Demand Shocks and Prices—Micro Evidence and Macro Implications

Christian Philip Hoeck
(University of Copenhagen and Danmarks Nationalbank)

Tobias Renkin
(Danmarks Nationalbank)

November 8, 2024

(PRELIMINARY DRAFT)

Abstract

We estimate the response of domestic prices and total output of Danish manufacturing firms to persistent firm level demand shocks resulting from heterogeneity in firms' exposure to different export destinations. Our results suggest supply curves are steep—a demand shock that increases output by 1% raises prices by 0.3%. We then augment the production side of a simple New Keynesian model with firm level demand shocks, and identify key parameters from matching the response of prices and output in the model to our estimates. We show that a model that can fit firm behavior in the cross-section produces a relatively steep Phillips curve in the aggregate. Our preferred estimate for the slope of the Phillips curve suggests that a monetary policy shock that increases output by 1% results in a 0.4% higher price level in the medium run.

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;

1 Introduction

The Phillips curve is an aggregate supply curve that determines the relationship between aggregate prices and the output gap. Its shape 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. The recent spike in inflation has led to some reconsideration of the "death" of the Phillips curve. In our view, there remains considerable uncertainty over the slope of the Phillips curve.

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 dynamic response of Danish manufacturing firms' prices and output to a plausibly exogenous cross-sectional demand shock and back out the slope of their supply curves. We then augment the production side of a New Keynesian model with persistent firm-level demand shocks, and fit the response of firms' relative prices and output to these shocks predicted by the model to match our empirical estimates. This identifies key parameters of the model and allows us to make predictions about the slope of the aggregate Phillips curve. Our analysis yields two key results. First, supply curves at the firm level are much steeper than implied by common parametrization of New Keynesian models. Second, this carries over to the aggregate—a model that matches firm-level behavior features a healthy aggregate Phillips curve that is substantially steeper than some recent estimates suggest. Our preferred estimate for the slope of the output Phillips curve is 0.11, which implies a moderately persistent monetary policy shock that increases output by 1% will raise the price level by 0.37% in the long run.

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 the dynamic response of prices and output to our 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 common supply shocks to Danish firms 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 sidesteps the most important identification issues that plague the estimation of the Phillips curve slope from aggregate data.

[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.

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 dynamics of wages and cost of capital. To make predictions about macroeconomic relationships, we combine our firm level estimates with a simple model of the variation 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 then fit the response of firms' relative prices and output to these shocks in the model to the impulse response we estimate from the data. This pins down key parameters of the model and allows us to make direct predictions about the component of the Phillips curve slope that arises from increasing supply curves at the firm level. The model does very well at fitting firm behavior. In the fitted model, the slope of firms' supply curve is substantially steeper than implied by common values of New Keynesian model parameters, while demand curves are flatter. In the aggregate, this contributes strongly to a steep Phillips curve. With additional estimates of the response of aggregate factor prices to aggregate demand shocks, we can make predictions about the overall slope of the Phillips curve as well. Our estimates suggest a slope of the output Phillips curve between 0.115 and 0.168 depending on our assumptions on factor price dynamics. This means that in the medium run, the aggregate price level increases by 0.36 to 0.53 percent in response to a moderately persistent monetary policy shock that increases output by 1%. Our model also suggests that the slope of the Phillips curve is dominated by the effects of increasing supply curves at the firm level—assuming real factor prices do not respond to aggregate demand shocks already gives us a Phillips curve slope of 0.119.

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 demand 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. This literature typically uses aggregate data and deals with inflation expectations using rational expectation assumptions with appropriate (typically internal) instruments. This vast literature is surveyed in detail in Mavroeidis et al. (2014). They conclude that the rational expectations 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.

