock. We then estimate local projections simultaneously on $\Delta z^{+}_{i,t}$ and $\Delta z^{-}_{i,t}$. Because conditioning on initial capacity utilization conditions on past shocks, we estimate this specification only using our specification that controls for lagged shocks and endogenous variables.

Figure 4 shows the results of the reduced form effects on output and prices. Output increases strongly in response to the below capacity component of a shock in the first two years—more strongly than in our baseline specification. In contrast, output does not respond to the above capacity component of a shock initially. This is consistent with "hard" short-run capacity constraints that result in convex supply curves, as in Boehm and Pandalai-Nayar (2022). However, output responds strongly to the above capacity component of a shock with a lag of two to three years. This is consistent with the idea that firms expand capacity when hit by a persistent demand shock. While this response is imprecisely estimated, it appears to be larger than the initial response to the below capacity component of a shock. Prices respond immediately to both the below and above capacity of a shock. However, the response to above capacity shocks is substantially larger and more persistent. Because the effect of the above capacity shock on initial output is zero, we can't estimate the slope of the supply curve using our IV specification (the supply curve in this case would have an infinite slope).

4.2 Robustness of main results

We test the robustness of our results to variations in the exact specification of our estimation and sample restrictions we apply.

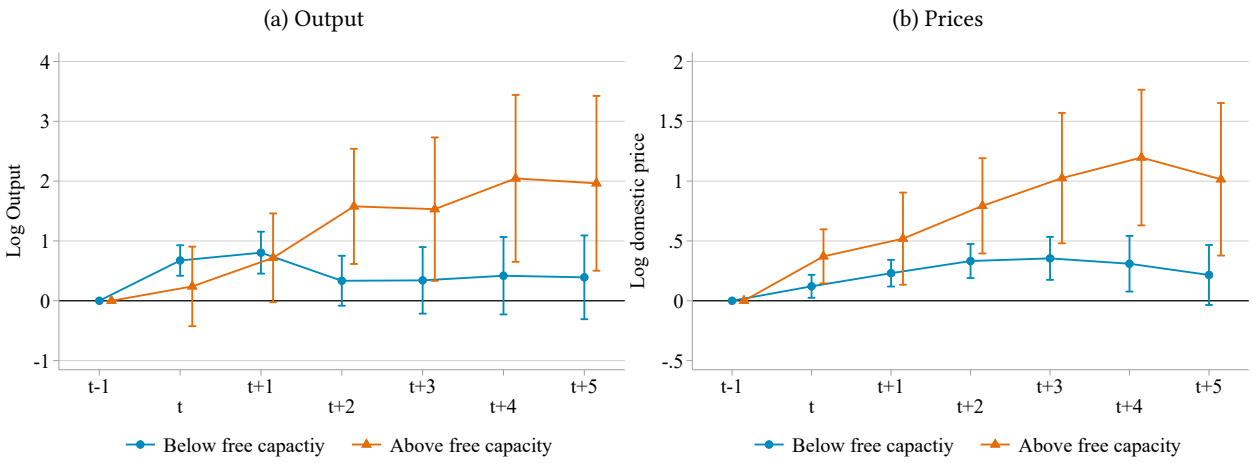


Figure 4: Effects on prices and output by initial capacity utilization

Specification checks. We use our more restrictive baseline specification with controls as the starting point and slightly vary the specifications we estimate. Our main robustness checks include specifications that add firm fixed effects that would absorb linear time trends in the local projections differences, replace 2-digit sector-time fixed effects with more finely grained 4-digit sector-time fixed effects or more coarse plain time fixed effects. Moreover, for our reduced form estimates, we present results using the Anderson and Hsiao (1982) estimator in which lags of first-differenced endogenous variables are instrumented using further lags of levels of the endogenous variable. The robustness checks on our main results on output and prices (Tables 4 and 10) as well as supplementary results on sales and capacity utilization (Table 6 and 8) show that our conclusion is robust to all of these variations. Table 12 shows robustness checks for our IV estimates. These effects are naturally less precisely estimated in our baseline due to the higher variance of the IV estimator. Including additional firm fixed effects or more finely grained sector-time fixed effects turns some of the estimates insignificant, but they remain of similar magnitude as in our baseline.

Sample checks. We also show that our results are robust to several variations of our sample restrictions. Our baseline restrictions are that firms have to be included in the PPI sample (also for output outcomes, which we observe for all manufacturing firms with more than 10 employees), have more than 20 employees for at least 5 years in a row during the sample period and are in a manufacturing sector (the PPI also includes some wholesalers and firms in agriculture). In addition, we restrict the sample of our local projections to firms that have an export share of at least 5% in the year they are hit by a given shock. We estimate the first-stage specifications for output, sales and capacity utilization for a sample that includes all firms in the PPI and all manufacturing firms (dropping all other restrictions). We also estimate specifications with tighter restrictions. We present results in a more strongly balanced panel that

only includes firms that have more than 20 employees for 8 years in a row, and in a panel of more export-oriented firms with an export share of at least 1/3 in the year they are hit by a given shock. Finally, we present results for a panel that only includes the period up to 2020. Our main reduced form results on the response of output and prices (Table 5 and 11), as well as supplementary results on sales and capacity utilization (Table 7 and 9) are robust in all of these samples. Our IV results (Table 13) are robust in terms of magnitude of the coefficients, and estimates mostly remain significant at least at the 10% level, although the precision is naturally lower in smaller samples.

5 Macroeconomic implications

In this section, we address the aggregate implications of the firm-level behavior we document in our empirical work. We show in section 4 that in the cross-section prices respond strongly to an exogenous demand shock. This is an all-else-equal estimate—ideally, all other factors that might affect prices are orthogonal to the shock or held constant by including fixed effects and controls. When the economy is hit by an aggregate shock, firms respond to a change in their own demand, but also to changes in inflation expectations that arise because other firms also experience an increase in demand. Moreover, changes in other aggregate variables, such as real wages, might also affect prices. The fixed effects in our empirical analysis absorb such general equilibrium effects. Our empirical results on their own are thus not sufficient to make predictions about the effect of an aggregate demand shock on the aggregate price level. For that, we need a more structural model. To derive macroeconomic implications of the firm behavior we estimate, we use a New Keynesian model augmented with firm-level demand shocks that mirror our empirical setting. Our strategy is to match the response to those shocks in the model to our estimates. This pins down key parameters of the model and implies an aggregate Phillips curve that is consistent with micro-level firm behavior.

In our empirical strategy we use the fact that Denmark is a small open economy and most Danish manufacturing firms are exporters. This would naturally also affect the slope of the Danish Phillips curve—in a small open economy, the Phillips curve is flatter than in a closed economy because the response to domestic shocks is muted by trade (Gali and Monacelli, 2005). We derive macroeconomic implications in a model of a closed economy in order to remain comparable to the empirical and quantitative literature. This model should be interpreted as a model of the Euro area, to which Denmark is connected through its long-standing currency peg to the Euro[5]. Our quantitative results thus require that Danish manufacturing firms are representative of the larger European manufacturing sector, and that European manufacturing firms therefore respond to demand shocks the same way as Danish firms.

[5]The Danish krone has been pegged to the Euro at a fixed exchange rate since the conception of the Euro area. Before that, it has been pegged to the Deutsche Mark since 1982.

5.1 Model setup

Final goods production. The final good in the model economy is produced by perfectly competitive producers with flexible prices, who turn intermediate goods into a final consumption good using a CES production function:

$$Y_t = \left(\int_i Z_{i,t}^{\frac{1}{\sigma}} Y_{i,t}^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} \tag{7}$$

$Z_{i,t}$ is an idiosyncratic demand shock for intermediate products that follows an AR(1) with persistence ρ and mean zero in logarithms. Since we model a closed economy, we abstract from the fact that the demand shocks in our empirical setting originate from abroad. The demand shock $Z_{i,t}$ resembles the idiosyncratic demand shock in section 3, and as we have shown there, the response to such a shock is the same as to weighted aggregate import demand shocks. Final goods producers minimize their expenditure $\int_i Y_{i,t} P_{i,t} di$ subject to (7):

$$\min_{\{Y_{i,t}\}} \int_i P_{i,t} Y_{i,t} di \quad \text{s.t. } Y_t = \left(\int_i Z_{i,t}^{\frac{1}{\sigma}} Y_{i,t}^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}. \tag{8}$$

This results in a CES demand curve for intermediate goods:

$$Y_{i,t} = Y_t Z_{i,t} \left(\frac{P_{i,t}}{P_t} \right)^{-\sigma}. \tag{9}$$

Since final good producers are competitive and use intermediates as their only input, the intermediate CES price index P_t corresponds to the price of final output.

In the baseline specification of our model, we use a simple CES aggregator for final output. In an extension, we use a Kimball aggregator (see Kimball (1995)) that results in a variable demand elasticity for intermediates. We describe the model with Kimball demand for intermediates in detail in Appendix D and discuss results in parallel to the CES case below. As we discuss in more detail below, the choice between CES and Kimball demand is not important for the slope of the Phillips curve in our setting—as long as we match the same empirical response of prices to demand shocks, the two specifications will produce different underlying dynamics of marginal cost and markups that add up to the same price response and an almost identical Phillips curve.

Intermediate goods production. Intermediates are produced by monopolistically competitive firms using a normalized CES production function, where labor and capital inputs are normalized with their steady-state values K_{SS}

and L_{SS}:

$$Y_{i,t} = \left(\alpha \left(\frac{K_{i,t}}{K_{SS}} \right)^{\psi} + (1-\alpha) \left(\frac{L_{i,t}}{L_{SS}} \right)^{\psi} \right)^{1/\psi}$$

We use a normalized CES, rather than a standard CES production function, to do comparative statics w.r.t. α and ψ without affecting the steady-state capital-to-labor ratio (Cantore et al., 2015). The production function converges to a standard Cobb-Douglas specification as $\psi \to 0$ and to a Leontief production function as $\psi \to -\infty$. We assume that capital inputs are fixed in the short run and that consequently marginal cost varies with output depending on the share parameter α and the substitution parameter ψ. Even with price-adjustment frictions, the slope of the flexible-price supply curve is an important determinant of the response of prices to aggregate and firm-level demand shocks. With CES demand, the slope of the log-linearized flexible-price supply curve equals the output elasticity of marginal cost:

$$\delta = \frac{(1-\psi)\alpha}{1-\alpha}. \tag{10}$$

Intermediate producers can reset their price with probability $1-\theta$ and discount the future at rate β. Firms maximize their future discounted profit whenever they have an opportunity to reset their price to a new price $P^*_{i,t}$. This yields the first-order condition:

$$\sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(Y_{i,t+k}(1-\sigma) \left(P^*_{i,t} - \frac{\sigma}{\sigma-1} MC_{i,t+k} \right) \middle| z_{i,t} \right) = 0, \tag{11}$$

where $Y_{i,t+k}$ and $MC_{i,t+k}$ depend on the chosen reset price, future states of idiosyncratic demand and aggregate state variables.

We log-linearize price-setting firms' FOC around a zero-inflation steady state in which idiosyncratic demand is equal to its mean. This implies there is no heterogeneity in prices or quantities in the steady state, like in a New Keynesian model without firm-specific demand shocks. We use lower-case letters to denote log deviations of a variable from its steady state value. We get the optimal reset price as a function of firms' current and expected future nominal marginal cost:

$$p^*_{i,t} = (1-\beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(mc^R_{i,t+k} + p_{t+k} \middle| z_{i,t} \right) \tag{12}$$

We use the production function and demand curve to derive an expression for the deviation of marginal cost from its steady state value, which is equal to $mc^R_{i,t+k} = w^R_{t+k} + \delta(z_{i,t+k} - \sigma(p^*_{i,t} - p_{t+k}))$. Plugging in and using the fact that $E(z_{i,t+k}|z_{i,t}) = \rho^k z_{i,t}$, we can pin down the optimal reset price as a function of the current realization of idiosyncratic demand, aggregate marginal cost and inflation expectations:

$$p^*_{i,t} - p_{t-1} = \frac{(1-\beta\theta)}{1+\sigma\delta} \sum_{k=0}^{\infty} (\beta\theta)^k \left(E_t \left(mc^R_{t+k} \right) + \delta\rho^k z_{i,t} \right) + \sum_{k=0}^{\infty} (\beta\theta)^k E_t(\pi_{t+k}), \tag{13}$$

Aggregate dynamics. Our model is in the aggregate identical to a textbook New Keynesian model as in Gali (2008) since idiosyncratic shocks cancel out on average. Around a zero inflation steady state, aggregate inflation is approximately equal to $\pi_t = (1-\theta)\left(\int_i p^*_{i,t} di - p_{t-1}\right)$. The marginal cost Phillips curve follows from combining this approximation with equation (13):

$$\pi_t = \lambda mc^R_t + \beta E(\pi_{t+1}), \tag{14}$$

where the marginal cost coefficient $\lambda = (1-\beta\theta)(1-\theta)/(\theta(1+\sigma\delta))$.

The deviation of aggregate real marginal cost from its steady state value is equal to $mc^R_t = w^R_t + \delta y_t$. To go from the marginal cost Phillips curve to the output Phillips curve, we thus need to define the relationship between real wages and output, i.e. the wage Phillips curve. In most New Keynesian models, this relationship is pinned down by households' labor supply decision and possibly subject to adjustment frictions as well. Here, we instead define a simple reduced form wage Phillips curve $w^R_t = \phi y_t$ that summarizes the co-movement of output and real wages in a single parameter ϕ. This leads to the output formulation of the Phillips curve:

$$\pi_t = \kappa y_t + \beta E(\pi_{t+1}). \tag{15}$$

We can decompose the slope of the output Phillips curve κ into two components reflecting dynamics in product and labor markets:

$$\kappa = \underbrace{\frac{(1-\theta\beta)(1-\theta)}{\theta} \frac{\delta}{1+\sigma\delta}}_{\text{Capacity pressure} \equiv \kappa^p} + \underbrace{\frac{(1-\theta\beta)(1-\theta)}{\theta} \frac{\phi}{1+\sigma\delta}}_{\text{Wage pressure} \equiv \kappa^w}. \tag{16}$$

The first component κ^P results from upward-sloping firm-level supply curves. Since this slope is a consequence of

production factors that are fixed in the short run, we call this part of the Phillips curve slope the "capacity pressure" channel. With a CES production function, the strength of this channel is determined by the importance of fixed factors for production α and the extent ψ to which flexible and fixed factors can be substituted. If capital and labor are perfect substitutes, i.e., $\psi = 1$, firms operate with a production technology that is linear in labor, have flat supply curves, and the capacity pressure channel doesn't contribute to the slope of the Phillips curve. As ψ decreases, this channel becomes stronger.[6] Our empirical estimates are directly informative about the importance of this channel.

The second component κ^W reflects the "wage pressure" channel that results from co-movement of real wages and output, i.e., the wage Phillips curve. Since wages are absorbed in fixed effects in our firm-level estimation, our estimates are not informative about the importance of this channel. The same is true for other estimates that use cross-sectional variation in demand, such as Gagliardone et al. (2024), without strong additional assumptions such as a lack of labor mobility between sectors or firms. While this is not our primary contribution, we will use aggregate data to estimate the value of ϕ and relate the magnitudes of the capacity and wage pressure channels in our results below.

The slope of the Phillips curve determines the response of inflation to an aggregate demand shock at fixed inflation expectations. In practice, we are often interested in the overall response including the effect coming from adjustments in expectations. To derive an expression for this overall effect, we need to complement the Phillips curve with a dynamic IS equation:

$$y_t = \eta_y \mathbb{E}_t y_{t+1} + \eta_\pi \mathbb{E}_t \pi_{t+1} + u_t. \tag{17}$$

In a fully specified RANK or TANK model, the coefficients of the dynamic IS equation depend on households' Euler equation and the monetary policy rule (see e.g., Bilbiie (2020)). The coefficients are needed to solve for the separate paths of aggregate output and inflation. But they are not necessary to characterize the output-inflation trade-off faced by policy makers, and hence we stick to a simple reduced form characterization. We assume that the aggregate demand shock u_t follows an AR(1) process with persistence ρ_u. We can then use the method of undetermined coefficients to solve:

$$\pi_{t+k} = \frac{\kappa}{1 - \beta\rho_u} \Lambda u_{t+k}$$

$$y_{t+k} = \Lambda u_{t+k}$$

[6] If $\psi \to -\infty$ then firms operate with Leontief production technology. This extreme case is discussed in Boehm and Pandalai-Nayar (2022), but since in this setting, firms' supply curves are highly non-linear, it is not described well by a log linearized model.

The parameter Λ depends on the coefficients of the dynamic IS-equation. However, the relative response of output and inflation only depends on κ, β and the persistence of the shock ρ_u. Consequently, we focus on two statistics to summarize the aggregate output-inflation trade-off implied by the model: The slope of the output Phillips curve κ and the output Phillips multiplier Ψ (Barnichon and Mesters, 2021). The output Phillips multiplier is the ratio of the response of inflation and output over horizon h[7]. We can see that the Phillips multiplier implied by the model is constant:

$$\Psi = \frac{\kappa}{1 - \beta\rho_u} \tag{18}$$

While the slope of the Phillips curve captures the trade-off between output and inflation keeping inflation expectations fixed, the Phillips multiplier captures the effect of increasing inflation expectations as well. However, it requires additional assumptions on the behavior of households—in particular, the existence of a dynamic IS-equation-like equation (17).

Response to idiosyncratic shocks Using equation (13), we can express the reset price of a firm with demand realization z relative to the average reset price as:

$$p_t^*(z) - \int p_t^*(z) dF(z) = \frac{1 - \beta\theta}{1 - \beta\theta\rho} \frac{\delta}{1 + \sigma\delta} z \tag{19}$$

Our empirical estimates identify the average path of prices conditional on a firm-level demand shock at time t relative to the sectoral average price path, but don't condition on prices being updated. To get a model counterpart to our estimates we thus need to average over relative prices of firms that are updated and those that are not. This average can be expressed as:

$$p_{t+k}(z_t) - p_{t+k} = (1 - \theta)(p_{t+k}^*(z_t) - p_{t+k}^*) + \theta(p_{t+k-1}(z_t) - p_{t+k-1})) \tag{20}$$

The previous period's average price conditional on period t demand, $p_{t+k-1}(z_t)$, will not be equal to the unconditional average price p_{t+k-1} because prices might have previously been updated in response to the time t demand shock. This is only true at horizon $k = 0$, because demand shocks in period t don't correlate with period $t-1$ prices. Hence we can iterate equation (20) backward and use the fact that $p_{t-1}(z_t) = p_{t-1}$. The average period $t + k$ price

[7]Barnichon and Mesters (2021) use the average response of inflation and the unemployment gap instead.

of firms hit with demand shock z_t in period t is equal to:

$$p_{t+k}(z_t) - p_{t+k} = \frac{(1-\beta\theta)(1-\theta)}{1-\beta\theta\rho}\frac{\theta^{k+1}-\rho^{k+1}}{\theta-\rho}\frac{\delta}{1+\sigma\delta}z_t, \tag{21}$$

Given equation (21), the corresponding response of average relative output to firm-level demand shocks is determined by log-linearizing the demand curve for intermediates given by equation (9):

$$y_{i,t+k}(z_t) - y_{t+k} = z_{i,t} - \sigma(p_{t+k}(z_t) - p_{t+k}) \tag{22}$$

Equation (21) is the model counterpart to our estimate of the response of prices to firm-level demand shocks. The response allows for three regimes. If idiosyncratic demand shocks are permanent, i.e. $\rho = 1$, then in response to a positive demand shock, the firm's log relative price will converge to a permanently higher level. The new relative price is determined by the flexible price supply and demand curves. In the empirically relevant case where demand shocks are transitory and $\rho < \theta$, demand shocks decay faster than prices are adjusted, and firms' log relative price increases in the first period and slowly returns to zero afterward. Finally, if $\rho > \theta$, the shock decays more slowly than prices are adjusted, and the response is hump-shaped—relative prices increase over several periods initially before slowly returning to zero.

The close connection between the response of prices to idiosyncratic demand shocks and the capacity pressure channel of the Phillips curve slope is immediately apparent. The magnitude of each is determined by the slope of the flexible-price supply curve δ and the demand curve σ. However, they are not the same, and there are important qualitative differences. Positive aggregate demand shocks generate a transitory increase in inflation and a permanent increase in the price level. Firm-level demand shocks generate a transitory increase in relative prices (unless the shocks themselves are permanent) that fades away over time. Our structural model is necessary to map the response of prices to firm-level demand shocks to the aggregate.

5.2 Model estimation

Estimation and identification. We estimate key parameters of the model by fitting the model impulse response of prices and quantities to a firm-level demand shock given by equations (21) and (22) to the local projections estimated in section 4. We formulate the model in quarterly frequency and aggregate prices and quantities to annual frequency consistent with the aggregation of the data (we use log deviations in the last quarter of a year for prices, and yearly

averages over quarterly log deviations for quantities).

Several parameters of the model can be directly observed in the data or have clear benchmark values. We fix the quarterly discount factor to $\beta = 0.99$, implying an annual interest rate of about 0.04. We set the frequency of price adjustment to $\theta = 0.66$, consistent with the frequency of price changes in the Danish PPI microdata. We set the persistence of idiosyncratic demand shocks to $\rho = 0.992$, which matches a persistence of 0.97 in our annual demand shock. Finally, we fix the cost share parameter of the CES production function to its standard value of $\alpha = 0.33$ to match the labor share in GDP.