The 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 a relatively steep Phillips curve. But deviations of monetary policy from policy rules are typically small, and using them as instruments provides 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 contributions is to use large, credible 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—even with a conservative approach to inference, the demand shocks we use move firm-level output with a t-statistic of about 5.6 (corresponding to an F statistic of $t^2 = 31.2$). The main disadvantage compared to estimating the slope of the Phillips curve directly in an aggregate setting is that we need to rely on additional assumptions embodied in our model to go from firm-level estimates to a macroeconomic relationship.

The paper most closely related to ours is Gagliardone et al. (2023), who use microdata on Belgian manufacturing firms to estimate the pass-through of current and expected future marginal cost into prices. This is equivalent to the marginal cost formulation of the Phillips curve, and pins down important Phillips curve parameters. However, it is not informative about the relationship between marginal cost and output, which is embodied in the slope of the supply curve, and is needed to get a relationship between output and prices. Gagliardone et. al. estimate this relationship using firm-level data and aggregate monetary policy shocks as instruments. Their results suggest that firms pass-through current and expected future marginal cost, but that marginal cost doesn't respond to changes in output, and that the Phillips curve is flat as a result. However, their estimates of the relationship between output and prices using monetary policy shocks as instruments are subject to the same concerns about weak identification highlighted in the aggregate Phillips curve literature[3]. Our main contribution relative to theirs is thus to identify firms' supply curves

[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).

[3] The problem is likely worse in their approach than in aggregate estimates, since they estimate one first-stage parameter per firm, resulting in a many-weak-instruments situation in which identification problems are typically worse than in just-identified IV with one weak instrument—see Hansen et al. (2008) or Mikusheva and Sun (2024) for an overview.

directly using a strong cross-sectional demand shock. This approach leads us to a quite different conclusion. We find a strong relationship between output and prices at the firm level which suggests a healthy Phillips curve in the aggregate.

Our paper is also very closely related to the literature estimating regional Phillips curves, most notably McLeay and Tenreyro (2019) and Hazell et al. (2022). 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 downward 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. There are several possible explanations for the discrepancy between their results and ours. First, our identification assumptions are weaker—Hazell et al. (2022) 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. A second possibility is that their instrument is valid but scaled incorrectly—the exogenous variation in overall state unemployment possibly doesn't represent a large shock to the non-tradeable sector. This would be the case, for example, if labor mobility between the tradeable and non-tradeable sector is limited, or if wages are partially set at the national level. Third, it is possible that prices of non-tradeable services simply respond differently to aggregate demand than prices of tradeable goods. In this case their estimates would underestimate and our estimates based on producer prices for manufacturing goods would overestimate the slope of the total aggregate 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 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 Combined Nomenclature (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 large 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[4]. 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

[4]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)

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 about 500 firms. Products are classified using 8-digit Harmonized System (HS) 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 strongly balanced, with very 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 code level. We construct leave-one-out 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.

Sample description Our baseline estimation sample covers manufacturing firms that report domestic prices in the PPI and report output in the VARS survey. Since all manufacturing firms with more than 10 employees participate in the VARS survey, 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 the cross the VARS survey 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 firm-level demand shifter to identify supply curves. We use demand shifters that arise from heterogeneity in exporting firms' exposure to aggregate fluctuations in different destination countries. This variation in exposure arises because firms persistently export their products to different destination countries, and aggregate demand fluctuations in these countries are not perfectly synchronized. We introduce a simple static framework describing the price-setting problem of a single firm to rigorously illustrate our strategy and identification assumptions involved and motivate the construction of the demand shifter we use in our empirical work. We keep this framework as simple possible for illustrative purposes, and show that our strategy is robust in more complicated settings in extensions in Appendix [To be written].

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:

$$Q_i = A_i L_i^{1-\alpha}.$$

Since the capital stock is fixed and $\alpha \in (0,1)$, firms operate with decreasing returns to scale in the short-run and the short-run supply curve is upward sloping. A_i 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, i.e. $A_i = A_h V_i$.