The remaining parameters are the production function substitution parameter ψ—which for a given value of α pins down the slope of the flexible-price supply curve δ—and the slope of the demand curve σ. We choose values of ψ and σ that minimize the sum of the mean squared error between the estimated and model-implied impulse responses for output and prices, weighted by their inverse standard errors. The intuition behind identification works as follows. Fitting the model impulse response of prices to firm-level demand shocks to the response in the data identifies the combined value of $\delta/(1 + \delta\sigma)$. The demand elasticity σ is identified from fitting the impulse response of output for a given response of prices. Combining these two values determines ψ. Finally, we estimate a normalization of the initial value of idiosyncratic demand shocks $z_{i,t}$, which is similar to the first-stage parameter in an IV estimation.[8]

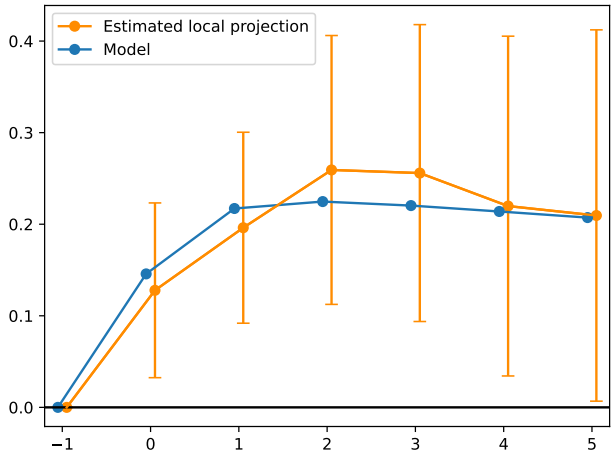
Baseline parameter estimates. Figure 5 shows the impulse responses of prices and output in the fitted model compared to the local projections estimated in section 4. Our simple model with only two free parameters does remarkably well at fitting the estimated responses to demand shocks both in terms of magnitude and timing. The model reproduces the increase in prices over 3 years and the slight decrease afterward. The output response in the model peaks in the year the shock hits, compared to the second year in the data and decreases slightly afterwards.

Table 2 shows the estimated structural parameters. Our baseline estimates for the CES demand case in column (1) are $\psi = -0.38$ and $\sigma = 4.13$. The negative value of ψ implies the supply curve is steeper than with a Cobb-Douglas production function. With a calibrated value of $\alpha = 0.33$, the implied slope of the flexible price supply curve is $\delta = 0.68$, compared to about 0.5 in the Cobb-Douglas case. Our estimate of the demand elasticity $\sigma = 4.13$ is close to benchmark values typically chosen in calibrated models to produce reasonable steady-state markups, which in our case would amount to $\sigma/(\sigma - 1) = 1.32$.

For comparison, we show the response to firm-level demand shocks in a model with a standard Cobb-Douglas pro-

[8] This parameter would be irrelevant if we target the MSE of the ratio of the price to output response directly. However, since the empirical output response comes close to zero in some periods, targeting this ratio results in less stable estimates than targeting the sum of the MSE of both responses.

(a) Impulse response of prices



(b) Impulse response of output

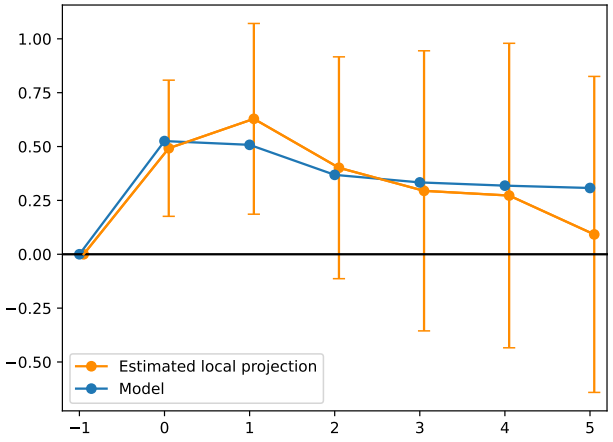


Figure 5: Impulse responses to an idiosyncratic demand shock in the data and our fitted model

duction function with $\alpha = 0.33$ and CES demand with $\sigma = 6$ in Figure 7 in the Appendix. This corresponds to the calibration of the textbook New Keynesian model in Gali (2008). The response of prices to a demand shock is smaller, and the response of output is slightly larger, but for the case of CES demand, the standard calibration roughly reproduces the patterns we observe in our empirical estimates.

CES vs Kimball demand. We also estimate an extension of the model that features Kimball demand for intermediates in line with Gagliardone et al. (2024). This more general model nests the CES baseline case, and is described in detail in Appendix D. With Kimball demand, pass-through of marginal cost is incomplete and markups fall after a positive demand shock even with flexible prices. The strength of the response of *desired* markups is governed by the parameter Γ, the elasticity of desired markups to the firms' relative price. The model with Kimball demand takes mostly the same form as the CES baseline described above, except that the slope of the flexible-price supply curve δ^{Kimball} now reflects incomplete pass-through of marginal cost governed by Γ:

$$\delta^{\text{Kimball}} = \frac{(1-\psi)\alpha}{(1+\Gamma)(1-\alpha)} \tag{23}$$

Fitting the price response in the model to the data still allows us to identify the slope of the flexible-price supply curve δ^{Kimball}, but not the parameters Γ and ψ separately—for any value of Γ, we can pick ψ to produce some given δ^{Kimball}. Naturally, different combinations of Γ and ψ vary in the contributions of marginal cost and markups to the response of prices. We therefore fix $\Gamma = 0.66$, roughly consistent with estimates for the pass-through of marginal

Table 2: Structual parameter estimates

	(1) Baseline (CES demand)	(2) Kimball demand	(3) Gali calibration
Estimated parameters			
Production function substitution ψ	-0.381	-1.293	0.000
Demand elasticity σ	4.125	4.125	6.000
Fixed parameters			
Production function share α	0.330	0.330	0.330
Firm-specific demand persistence ρ	0.992	0.992	0.992
Frequency of price adjustment θ	0.660	0.660	0.660
Elasticity of markups Γ	0.000	0.660	0.000
Slope of flex-price supply curve δ	0.680	0.680	0.493
Mean squared error	1.398	1.398	3.293

Notes: The table plots the structural parameter estimates that match the impulse response of prices and output in the model (equations (21) and (22)) to the local projections estimated in section 4. (1) Estimates the parameters under CES demand ($\Gamma = 0$) and (2) under Kimball demand ($\Gamma = 0.5$).

cost shocks in Amiti et al. (2019) and Gagliardone et al. (2024), and choose ψ to fit the price response we observe in the data.

The resulting parameter estimate for the Kimball case is $\psi = -1.29$ (column 2 of Table 2), resulting in the exact same slope of the flexible price supply curve as in the CES demand case, i.e. $\delta = \delta^{\text{Kimball}} = 0.68$. Consequently, the estimated slope of the demand curve σ and the MSE to our empirical estimates is also the same with both specifications of demand. The differences between CES and Kimball demand are illustrated in Figure 6. In both cases, marginal cost increases and markups fall after a positive demand shock. However, with Kimball demand pass-through of marginal cost is incomplete even after prices have been updated, and the model generates a stronger increase in marginal cost to produce a price response that matches our empirical estimate. The elasticity of marginal cost to output is now 1.13, i.e. almost double the elasticity in the model with CES demand. Markups return to their steady-state value after 2 years with CES demand, while with Kimball demand, markups are lower as long as a firm's price is elevated. Finally, an important difference between the CES and Kimball models is that with Kimball demand, the model wouldn't generate a price response close to the data with a standard Cobb-Douglas production function, and the CES production function becomes necessary to reproduce the empirical patterns.

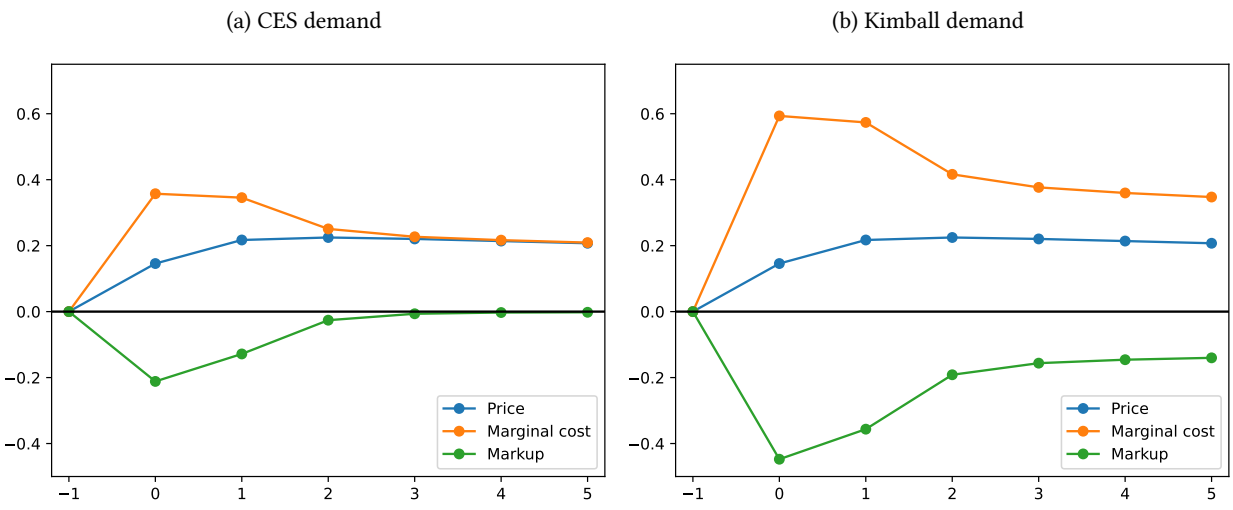


Figure 6: Impulse response of prices, marginal cost and markups to a firm-level demand shock in the fitted model

5.3 Aggregate implications

We now turn to the slope of the Phillips curve implied by a model that fits the estimated response to firm-level demand shocks. We separately calculate the components of the slope κ^P ("capacity pressure", see equation (16)) and, using additional assumptions on the response of real wages ϕ, κ^W ("wage pressure"). We also calculate Phillips multipliers Ψ, using $\rho_u = 0.69$ as the persistence of aggregate demand shocks to match the properties of ECB monetary policy surprises (Jarocinski and Karadi, 2020). The results of our calculations are summarized in Table 3.

Capacity pressure channel. Our estimates suggest that $\kappa^P = 0.032$. That means that after an aggregate demand shock that increases output by 1%, upward-sloping supply curves lead to an initial increase in inflation by 0.032, holding real wages and inflation expectations fixed. Taking into account the effects of changes in inflation expectations, the initial increase in inflation corresponds to $\Psi^P = 0.101$.

These numbers are the same whether we assume CES or Kimball demand since both produce the same supply curve. This highlights an important insight. The slope of the Phillips curve depends on δ, not on the separate elasticities of marginal cost and markups. δ is pinned down by the response of prices to firm-level demand shocks we estimate in the data, not by assumptions on the hard-to-observe behavior of markups and marginal cost.

We compare these values to a benchmark calibration of the New Keynesian model in Gali (2008) which uses a Cobb-Douglas production function (see Table 2) and produces values $\kappa^P = 0.021$ and $\Psi^P = 0.067$. The Phillips curve that is consistent with our estimates features a capacity pressure channel that is more important than suggested by

this calibration—in particular, the capacity pressure channel contributes about 1.5 times these values to the slope of the Phillips curve.