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 elasticity of demand ε in each market k, i.e. $Q_{i,k} = (P_i/\overline{P}_k)^{-\varepsilon} Z_k$. $\overline{P}_k$ is the price index of the firms' competitors in market k and Z_k is an aggregate demand shifter in market k. Profit maximization results in a constant markups over marginal cost price-setting policy. Taking logs and differentiating, we get the following supply-demand system that determines changes in prices and quantities:

$$\Delta p_i = \frac{\alpha}{1-\alpha}\Delta q_i + \Delta w_h - \frac{1}{1-\alpha}(\Delta a_h + \Delta v_i) \qquad \text{Inverse supply} \qquad (2)$$

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

The variable $\gamma_{i,k} = Q_{i,k}/Q_i$ measures the share of the firm's output sold in market k and 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 entirely 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:

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

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 will absorb the home market supply shifters Δw_h and Δa_h and all 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 not correlated with 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:

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

The coefficient is a mix of the slope of the inverse supply curve α and the slope of the demand curve ε. We also 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:

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

Our 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, 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. 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 $COV(\Delta p_i, \Delta\tilde{z}_i)$ 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 our IV estimate recovers the slope of Danish firms' supply curve.

In Appendix [To be written] we discuss several extensions of this simple framework. We show that our identification strategy identifies firms' supply curve with heterogeneous demand elasticities, strategic complementarity, pricing-to-market and in the case of multi-product firms. 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 demand variation 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. We now discuss the construction of these demand shifters.

Construction of demand shifters. We construct shift-share demand shifters in line with the prescriptions of section 3.1—we use annual import growth as a proxy for aggregate 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 demand for Danish products, than for example GDP growth or measures of the output gap. Second, both the country-level import data and the 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.[5]

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

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

[5] 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)

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 in time t 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 on average positive and varies over international business cycles. Our estimation will only use cross-sectional within-sectors variation in the demand shifter, 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. Even though they are not residualized, 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.

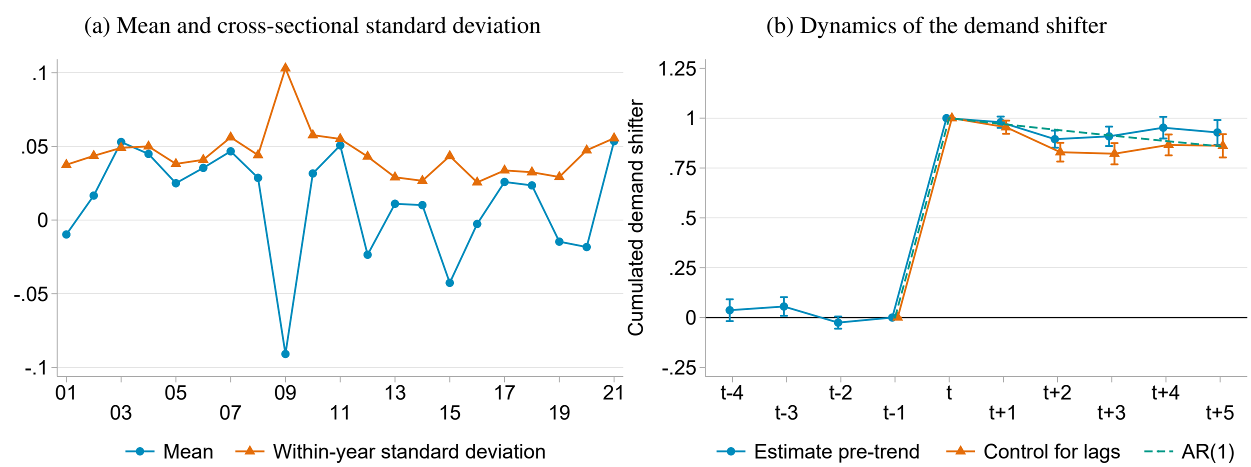

Figure 1: Properties of the demand shifter

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 sales and other firm level outcomes) or firm-product level (for price outcomes). Our main reduced-form specification has domestic prices in Denmark on the left hand side:

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

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 solve the problem of 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 that 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. To do so, we estimate an IV-local projection of firms output growth in period t on changes in prices over different horizons $t+h$:

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

In line with our analytical framework, we use $\Delta\tilde{z}_{i,t}$ as an instrument for output growth $\Delta q_{i,t}$. This estimates the response of prices at horizon $t+h$ to a demand shock that increases log output by one in period h, and can therefore be directly interpreted as the demand elasticity of prices. The parallel to the alternative dynamic panel specification we estimate for our reduced form specification is to include lagged values of output growth and add lagged shocks as additional instruments.

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 internally 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 a 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 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 variation in source countries. We include this control in our second baseline specification and none of our results are meaningfully affected by it.

4 Empirical Results

In this section, we report our baseline 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. 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. Figure 2 presents our estimates. The figure includes 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 0.6 log points contemporaneously for every unit increase in export-weighted demand, stays at about the same level the year after and then decays. The shock decays slightly faster in our specification with controls. We estimate no significant placebo pre-trend coefficients. Panel (b) shows that sales increase by about 0.7 log points in the first year. This effect persists at a similar level in the year after and then slowly decays. Similar to the response of output, the effect of sales decays slightly faster in our estimates with controls. We estimate one significant pre-treatment dip for sales—this is the main reason we estimate specifications with lagged controls for all outcomes. The difference between the results on sales and output is already indicative for an increase in prices in response to the demand shock.

Finally, for a subset of firms, we observe self-reported capacity utilization (in %). In panel (c), we show that capacity utilization of these firms 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 medium run, the effect on capacity utilization appears to be slightly more short-lived than the effect on output. Coefficient estimates, standard errors and summary statistics for all three output measures are also presented in Tables 4 (output), 6 (sales) and 8 (capacity utilization).