Table 3: Slope of the Phillips curve and Phillips multipliers

	κ^P	κ^W	κ	Ψ^P	Ψ
Based on our estimates (CES)					
with $\phi = -0.15$	0.032	-0.007	0.025	0.101	0.079
with $\phi = 2.50$	0.032	0.117	0.149	0.101	0.471
Based on our estimates (Kimball)					
with $\phi = -0.15$	0.032	-0.005	0.027	0.101	0.086
with $\phi = 2.50$	0.032	0.078	0.11	0.101	0.347
Comparison to other estimates					
Gali (2008) calibration	0.021	0.106	0.128	0.067	0.402
Barnichon and Mesters (2020)			0.12		0.235
Barnichon and Mesters (2021)			0.181		0.181
Hazell et. al. (2022)			0.008		0.013
Gagliardone et. al. (2024) - Output PC			0.017		0.054

Notes: The table compares the slope of the Phillips curve κ^P, κ^W and κ, as well as the Phillips multipliers Ψ^P and Ψ for different sets of parameters. Rows 1 and 2 use our estimates under CES demand. Rows 3 and 4 use our estimates under Kimball demand. $\phi = 2.50$ corresponds to the cyclicality of real wages in the calibration of Gali (2008). $\phi = -0.15$ corresponds to the cyclicality of real wages in Danish data. Rows 5–10 compare our estimates with different benchmarks in the Literature that are made comparable to our estimates. Barnichon and Mesters (2021) estimate the trade-off between unemployment and inflation, which we convert to an output-inflation trade-off using the correlation of the US output and unemployment gaps. We do the same for the estimates from Hazell et al. (2022). Additionally, we also adjust the Phillips multiplier for the differences in shock persistence. Therefore, these numbers differ from the ones in their paper.

Overall slope. To compare the importance of the capacity pressure channel to the slope of the overall Phillips curve, we consider two alternative scenarios for the behavior of real wages after an aggregate demand shock. For our preferred scenario, we estimate the response of Danish real wages and output to ECB monetary policy surprise taken from the dataset provided by Jarocinski and Karadi (2020). This estimation is described in detail in Appendix E. We use the average of the impulse response over a horizon of 8 quarters to summarize the response in a single parameter. Danish real wages fall after a monetary policy shock that increases output, with an output elasticity of $\phi = -0.15$. Combined with the other parameters of the model, this suggests $\kappa^W = -0.007$ and a slope of the overall Phillips curve equal to $\kappa = 0.025$ in the CES case.[9] This corresponds to a relatively flat Phillips curve, whose slope is entirely

[9]Note that in this case the Kimball specification produces slightly different values, since the parameters ψ and Γ impact κ^W slightly differently (see Equation 16). However, the difference between the two specifications is not economically significant.

dominated by the capacity pressure channel.

For comparison, we take the relationship of real wages and output suggested by the calibration in Gali (2008), which uses log utility, a Frisch elasticity of 1, and an elasticity of output to employment of $\alpha = 0.33$. This results in a strong co-movement of real wages and output, with $\phi = 2.5$. While this case presents an interesting benchmark, this extent of real wage cyclicality is clearly at odds with Danish aggregate data. This results in a contribution of the wage pressure channel to the slope of the Phillips curve of $\kappa^W = 0.12$ and an overall slope of $\kappa = 0.15$. The Phillips multiplier Ψ in this case amounts to $\Psi = 0.47$. With this relatively steep Phillips curve, the product market channel contributes about one fifth of the slope.

Our analysis in this section yields two key takeaways. First, in a model that is consistent with our firm-level estimates, the capacity pressure channel is an important contributor to the slope of the Phillips curve. In our preferred calibration that also fits the slight counter-cyclicality of Danish real wages, it contributes all of the slope of the Phillips curve. Second, the capacity pressure channel is slightly more important in a model that fits firm behavior than suggested in benchmark calibrations. The overall Phillips curve is still rather flat, and this is due to a flat wage Phillips curve.

Discussion of alternative estimates We compare the slope of the Phillips curve implied by our estimates to several values from the recent literature in Table 3. We provide both the slope estimates as well as a Phillips multiplier, which we calculate as the response to a monetary policy shock with persistence 0.69, the persistence of ECB monetary policy shocks in Jarocinski and Karadi (2020).

As a benchmark, we first compare our estimates to the textbook calibration in Gali (2008). This calibration produces a Phillips curve that is much steeper than our preferred estimate, with a slope of $\kappa = 0.13$ and a Phillips multiplier of $\Psi = 0.40$. The contribution of the capacity pressure channel to the slope of the Phillips curve suggested by our preferred estimate is actually about 1.5 times larger than in the Gali calibration. The large difference in the overall slope arises from the fact that the Gali calibration assumes a steep wage Phillips curve that is at odds with Danish aggregate data.

We can also compare our results to recent macro estimates of the US Phillips curve slope in Barnichon and Mesters (2020) and Phillips multiplier in Barnichon and Mesters (2021). Both papers find a trade-off between output and inflation[10] that is flatter post-1990 compared to earlier decades. The magnitude of the output Phillips curve slope they report for the 1990–2008 period is 0.12 and 0.18 respectively, i.e., of a similar magnitude as suggested by the

[10]Barnichon and Mesters (2021) estimate the trade-off between unemployment and inflation, which we convert to an output-inflation trade-off using the correlation of the US output and unemployment gaps.

Gali calibration. However, the confidence bands of their estimates are very wide, and it is not possible to rule out a flat Phillips curve based on their estimates.

Two recent papers use regional and firm-level panel data to estimate Philips curves from cross-sectional data. Gagliardone et al. (2024) estimate pass-through of marginal cost into prices of Belgian manufacturing firms. This approach implies a slope of the marginal cost Phillips curve of around 0.05. However, going from pass-through estimates to the slope of the output Phillips curve, requires an estimate of the elasticity of marginal cost to output. This elasticity is embodied in firms' supply curve and hence can't be identified without a demand shock. Gagliardone et al. (2024) estimate an output Phillips curve using aggregate monetary policy surprises allowed to have heterogeneous effects in different sectors as sectoral demand shifters. This is similar to our approach, and yields a slope estimate of 0.017. However, their approach is subject to the same concerns over the weakness as monetary policy surprises as instruments as the aggregate time-series literature using the same shocks. We replicate their estimation using Danish data in Appendix F. We show that once we correctly take into account the estimation uncertainty surrounding the heterogeneous first-stage parameters, the first stage of the monetary surprises IV becomes essentially insignificant, with an F-statistic of 1.2. A LIML estimate that is more robust to many weak instruments is larger, but imprecise and insignificant. A second important difference between Gagliardone et. al. and our work concerns the interpretation of the estimated parameters. While the regressions in both papers include sector-time fixed effects, Gagliardone et. al. argue that wage dynamics vary within 2-digit sectors, and that their estimates capture both the capacity-pressure and wage-pressure channel (in our terminology). We view segmentation of labor markets at such a fine level as unlikely, and would argue that 2-digit-sector-time fixed effects would absorb most if not all cyclical variation in wages.

Finally, Hazell et al. (2022) estimate Phillips curves using state-level data on prices of non-tradeables. They use variation in the local exposure to aggregate shocks to US tradeable industries as a source of variation in local aggregate demand. They estimate[11] a slope of 0.008. This is even flatter than the slope implied by our estimates, even though real wages in the US seem more pro-cyclical than in Denmark. Our estimates relax the most important identification assumptions that are required in their approach. First, they require perfect labor mobility between the tradeable and non-tradeable sectors to make sure that the response of state unemployment to tradeable industry shocks is a good measure of slack for the non-tradeable industry. If labor mobility is limited, unemployment rates in the two sectors are not equalized, and using the overall state unemployment rate as the forcing variable in non-tradeable firms' Phillips curve scales their price response incorrectly and biases their estimates downward. Second, there could be

[11] We make their estimates comparable to ours by converting their unemployment Phillips curve to an output gap equivalent using the covariance of aggregate US unemployment with the output gap. We also adjust the Phillips multiplier for the differences in shock persistence. Therefore, these numbers differ from the ones in their paper.

labor mobility between states. In this case, a shock to local aggregate demand would lead to inflows of workers and both aggregate demand and aggregate supply would be affected.

6 Conclusion

Our paper makes several contributions to understanding price dynamics at both the firm and aggregate levels. Using detailed Danish microdata and an identification strategy based on cross-sectional variation in firms' exposure to foreign demand shocks, we estimate how firms adjust their prices in response to firm-level demand shocks. Our findings show that supply curves are steeper than commonly assumed in macroeconomic models—a demand shock that increases output by 1% leads to a 0.5% increase in prices.

By augmenting a standard New Keynesian model with firm-level demand shocks, we show how our firm-level estimates map to aggregate price dynamics through a "capacity pressure" channel. Our fitted model suggests that for realistic values of wage cyclicality, the slope of the Phillips curve is mostly determined by the "capacity pressure" channel. Our estimates suggest a flat output Phillips curve with a slope of 0.027.

Our findings contribute to the ongoing debate about slope of the Phillips curve and suggest that the flattening of the Phillips curve is the result of a lack of meaningful cyclicality in real wages, while the "capacity pressure" channel contributes more than a textbook calibration of the New Keynesian model would predict. Our work highlights the possibilities of using microdata and credible identification strategies to understand aggregate price dynamics.

References

Amiti, Mary, Oleg Itskhoki, and Jozef Konings, "International Shocks, Variable Markups, and Domestic Prices," *Review of Economic Studies*, February 2019.

Anderson, T.W. and Cheng Hsiao, "Formulation and Estimation of Dynamic Models Using Panel Data," *Journal of Econometrics*, January 1982, *18* (1), 47–82.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, December 2008.

Arellano, Manuel and Stephen Bond, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economic Studies*, 1991, *58* (2), 277.

Barnichon, Regis and Geert Mesters, "Identifying Modern Macro Equations with Old Shocks," *Quarterly Journal of Economics*, November 2020, *135* (4), 2255–2298.

__ **and** __ , "The Phillips Multiplier," *Journal of Monetary Economics*, January 2021, *117*, 689–705.

Bilbiie, Florin O., "The New Keynesian Cross," *Journal of Monetary Economics*, October 2020, *114*, 90–108.

Boehm, Christoph E. and Nitya Pandalai-Nayar, "Convex Supply Curves," *American Economic Review*, December 2022, *112* (12), 3941–3969.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, January 2022, *89* (1), 181–213.

Cantore, Cristiano, Paul Levine, Joseph Pearlman, and Bo Yang, "CES Technology and Business Cycle Fluctuations," *Journal of Economic Dynamics and Control*, December 2015, *61*, 133–151.

Dedola, Luca, Kristoffersen, Mark Strøm, and Zuellig, Gabriel, "Price Synchronization and Cost Pass-Through in Multiproduct Firms: Evidence from Danish Producer Prices," Working Paper July 2019.

Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens, "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," June 26 2024. Available at SSRN: `https://ssrn.com/abstract=4481590`. (Previous version: NBER Working Paper, June 2023, No. 31382.).

Gali, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press, 2008.

__ **and Tommaso Monacelli**, "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, July 2005, *72* (3), 707–734.

Garin, Andrew and Filipe Silvério, "How Responsive Are Wages to Firm-Specific Changes in Labour Demand? Evidence from Idiosyncratic Export Demand Shocks," *Review of Economic Studies*, May 2024, *91* (3), 1671–1710.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 44–76.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, August 2020, *110* (8), 2586–2624.

Hansen, Christian, Jerry Hausman, and Whitney Newey, "Estimation With Many Instrumental Variables," *Journal of Business & Economic Statistics*, October 2008, *26* (4), 398–422.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from U.S. States," *Quarterly Journal of Economics*, August 2022, *137* (3), 1299–1344.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, June 2014, *104* (6), 1597–1629.

Inoue, Atsushi and Gary Solon, "Two-Sample Instrumental Variables Estimators," *Review of Economics and Statistics*, August 2010, *92* (3), 557–561.

Jarocinski, Marek and Peter Karadi, "Deconstructing Monetary Policy Surprises— The Role of Information Shocks," *American Economic Journal: Macroeconomics*, April 2020, *12* (2), 1–43.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, March 2005, *95* (1), 161–182.

Kimball, Miles S., "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, November 1995, *27* (4), 1241.

Klenow, Peter J. and Jonathan L. Willis, "Real Rigidities and Nominal Price Changes," *Economica*, 2016, *83* (331), 443–472.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock, "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 2014, *52* (1), 124–188.

McLeay, Michael and Silvana Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 2019, *34*, 199–255.

Mikusheva, Anna and Liyang Sun, "Weak Identification with Many Instruments," *The Econometrics Journal*, June 2024, *27* (2), C1–C28.

Nakamura, Emi and Jón Steinsson, "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect*," *The Quarterly Journal of Economics*, August 2018, *133* (3), 1283–1330.

Pacini, David and Frank Windmeijer, "Robust Inference for the Two-Sample 2SLS Estimator," *Economics Letters*, September 2016, *146*, 50–54.

United Nations Statistics Division, "Correlation and Conversion Tables Used in UN Comtrade," August 2022.

Appendix

A Additional Empirical Results

Table 4: Effect on output — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Anderson-Hsiao	(4) Firm FE	(5) 4d-nace FE	(6) No sector FE
t	0.50***	0.55***	0.46***	0.46***	0.33**	0.66***
	(0.091)	(0.082)	(0.13)	(0.12)	(0.16)	(0.100)
t+1	0.60***	0.66***	0.50***	0.55***	0.51***	0.75***
	(0.12)	(0.11)	(0.16)	(0.13)	(0.19)	(0.12)
t+3	0.28*	0.34**	0.22	0.35**	0.30	0.48***
	(0.15)	(0.14)	(0.22)	(0.17)	(0.26)	(0.15)
t+5	0.12	0.29*	-0.017	0.12	0.19	0.30
	(0.20)	(0.18)	(0.30)	(0.16)	(0.31)	(0.18)
Firms	747	785	735	704	690	751
N	7,904	9,558	7,636	7,861	7,068	7,925
F	9.083	44.828	3.785	6.139	2.184	11.022

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline specification. (2) Baseline specification without controls. (3) Anderson-Hsiao estimator using lagged levels as instruments for lagged growth rates. (4) Adds a firm FE that controls for trends. (5) Replaces 2d-sector X time FE with 4d-sector X time FE. (6) Replaces 2d-sector X time FE with time FE.

Table 5: Effect on output — different sample restrictions

	(1) Baseline	(2) All PPI	(3) All MFG	(4) More balanced	(5) Higher exports	(6) Pre 2020
t	0.50***	0.55***	0.55***	0.43***	0.42***	0.38***
	(0.091)	(0.11)	(0.096)	(0.13)	(0.13)	(0.14)
t+1	0.60***	0.49***	0.50***	0.57***	0.57***	0.58***
	(0.12)	(0.13)	(0.12)	(0.16)	(0.16)	(0.18)
t+3	0.28*	0.38**	0.42**	0.42**	0.44**	0.44**
	(0.15)	(0.17)	(0.17)	(0.19)	(0.20)	(0.20)
t+5	0.12	0.19	0.26	0.19	0.22	0.22
	(0.20)	(0.21)	(0.19)	(0.23)	(0.23)	(0.23)
Firms	747	825	1,898	654	581	573
N	7,904	8,955	15,141	7,193	5,951	5,311
F	9.083	9.235	9.749	3.314	3.158	2.630

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline sample. (2) Includes all firms that appear in the PPI w/o further restrictions. (3) Includes all mfg. firms that fulfill sample restrictions (also those not in the PPI). (4) Adds restriction of >=10 years uninterrupted activity. (5) Adds restrictions of export share >0.33 in the previous year. (6) Restricts sample to the pre-2020 period.

Table 6: Effect on sales — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Anderson-Hsiao	(4) Firm FE	(5) 4d-nace FE	(6) No sector FE
t	0.69***	0.74***	0.59***	0.66***	0.45***	0.90***
	(0.081)	(0.072)	(0.078)	(0.074)	(0.090)	(0.068)
t+1	0.80***	0.80***	0.60***	0.68***	0.56***	0.95***
	(0.10)	(0.092)	(0.10)	(0.092)	(0.12)	(0.089)
t+3	0.51***	0.52***	0.32**	0.30***	0.29*	0.60***
	(0.13)	(0.13)	(0.13)	(0.10)	(0.15)	(0.11)
t+5	0.34**	0.42***	0.10	0.18	0.17	0.47***
	(0.17)	(0.16)	(0.18)	(0.12)	(0.21)	(0.14)
Firms	815	848	819	783	777	821
N	9,678	11,880	9,607	9,726	8,880	9,774
F	14.118	105.990	18.051	20.411	8.888	40.516

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline specification. (2) Baseline specification without controls. (3) Anderson-Hsiao estimator using lagged levels as instruments for lagged growth rates. (4) Adds a firm FE that controls for trends. (5) Replaces 2d-sector X time FE with 4d-sector X time FE. (6) Replaces 2d-sector X time FE with time FE.

Table 7: Effect on sales — different sample restrictions

	(1) Baseline	(2) All PPI	(3) All MFG	(4) More balanced	(5) Higher exports	(6) Pre 2020
t	0.69***	0.74***	0.60***	0.61***	0.61***	0.58***
	(0.081)	(0.068)	(0.056)	(0.074)	(0.074)	(0.078)
t+1	0.80***	0.74***	0.66***	0.69***	0.67***	0.66***
	(0.10)	(0.087)	(0.070)	(0.096)	(0.099)	(0.10)
t+3	0.51***	0.53***	0.53***	0.42***	0.44***	0.44***
	(0.13)	(0.10)	(0.093)	(0.12)	(0.12)	(0.12)
t+5	0.34**	0.33**	0.49***	0.31*	0.36**	0.36**
	(0.17)	(0.14)	(0.12)	(0.17)	(0.17)	(0.17)
Firms	815	912	2,291	712	627	618
N	9,678	11,198	20,340	8,867	7,251	6,501
F	14.118	26.765	31.306	18.862	18.899	16.131

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline sample. (2) Includes all firms that appear in the PPI w/o further restrictions. (3) Includes all mfg. firms that fulfill sample restrictions (also those not in the PPI). (4) Adds restriction of >=10 years uninterrupted activity. (5) Adds restrictions of export share >0.33 in the previous year. (6) Restricts sample to the pre-2020 period.

Table 8: Effect on capacity utilization — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Anderson-Hsiao	(4) Firm FE	(5) 4d-nace FE	(6) No sector FE
t	0.19***	0.23***	0.20***	0.26***	0.13	0.22***
	(0.063)	(0.055)	(0.075)	(0.068)	(0.098)	(0.054)
t+1	0.22***	0.24***	0.15	0.24***	0.27*	0.21***
	(0.079)	(0.078)	(0.10)	(0.082)	(0.14)	(0.062)
t+3	0.0050	-0.061	0.043	0.084	0.014	-0.032
	(0.090)	(0.092)	(0.12)	(0.089)	(0.19)	(0.089)
t+5	0.038	0.086	0.13	0.17	-0.11	0.11
	(0.11)	(0.11)	(0.15)	(0.12)	(0.20)	(0.090)
Firms	410	449	389	366	346	413
N	3,249	3,951	3,015	3,202	2,464	3,309
F	27.201	17.701	8.884	30.811	15.578	32.297

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline specification. (2) Baseline specification without controls. (3) Anderson-Hsiao estimator using lagged levels as instruments for lagged growth rates. (4) Adds a firm FE that controls for trends. (5) Replaces 2d-sector X time FE with 4d-sector X time FE. (6) Replaces 2d-sector X time FE with time FE.

Table 9: Effect on capacity utilization — different sample restrictions

	(1) Baseline	(2) All PPI	(3) All MFG	(4) More balanced	(5) Higher exports	(6) Pre 2020
t	0.19***	0.25***	0.19***	0.18***	0.16**	0.14*
	(0.063)	(0.057)	(0.057)	(0.070)	(0.074)	(0.079)
t+1	0.22***	0.18***	0.12	0.16*	0.13	0.16
	(0.079)	(0.069)	(0.076)	(0.089)	(0.095)	(0.10)
t+3	0.0050	0.048	0.083	-0.021	-0.027	-0.027
	(0.090)	(0.085)	(0.085)	(0.098)	(0.10)	(0.10)
t+5	0.038	0.020	0.035	-0.031	-0.097	-0.097
	(0.11)	(0.11)	(0.12)	(0.13)	(0.14)	(0.14)
Firms	410	450	653	373	331	305
N	3,249	3,636	4,198	2,983	2,425	2,113
F	27.201	24.137	32.580	25.189	22.008	19.521

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline sample. (2) Includes all firms that appear in the PPI w/o further restrictions. (3) Includes all mfg. firms that fulfill sample restrictions (also those not in the PPI). (4) Adds restriction of >=10 years uninterrupted activity. (5) Adds restrictions of export share >0.33 in the previous year. (6) Restricts sample to the pre-2020 period.

Table 10: Reduced form effect on prices — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Anderson-Hsiao	(4) Firm FE	(5) 4d-nace FE	(6) No sector FE
t	0.13***	0.19***	0.16**	0.11**	0.098*	0.14***
	(0.049)	(0.044)	(0.061)	(0.053)	(0.055)	(0.054)
t+1	0.20***	0.23***	0.27***	0.15***	0.21***	0.11**
	(0.051)	(0.057)	(0.074)	(0.052)	(0.081)	(0.048)
t+3	0.26***	0.31***	0.35***	0.12	0.25**	0.21***
	(0.084)	(0.073)	(0.094)	(0.078)	(0.12)	(0.078)
t+5	0.21**	0.36***	0.40***	0.062	0.065	-0.026
	(0.10)	(0.097)	(0.15)	(0.083)	(0.15)	(0.14)
Firms	590	716	584	582	582	590
N	17,175	25,452	17,160	17,167	17,024	17,178
F	7.541	17.748	1.802	9.062	5.557	7.307

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline specification. (2) Baseline specification without controls. (3) Anderson-Hsiao estimator using lagged levels as instruments for lagged growth rates. (4) Adds a firm FE that controls for trends. (5) Replaces 2d-sector X time FE with 4d-sector X time FE. (6) Replaces 2d-sector X time FE with time FE.