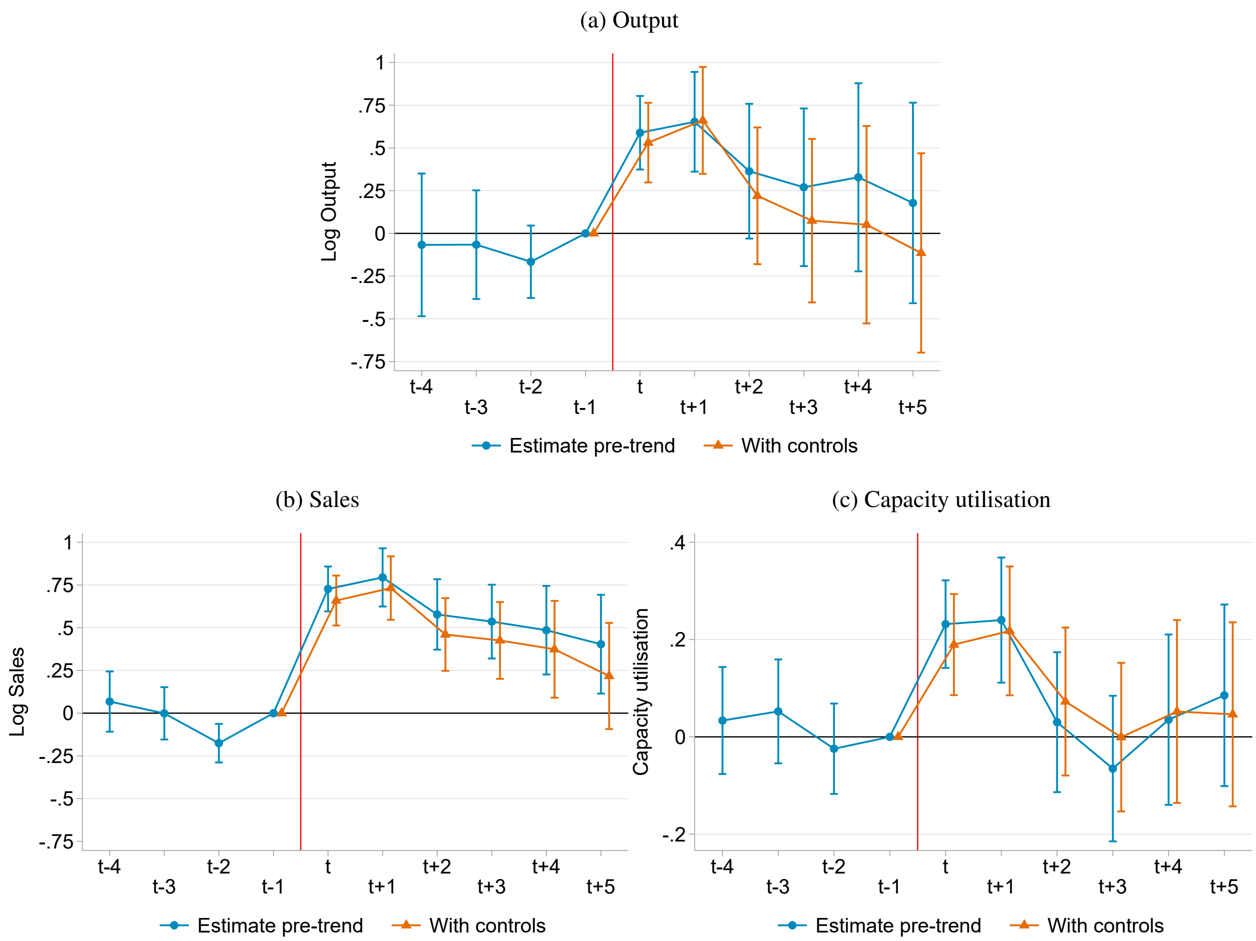

Figure 2: First stage results—effects of demand shifter on output and capacity utilisation

Reduced form results on prices. In Figure 3, panel(a), we show the reduced form response of prices to demand shocks. Coefficient estimates, standard errors and other summary statistics are also shown in Table 10. Prices increase by about 0.15 log points for every unit increase in the demand shifter on impact and continue to rise slowly for another 3 years up to a maximum increase of about 0.3 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.