Table 11: Reduced form effect on prices — sample restrictions

	(1) Baseline	(2) All PPI	(3) More balanced	(4) Higher exports	(5) Pre 2020
t	0.13***	0.14***	0.13***	0.11**	0.13**
	(0.049)	(0.046)	(0.050)	(0.050)	(0.055)
t+1	0.20***	0.21***	0.20***	0.17***	0.22***
	(0.051)	(0.047)	(0.051)	(0.051)	(0.052)
t+3	0.26***	0.30***	0.27***	0.20**	0.26***
	(0.084)	(0.078)	(0.085)	(0.081)	(0.084)
t+5	0.21**	0.24***	0.22**	0.17*	0.21**
	(0.10)	(0.091)	(0.10)	(0.099)	(0.10)
Firms	590	649	545	510	570
N	17,175	19,603	16,716	13,879	14,803
F	7.541	3.656	7.643	7.146	10.776

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline sample. (2) Includes all firms that appear in the PPI w/o further restrictions. (3) Includes all mfg. firms that fulfill sample restrictions (also those not in the PPI). (4) Adds restriction of >=10 uninterrupted observations per firm. (5) Adds restrictions of export share > 0.33 in the previous year. (6) Restricts sample to the pre-2020 period.

Table 12: IV effect on prices — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Firm FE	(4) 4d-nace FE	(5) No sector FE
t	0.28**	0.35***	0.24*	0.32	0.32***
	(0.14)	(0.12)	(0.13)	(0.25)	(0.10)
t+1	0.39**	0.44***	0.30**	0.48	0.27***
	(0.16)	(0.15)	(0.14)	(0.35)	(0.090)
t+3	0.56**	0.58***	0.41**	0.68	0.41***
	(0.23)	(0.19)	(0.20)	(0.50)	(0.13)
t+5	0.48*	0.69***	0.23	0.33	0.24*
	(0.25)	(0.23)	(0.21)	(0.49)	(0.14)
Firms	549	716	539	542	549
N	16,779	25,452	16,769	16,629	16,782
First-stage F	8.088	45.266	14.894	3.630	18.326

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline specification. (2) Baseline specification without controls. (4) Adds a firm FE that controls for trends. (5) Replaces 2d-sector X time FE with 4d-sector X time FE. (6) Replaces 2d-sector X time FE with time FE.

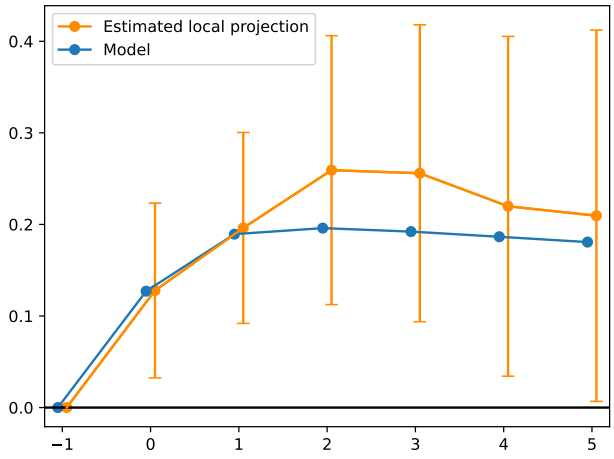
Table 13: IV effect on prices — sample restrictions

	(1) Baseline	(2) All PPI	(3) All MFG	(4) More balanced	(5) Higher exports	(6) Pre 2020
t	0.28**	0.26**	0.29**	0.29*	0.26*	0.27*
	(0.14)	(0.11)	(0.13)	(0.15)	(0.14)	(0.15)
t+1	0.39**	0.36***	0.40***	0.38**	0.36**	0.44**
	(0.16)	(0.13)	(0.15)	(0.16)	(0.16)	(0.19)
t+3	0.56**	0.51***	0.57***	0.56**	0.50**	0.58**
	(0.23)	(0.18)	(0.22)	(0.24)	(0.23)	(0.25)
t+5	0.48*	0.41**	0.49**	0.51*	0.46*	0.49*
	(0.25)	(0.19)	(0.25)	(0.27)	(0.25)	(0.26)
Firms	549	606	549	502	478	535
N	16,779	19,203	16,779	16,321	13,667	14,594
First-stage F	8.088	11.912	13.871	6.798	7.774	6.776

Notes: SE in parenthesis are clustered at the firm level. * p<0.1 ** p<0.05 *** p<0.01. Model summary statistics are reported for horizon h=0. (1) Baseline sample. (2) Includes all firms that appear in the PPI w/o further restrictions. (3) Includes all mfg. firms that fulfill sample restrictions (also those not in the PPI). (4) Adds restriction of >=10 years uninterrupted activity. (5) Adds restrictions of export share >0.33 in the previous year. (6) Restricts sample to the pre-2020 period.

B Additional Model Results

(a) Impulse response of prices



(b) Impulse response of output

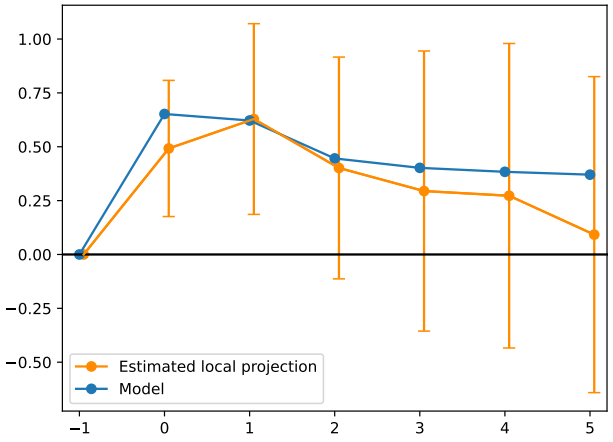


Figure 7: Counterfactual responses to a firm-level demand shocks with standard Cobb-Douglas production and CES demand ($\psi = 0$, $\alpha = 0.33$, $\sigma = 4$ and $\Gamma = 0$)

C Appendix: Identification with strategic complementarities and pricing-to-market

In this section, we discuss our estimation strategy in a more general setting in which firms are able to price discriminate between customers in different markets, and face a market demand curve that features strategic complementarities in price-setting. Like in the baseline case for the static model, we use a Cobb-Douglas production function $Q_i = A_h V_i L_i^{1-\alpha}$. Like in our extended NK model below, we use the Klenow and Willis (2016) specification of the Kimball (1995) aggregator that results in the following demand curve:

$$Q_{i,k} = Z_k \Gamma_{i,k} \Upsilon \left(\frac{P_{i,k}}{\overline{P}_k} D_k \right) \text{ with } \Upsilon(x) = \left(1 - \tau \log \left(\frac{\sigma}{\sigma - 1} x \right) \right)^{\sigma/\tau} \tag{24}$$

This leads to a supply curve with markups over marginal cost that vary based on local conditions:

$$P_{i,k} = \frac{\epsilon_{i,k}}{\epsilon_{i,k} - 1} MC_i \quad \text{with} \quad \epsilon_{i,k} = \frac{\sigma}{1 - \tau \log \left(\frac{\sigma}{\sigma - 1} x_{i,k} \right)} \tag{25}$$

Where the elasticity of markups w.r.t. relative prices $\Gamma_{i,k}$ is defined as:

$$\Gamma_{i,k} = \frac{\tau}{\sigma - 1 + \tau \log \left(\frac{\sigma}{\sigma - 1} x_{i,k} \right)} \tag{26}$$

We consider a case with no ex-ante heterogeneity in relative prices between firms. In this case $P_{i,k} = \overline{P}_k$, $D_k = (\sigma - 1)/\sigma$ and $\Gamma_{i,k} = \tau/(\sigma - 1)$ in all markets. We take logs and differentiate to get:

$$\Delta p_{i,k} = \frac{\alpha}{(1-\alpha)(1+\Gamma)} \Delta q_i + \frac{\Gamma}{1+\Gamma} \Delta \overline{p}_k + \frac{1}{1+\Gamma} \Delta w_h - \frac{\Delta a_h + \Delta v_i}{(1-\alpha)(1+\Gamma)} \qquad \text{Inverse supply} \tag{27}$$

$$\Delta q_i = -\sigma \sum_{k=1}^{K} \gamma_{i,k} (\Delta p_{i,k} - \Delta \overline{p}_k) + \Delta \tilde{z}_i \qquad \text{Total demand} \tag{28}$$

Note that even though firms respond to the prices of local competitors, the slope of their supply curve is the same in each market. Combining equations (27) and (28) we can express firms' prices and quantities sold in each market

as a function of exogenous variables:

$$\Delta p_{i,k} = \frac{\Gamma}{1+\Gamma}\Delta\overline{p}_k + \frac{\alpha}{(1-\alpha)(1+\Gamma)+\alpha\sigma}\Delta\tilde{z}_i + \frac{\Delta a_h + \Delta v_i - (1-\alpha)\Delta w_h}{(1-\alpha)(1+\Gamma)+\alpha\sigma} + \frac{\sigma\alpha}{(1+\Gamma)((1-\alpha)(1+\Gamma)+\alpha\sigma)}\Delta\tilde{\overline{p}}_i \tag{29}$$

and

$$\Delta q_i = \frac{(1-\alpha)(1+\Gamma)}{(1-\alpha)(1+\Gamma)+\alpha\sigma}\Delta\tilde{z}_i + \frac{\sigma(1-\alpha)}{(1-\alpha)(1+\Gamma)+\alpha\sigma}\Delta\tilde{\overline{p}}_i + \frac{\sigma(\Delta a_h + \Delta v_i - (1-\alpha)\Delta w_h)}{(1-\alpha)(1+\Gamma)+\alpha\sigma} \tag{30}$$

Firms charge different prices whenever local price levels $\Delta\overline{p}_k$ develop differently in different locations. In our empirical analysis, we consider the response of domestic prices $\Delta p_{i,h}$. This means that the domestic pricing-to-market term $\Gamma/(1+\Gamma)\Delta\overline{p}_h$ doesn't vary between firms in a sector and is absorbed in time fixed effects. A reduced form regression of domestic prices $\Delta p_{i,h}$ on $\Delta\tilde{z}_i$ identifies:

$$\beta_{\tilde{z}} = \frac{\alpha}{(1-\alpha)(1+\Gamma)+\alpha\sigma}\left(1 + \frac{\sigma}{(1+\Gamma)}\frac{COV(\Delta\tilde{z}_i, \Delta\tilde{\overline{p}}_i)}{VAR(\Delta\tilde{z}_i)}\right).$$

IV estimation normalizes with the effect on output and recovers the slope of the supply curve:

$$\beta_{\tilde{q}}^{IV} = \frac{COV(\Delta p_i, \Delta\tilde{z}_i)}{COV(\Delta q_i, \Delta\tilde{z}_i)} = \frac{\alpha}{(1-\alpha)(1+\Gamma)}$$

Our example covers a case with no ex-ante heterogeneity in prices and thus imposes a pass-through that is the same for all firms, but ex-ante heterogeneity might be present in the data. In this case, the elasticity of markups Γ would vary between firms, and the slope of supply curves would depend on initial relative prices. In this case, our estimation recovers an average supply curve.