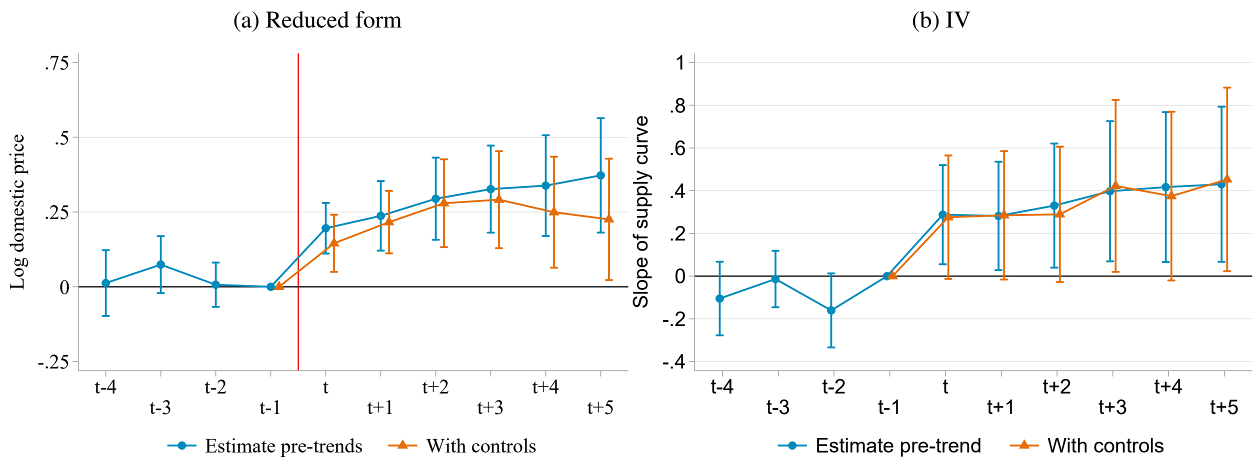

Figure 3: Effects on prices

IV results. We present our baseline IV results in Figure 3, panel(b). 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 estimates suggest that a 10% shift in demand (at constant prices) translates into a 3-4% increase in prices. Like in our reduced form estimates, we find no significant pre-trends. The elasticity increases over the first four years after a shock and then remains fairly stable at a level of 0.4. As with all our estimates our estimates, the coefficient estimates from the two baseline specifications with and without controls are very similar. Table 12 column (1) and (2) present the coefficients for our two baseline specifications, as well as other summary statistics. The F-statistic for our baseline specification without controls amounts to 31, suggesting our demand shocks are a strong instrument. The baseline specification with additional controls gives a lower F-statistic of 8.5 (explaining the higher standard errors). Since both specifications are just-identified, we are not concerned about bias arising from this weaker first-stage (Angrist and Pischke, 2008).

Effects by initial capacity utilization. For a subset of about 450 firms, we observe self-reported capacity utilisation. For these firms, we test how the effect of a demand shock depends on initial capacity utilisation 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 utilisation conditions on past shocks, we estimate this specification only using our specification that controls for lagged shocks and endogenous variables.

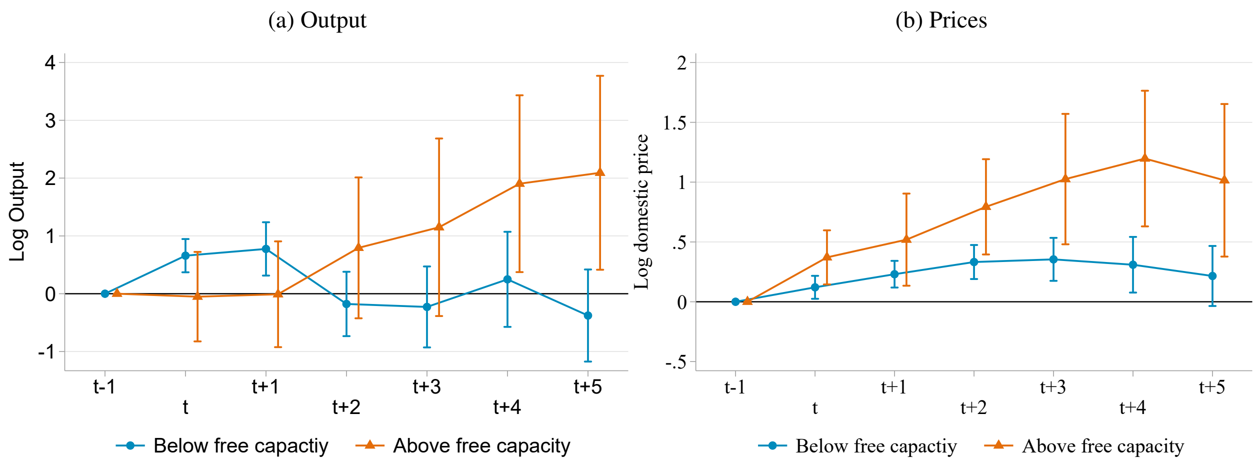

Figure 4: Effects on prices and output by initial capacity utilisation

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 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.

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 are 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 for 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 of course lower in the smaller samples.