The extension of the baseline model to a more general setting with pricing-to-market and strategic complementarities also shows that it is preferable to estimate the response of domestic prices to a weighted export demand shock, rather than estimate the response of prices in a given destination to variation in local demand. To do so, one would need to control for the prices of local competitors in each destination. Moreover, even though this is not explicit in our model, local supply factors (such as distribution cost) could be a part of local prices that would vary with local shocks.

D Appendix: New Keynesian model with Kimball Demand

In this section, we describe a New Keynesian model with idiosyncratic shocks and Kimball demand. This model nests the model with CES demand described in the main text, which corresponds to the limit case in which the super-elasticity parameter τ of the Kimball aggregator approaches 0 and consequently the elasticity of markups w.r.t. relative prices $\Gamma = 0$.

Final goods production. Like in the CES model, the final good is produced by perfectly competitive producers with flexible prices. Final goods producers turn intermediate goods into the final good using the Kimball aggregator Υ:

$$\int_i Z_{i,t} \Upsilon \left(\frac{Y_{i,t}}{Z_{i,t} Y_t} \right) di = 1. \tag{31}$$

We use the Klenow and Willis (2016) specification of Kimball demand.[12] $Z_{i,t}$ is a demand shifter for intermediate products that follows an AR(1) process with persistence ρ and mean zero in logarithms. Final goods producers minimize their expenditure $\int_i Y_{i,t} P_{i,t} di$ subject to (31):

$$\min_{\{Y_{i,t}\}} \int_i P_{i,t} Y_{i,t} di \quad \text{s.t.} \int_i Z_{i,t} \Upsilon \left(\frac{Y_{i,t}}{Z_{i,t} Y_t} \right) di = 1. \tag{32}$$

This results in the Kimball demand curve for intermediate goods:

$$Y_{i,t} = Y_t Z_{i,t} \Psi \left(\frac{P_{i,t}}{P_t} D_t \right) \text{ with } \Psi(x) = \left(1 - \tau \log \left(\frac{\sigma}{\sigma - 1} x \right) \right)^{\sigma/\tau}, \tag{33}$$

P_t is the cost-minimizing price index for intermediates and $D_t = \int_i Y_{i,t}/Y_t \Upsilon'(Y_{i,t}/(Y_t Z_{i,t})) di$ summarizes heterogeneity of intermediate producers. The demand elasticity ϵ is defined as:

$$\epsilon(x_{i,t}) = \frac{\sigma}{1 - \tau \log \left(\frac{\sigma}{\sigma - 1} x_{i,t} \right)}, \tag{34}$$

[12] This corresponds to $\Upsilon(x) = 1 + \frac{\sigma - 1}{\tau} \tau^{\sigma/\tau} e^{1/\tau} \left(\Gamma(\frac{\sigma}{\tau}, \frac{1}{\tau}) - \Gamma(\frac{\sigma}{\tau}, \frac{x^{\tau/\sigma}}{\tau}) \right)$, where Γ is the incomplete Gamma function.

and the inverse elasticity of markups w.r.t. the relative price of a firm is given by

$$\Gamma(x_{i,t}) = \frac{\tau}{\sigma - 1 + \tau \log\left(\frac{\sigma}{\sigma-1} x_{i,t}\right)}. \tag{35}$$

Note that the final output price index P_t implied by the cost function of final goods producers is in general not the same as the competitor price index P_t/D_t considered by intermediate producers when setting prices. However, the two are equal in a linear approximation around a symmetric steady-state[13] and we will use them interchangeably below. In the CES special case of $\tau \to 0$, D_t is a equal to $(\sigma - 1)/\sigma$ the demand curve reduces to $Y_{i,t} = Y_t Z_{i,t} (P_{i,t}/P_t)^{-\sigma}$ with constant elasticity σ.

Intermediate goods production. Like in the baseline case with CES demand, intermediates are produced using a normalized CES production function:

$$Y_t = \left(\alpha \left(\frac{K_t}{K_{SS}} \right)^{\psi} + (1 - \alpha) \left(\frac{L_t}{L_{SS}} \right)^{\psi} \right)^{1/\psi} \tag{36}$$

Note that this production function converges to a standard Cobb-Douglas specification as $\psi \to 0$, to a Leontief production function as $\psi \to -\infty$, and to a linear production function as $\psi \to 1$. Given this production function, the slope of the flexible-price supply curve is now equal to:

$$\delta^{\mathrm{K}} = \frac{(1 - \psi)\alpha}{(1 + \Gamma)(1 - \alpha)} \tag{37}$$

Intermediate producers can reset their price with probability $1 - \theta$ and discount the future at rate β. Firms maximize their future discounted profit whenever they have an opportunity to reset their price to a new optimal price P_t^*. For any specification of demand and production functions, this yields the first-order condition:

$$\sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(Y_{i,t+k}(1 - \epsilon_{i,t+k}) \left(P_{i,t}^* - \frac{\epsilon_{i,t+k}}{\epsilon_{i,t+k} - 1} MC_{i,t+k} \right) \middle| Z_{i,t} \right) = 0 \tag{38}$$

, where output $Y_{i,t+k}$, the demand elasticity $\epsilon_{i,t+k}$ and marginal cost $MC_{i,t+k}$ are functions of its optimal reset price and ultimately the realization of idiosyncratic demand shocks and aggregate state variables.

We approximate price-setting firms' FOC using the demand and supply specifications from above around a zero-

[13] See also Amiti et al. (2019), Appendix D.

inflation steady state in which all idiosyncratic demand shocks are equal to their mean. This implies there is no heterogeneity in prices or quantities in the steady state, like in a New Keynesian model without firm-specific demand shocks. This simplification is important for the Kimball case, as it allows us to derive analytic expressions for the aggregate Phillips curve and the response of prices to firm-specific demand shocks. In such a steady-state, $D_{SS} = (\sigma-1)/\sigma$, $\epsilon(x_{SS}) = \sigma$ and $\Gamma(x_{SS}) = \tau/(\sigma-1)$. We divide the FOC by P_{t-1} to express it in terms of variables with constant steady-state values:

$$\sum_{k=0}^{\infty}(\beta\theta)^k E_t\left(Y_{i,t+k}(1-\epsilon_{i,t+k})\left(\frac{P^*_{i,t}}{P_{t-1}} - \frac{\epsilon_{i,t+k}}{\epsilon_{i,t+k}-1}\frac{MC_{i,t+k}}{P_{t+k}}\Pi_{t+k,t-1}\right)\middle| Z_{i,t}\right) = 0. \tag{39}$$

We then log-linearize. We use lower-case letters to denote log deviations of a variable from its steady state value and use the fact that $\epsilon_{i,t+k}/(\epsilon_{i,t+k}-1) \approx \mu(1-\Gamma(p^*_{i,t}-p_t))$, where $\mu = \sigma/(\sigma-1)$ is the steady state markup. This yields the optimal reset price as a function of expected nominal marginal cost and the aggregate price level:

$$p^*_{i,t} = (1-\beta\theta)\sum_{k=0}^{\infty}(\beta\theta)^k E_t\left(\frac{1}{1+\Gamma}(mc^R_{i,t+k}+p_{t+k}) + \frac{\Gamma}{1+\Gamma}p_{t+k}\middle| z_t\right) \tag{40}$$

Analogous to the static case with flexible prices, reset prices with Kimball demand are a weighted average of discounted expected future nominal marginal cost and competitor prices. We use the production function and demand curve to derive an expression for the deviation of marginal cost from its steady state value, which is equal to $mc^R_{i,t+k} = w^R_{t+k} + (1-\psi)\alpha/(1-\alpha)(z_{i,t+k} - \sigma(p^*_{i,t}-p_{t+k}))$. Plugging in and simplifying, we can express the reset price as a function of aggregate marginal cost, idiosyncratic demand and inflation expectations:

$$p^*_{i,t} - p_{t-1} = \frac{1-\beta\theta}{(1+\Gamma)(1+\sigma\delta^K)}\sum_{k=0}^{\infty}(\beta\theta)^k\left(E_t\left(mc^R_{t+k}\right) + (1+\Gamma)\delta^K\rho^k z_{i,t}\right) + \sum_{k=0}^{\infty}(\beta\theta)^k E_t(\pi_{t+k}), \tag{41}$$

Aggregate dynamics. The model follows the same aggregate dynamics as a textbook model without idiosyncratic demand shocks. Around a zero inflation symmetric steady state, aggregate inflation can be approximated as $\pi_t = (1-\theta)\left(\int_i p^*_{i,t}di - p_{t-1}\right)$. We combine this definition with equation (41) to derive the marginal cost Phillips curve:

$$\pi_t = \lambda mc^R_t + \beta E_t(\pi_{t+1}) \tag{42}$$

, where $\lambda = (1-\theta)(1-\beta\theta)/(\theta(1+\Gamma)(1+\sigma\delta^K))$. We can use the definition of deviations of marginal cost from their steady-state value $mc_t = w^R_t + (1-\psi)\alpha/(1-\alpha)y_t$ to go from equation (42) and the reduced form wage

Phillips curve discussed in the main text $w_t^R = \phi y_t$, to get the output Phillips curve:

$$\pi_t = \kappa y_t + \beta E_t(\pi_{t+1}) \tag{43}$$

Like in our CES baseline model, the output Phillips curve can be decomposed in a "capacity pressure" and a "wage pressure" channel:

$$\kappa = \underbrace{\frac{(1-\theta\beta)(1-\theta)}{\theta}\frac{\delta^K}{1+\sigma\delta^K}}_{\text{Capacity pressure}\equiv\kappa^p} + \underbrace{\frac{(1-\theta\beta)(1-\theta)}{\theta}\frac{\phi}{(1+\sigma\delta^K)(1+\Gamma)}}_{\text{Wage pressure}\equiv\kappa^w}. \tag{44}$$

The capacity pressure channel is given by the same expression as in the CES baseline, but the flexible-price supply curve δ^K now reflects a different combination of underlying parameters due to Kimball demand. The wage pressure channel is muted because with strategic complementarity, firms who update their price do not pass-through wage increases to the same extent as in the CES case. This highlights the result that as long as the model is fit to match the empirical price response to demand shocks, it will produce the same capacity pressure channel, no matter the underlying values of Γ and ψ.