5 Macroeconomic implications

Our empirical estimates of section 4 show that in the cross-section firms' prices respond strongly to an exogenous demand shock. This is an all-else-equal estimate—ideally, all other factors that might affect prices are held constant by including fixed effects and controls, or are orthogonal to the shock. When an economy is hit by an aggregate shock, firms respond to a change in their own demand, but also to changes in inflation expectations that arises because other firms are also experience an increase in demand. Moreover, other aggregate variables, such as real wages, might also respond and 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. In this section we augment the supply side of a standard New Keynesian model with firm level demand shocks that mirror our empirical setting. We illustrate how the response of prices to idiosyncratic demand shocks in the cross-section is driven by parameters that also determine the slope of the aggregate Phillips Curve. We then identify these parameters by fitting the response of firm level prices and quantities in the model to the responses estimated in our empirical analysis. Finally, we make predictions about the aggregate Phillips curve based on these parameter values.

5.1 A simple dynamic model

We use the fact that Denmark is a small open economy with most Danish manufacturing firms exporting a large share of their output to different destination countries in our empirical strategy. This also affects the slope of the Danish Phillips curve—in a small open economy the Phillips curve flatter than in a closed economy because the response to domestic shocks is muted by trade (Gali and Monacelli, 2005). In this section, we map our empirical results to a model of a closed economy. We interpret this model as a model of the Euro area, of which Denmark is implicitly a part due to its long-standing currency peg to the Euro[6]. This means we assume European manufacturing firms respond to demand shocks like Danish manufacturing firms.

Our model is in the aggregate identical to the production side of a textbook New Keynesian model. As in our illustrative static model in section 3.1, firms have a short-run Cobb-Douglas production function with decreasing returns to scale, i.e. $Q = N^{1-\alpha}$. The decreasing returns to scale capture the idea that some production factors are fixed in the short run, and consequently individual firms' supply curves are increasing in output. We encourage readers to interpret α as the determinant of the slope of firms' flexible price supply curve rather than literally as the output elasticity of employment.

Since we model a closed economy, we abstract from the fact that the demand shocks in our empirical setting originate from abroad. We model total firm demand, i.e. $Q = ZY(P/\overline{P})^{-\varepsilon}$ with constant demand elasticity ε. $\overline{P}$ denotes the aggregate price level, Y aggregate demand, and Z an exogenous firm-specific demand shifter. As we show in section 3.1 the response of firm prices to such a shock is the same as to our export-weighted demand shifter in a static setting. $\log Z \equiv z$ is mean zero and follows an AR(1) process with persistence ρ. Firms discount the future at a constant rate β

[6]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.

and can reset their price with probability $1-\theta$. We do not explicitly model the behavior of households or the central bank, since the slope of the output-inflation trade-off is (mostly) determined by firm behavior. Household behavior only matters for the response of real wages to an aggregate demand shock, which we instead model in a reduced form way described below.

Idiosyncratic demand shocks Firms set their price to maximize future profits whenever they have an opportunity to do so. We can distinguish between the average reset price p_t^* at $z_t = 0$, and the reset price conditional on the current demand realization z_t, which we denote $p_t^*|z_t$. We show in the appendix that the conditional reset price $p_t^*|z_t$ is equal to the aggregate unconditional reset price p_t^* plus the additional component of current and future marginal cost caused by higher than average demand:

$$p_t^*|z_t = p_t^* + (1-\theta\beta)\frac{\alpha}{1-\alpha+\alpha\varepsilon}\sum_{k=0}^{\infty}(\beta\theta)^k \mathbb{E}_t z_{t+k}. \tag{7}$$

Using the definition of the current price level as an average of the reset price and the price level in the previous period, $p_t = (1-\theta)p_t^* + \theta p_{t-1}$, and the equivalent definition for the current mean price conditional on z_t, the price of firms with demand realization z_t relative to the aggregate price level is determined by:

$$p_t|z_t - p_t = (1-\theta\beta)(1-\theta)\frac{\alpha}{1-\alpha+\alpha\varepsilon}\sum_{k=0}^{\infty}(\beta\theta)^k \mathbb{E}_t z_{t+k} + \theta(p_{t-1}|z_t - p_{t-1})$$