Response to idiosyncratic shocks Using equation (41) and the definition of the firm-level demand process, we can express the reset price of a firm with demand realization z relative to the average reset price as just like in the CES case:

$$p_t^*(z) - \int p_t^*(z)dF(z) = \frac{1-\beta\theta}{1-\beta\theta\rho}\frac{\delta^K}{1+\sigma\delta^K}z \tag{45}$$

To get the average period $t+k$ relative price of firms hit by demand shock z_t in period t, we average over those that reset their price and those that don't:

$$p_{t+k}(z_t) - p_{t+k} = (1-\theta)(p_{t+k}^*(z_t) - p_{t+k}^*) + \theta(p_{t+k-1}(z_t) - p_{t+k-1})) \tag{46}$$

Recall that the period $t+k-1$ average price conditional on period t demand, $p_{t+k-1}(z_t)$, will not be equal to the unconditional average price p_{t+k-1} because prices might have previously been updated in response to the time t demand shock. This is only true at horizon $k=0$, because demand shocks in period t don't correlate with period $t-1$ prices. We iterate equation (46) backward and use the fact that $p_{t-1}(z_t) = p_{t-1}$. The average period $t+k$ price

of firms hit with demand shock z_t in period t is equal to:

$$p_{t+k}(z_t) - p_{t+k} = \frac{(1-\beta\theta)(1-\theta)}{1-\beta\theta\rho} \frac{\theta^{k+1} - \rho^{k+1}}{\theta - \rho} \frac{\delta^K}{1+\sigma\delta^K} z_t, \tag{47}$$

Given equation (47), the corresponding response of average relative output to firm-level demand shocks is determined by log-linearizing the demand curve for intermediates given by equation (9):

$$y_{i,t+k}(z_t) - y_{t+k} = z_{i,t} - \sigma(p_{t+k}(z_t) - p_{t+k}) \tag{48}$$

Overall, the model mostly follows the same dynamics with Kimball and CES demand at the aggregate and at the firm level. The main difference is that the flexible-price supply curve δ^K reflects incomplete pass-through of marginal cost changes governed by Γ.

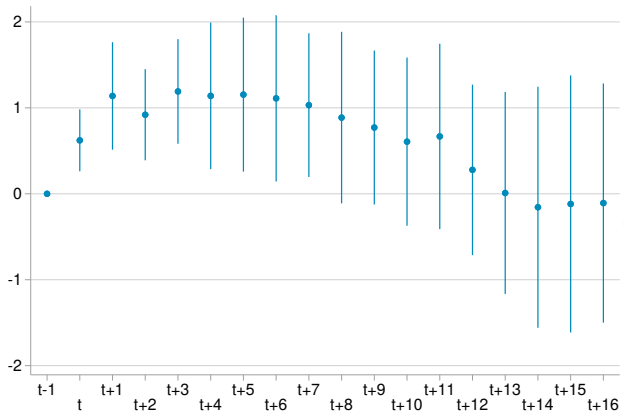
E Appendix: Response of real wages to an aggregate demand shock

In this section, we describe how we estimate the response of real wages to a cyclical demand shock. We use quarterly average wages in manufacturing deflated by the CPI as our measure of real wages. In addition, we use data on real Danish GDP and ECB monetary policy surprises cleaned from information effects estimated in Jarocinski and Karadi (2020). Because the Danish krone is pegged to the Euro, Danmarks Nationalbank usually closely follows ECB policy. Because ECB policy does not aim to offset Danish aggregate demand (which of course correlates with Euro area aggregate demand), the shocks should arguably produce a stronger effect in Denmark than in the Euro area. We aggregate monetary policy surprises to quarters by summing up all observations within a quarter and scale the shock to produce a unit effect on real GDP after four quarters. We then estimate the following local projection:

$$\Delta^h y_t = \beta^h u_t + \sum_{k=1}^{4} u_{t-k} + \sum_{k=1}^{4} \Delta y_{t-k} \tag{49}$$

with real wages and real GDP as outcomes y and the monetary policy shock as u. The results are shown in Figure 8. Danish real GDP increases over one year, then stays elevated for three years and returns to its initial level. At the same time, Danish real wages decline with an elasticity of up to about -0.3. We take the ratio of the two coefficient vectors and calculate the average over the first 8 quarters where we observe a strong effect of output. This suggests an output elasticity of real wages of about -0.15.

(a) Danish real GDP



(b) Danish manufacturing real wage

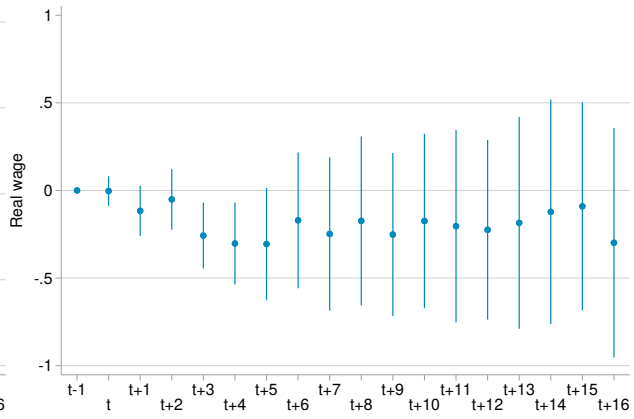


Figure 8: Effects of an ECB monetary policy surprise

F Appendix: Gagliardone et. al. (2024) specification with Danish data

Gagliardone et al. (2024) (henceforth GGLT) examine the firm-level response of prices to demand shocks.[14] Specifically, GGLT are interested in estimating the following empirical model, here presented using their notation:[15]

$$p_{ft} = \Psi^y \sigma^y y^n_{ft} + \theta p_{ft-1} + \alpha_f + \alpha_{s\times t} + \varepsilon^p_{ft} \tag{50}$$

, where p_{ft} is the price charged by firm f in period t and y^n_{ft} is the nominal firm-level output. α_f and $\alpha_{s\times t}$ are firm and sector-by-time fixed effects (2 Digit NACE Code).[16]

Estimation of $\Psi^y \sigma^y$ using OLS is biased because p_{ft} and y^n_{ft} are determined simultaneously and directly affected by idiosyncratic demand and supply shocks contained in ε^p_{ft}. To address this issue, GGLT use industry-specific demand shocks triggered by monetary policy surprises as an instrument. They estimate industry-specific sensitivities to monetary surprises using the following equation:

$$y^n_{ft} = \alpha_f + \sum_i \psi_i D^i_f MS_{t-4} + \epsilon^m_{ft} \tag{51}$$

MS_{t-4} denotes the monetary surprises (lagged 4 quarters), ψ_i represents the industry-specific sensitivity of nominal output to monetary policy surprises, D^i_f is a binary variable indicating whether firm f belongs to industry i, and α_f is a firm fixed effect. They then construct an instrument using the estimates $\hat{\psi}_i$:

$$y^{IV}_{ft} = \hat{\psi}_i MS_{t-4} \tag{52}$$

y^{IV}_{ft} is used as an instrument for y^n_{ft} when estimating (50). In this specification, the constructed variable y^{IV}_{ft} is treated as a single instrument. This ignores the estimation uncertainty from surrounding $\hat{\psi}_i$ which is itself estimated in equation (51) and overstates of the strength of the first stage. In contrast, if we correctly treat $D^i_f MS_{t-4}$ as the instruments, then $\Psi^y \sigma^y$ is no longer just-identified by a single instrument but is instead over-identified by many (potentially weak) instruments which is associated with various estimation problems (see Hansen et al., 2008, Mikusheva and Sun, 2024, for reviews). To demonstrate how different the conclusions can be when monetary policy surprises interacted with industry dummies are properly treated as instruments, we replicate the exercise in GGLT using our data.

[14] See Section 7 in Gagliardone et al. (2024)

[15] See Model D in GGLT. They include the prices of a firm's competitors as a control. For the purpose of this exercise, we ignore this.

[16] Having already estimated most of the parameters that make up Ψ^y, the authors are primarily interested in pinning down σ^y. This distinction is irrelevant for the point of our current exercise, and we just focus on estimating $\Psi^y \sigma^y$.

For the monetary policy shocks, MS_{t-4}, we use ECB monetary policy surprises estimated in Jarocinski and Karadi (2020). As mentioned in Appendix E, the Danish krone is pegged to the Euro, and Danmarks Nationalbank usually closely follows ECB policy.

Column (2) of Table 14 presents the resulting 2SLS estimates of $\Psi^y \sigma^y$ when using y_{ft}^{IV} as the instrument as GGLT do—i.e. the first stage is given by:

$$y_{ft}^n = \beta y_{ft}^{IV} + \theta^{FS} p_{ft-1} + \alpha_f^{FS} + \alpha_{s\times t}^{FS}. \tag{53}$$

The F-statistic for this regression amounts to 37 and one might conclude that y_{ft}^{IV} is a powerful instrument. However, this ignores the fact that y_{ft}^{IV} is based on estimated parameters.

Column (3) of Table 14 presents the results from using a first-stage specification that includes the estimation of sectoral sensitivities in the IV estimation directly:

$$y_{ft}^n = \sum_i \psi_i D_f^i MS_{t-4} + \theta^{FS} p_{ft-1} + \alpha_f^{FS} + \alpha_{s\times t}^{FS} \tag{54}$$

The IV estimates and estimated standard errors will, for mechanical reasons, be practically identical using this method—the only difference comes from using the correct degrees of freedom. However, the first stage now appears critically weak, with an F-statistic of 1.2. With many weak instruments, it is well known that TSLS will be biased toward OLS Mikusheva and Sun (2024). In Column (4), we present estimates obtained using the Limited Information Maximum Likelihood (LIML) estimator.[17] LIML remains median-unbiased even when using many weak instruments. The estimates are indeed larger than the TSLS estimates. The estimated standard errors are considerably larger and the coefficient is not significant.[18]

[17] See (Mikusheva and Sun, 2024)

[18] Note that in Column (4) we also cluster standard errors at the product-firm level to account for autocorrelation in the residuals.

Table 14: Monetary Policy Surprises as Demand Shifters

	(1) OLS	(2) TSLS - Pre-Estimated	(3) TSLS	(4) LIML
Log Revenue	0.00309***	0.0163*	0.0163*	0.0302
	(0.000807)	(0.00961)	(0.00961)	(0.0220)
Log Price (Lagged 1 quarter)	0.909***	0.908***	0.908***	0.907***
	(0.00476)	(0.00486)	(0.00486)	(0.00524)
Constant	0.496***			
	(0.0286)			
N	127,463	127,463	127,463	127,463
Kleibergen-Paap F		37.361	1.188	1.329

Notes: SE in parenthesis are clustered at the Industry-Year level ((4) is additionally clustered at the Firm-Product level). * $p<0.1$ ** $p<0.05$ *** $p<0.01$. All specifications include Sector-Year fixed effects and Firm-Product fixed effects. Log prices and log revenues are winsorized at the 2.5th and 97.5th percentile. (1) OLS. (2) TSLS with heterogenous first stage-parameters estimated separately. (3) TSLS with monetary policy shocks interacted with industry dummies as instruments. (4) LIML with monetary policy shocks interacted with industry dummies as instruments.