Since innovations to the idiosyncratic shock are i.i.d., initial prices are not related to the current shock, and $p_{t-1}|z_t = p_{t-1}$. Iterating forward using the dynamics of the shock, we can describe the path of expected relative prices following a firm level shock as follows:

$$p_{t+k}|z_t - p_{t+k} = \frac{(1-\theta\beta)(1-\theta)}{1-\theta\beta\rho}\frac{\alpha}{1-\alpha+\alpha\varepsilon}\frac{\theta^{k+1}-\rho^{k+1}}{\theta-\rho}z_t$$

This expression allows for three "regimes". If the idiosyncratic demand shock is permanent, i.e. $\rho = 1$, then the firm's price will converge slowly to a new, permanently higher price. This new price is determined by the flexible-price supply curve and demand. If the demand shock is transitory and $\rho < \theta$, the demand shock decays faster than prices are adjusted, and firms' relative price increases in the first period and then slowly returns to zero as the firms price converges to the aggregate price level. 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.

Aggregate dynamics. Apart from idiosyncratic demand shocks, our setup corresponds to the production side of a standard New Keynesian model. This standard model already features cross-sectional heterogeneity in prices and quantities due to price stickiness—firms' prices and quantities depend on when they last reset their price. Idiosyncratic demand shocks thus add no new feature to the aggregate representation of the model, and aggregate price dynamics are described by the standard New Keynesian Phillips curve that relates inflation to the deviation of real aggregate marginal cost from it's steady-state value and firms' inflation expectations:

$$\pi_t = \lambda mc_t^R + \beta \mathbb{E}_t(\pi_{t+1}) \tag{8}$$

Aggregate marginal cost is related to output through two channels[7]. The first channel works through the aggregation of firm level supply curves. In the short run, firms operate under decreasing returns to scale and their marginal cost increases when output is high. We call this the "product market slack" channel. Second, the real wage w^R (and other real factor prices), might respond to aggregate demand. We call this the "labor market slack" channel. In a fully fledged model, the relationship between real wages and output is determined by households' labor supply decision and the structure of labor markets. Here, we instead use a reduced form relationship between real wages and output, $w^R = \phi y$. We can then express the Philips curve as the relationship between inflation and output:

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t(\pi_{t+1}), \text{ where } \qquad \kappa = \underbrace{\frac{(1-\theta\beta)(1-\theta)}{\theta}\frac{\alpha}{1-\alpha+\alpha\varepsilon}}_{\text{Product market slack}\equiv\kappa^p} + \underbrace{\frac{(1-\theta\beta)(1-\theta)}{\theta}\frac{(1-\alpha)}{1-\alpha+\alpha\varepsilon}\phi}_{\text{Labor market slack}\equiv\kappa^w} \tag{9}$$

Note for a given cyclical behavior of real wages, the importance of the "product market slack" channel for inflation is increasing in the slope of firms' supply curve, while the importance of the labor market slack channel is decreasing in the slope of firms' supply curve. This is due to the fact that a steep supply curve increases the response of prices to an aggregate demand shock, but mutes the response of output and thus real wages.

Our estimates use cross-sectional variation in product market slack, but all factor prices are absorbed into time-sector fixed effects. Our estimates are thus not informative about the cyclicality of wages or their pass-through into firms' prices. We can thus use our firm-level estimates to determine the magnitude of κ^P, but need an additional estimate of ϕ, the response of real wages to aggregate demand shocks, to determine the importance of the labor market slack channel κ^W.

[7] We assume that potential output is not affected by any of the shocks we consider, and use output interchangeably with the output gap.

We complement our Phillips curve with a reduced form dynamic IS equation:

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t(\pi_{t+1}) \tag{10}$$

$$y_t = \gamma_y \mathbb{E}_t y_{t+1} + \gamma_\pi \mathbb{E}_t \pi_{t+1} + u_t. \tag{11}$$

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. These coefficients are needed to solve for the path of output and inflation. However, they are not needed to characterize the output-inflation trade-off faced by policy makers. We assume that the aggregate demand shock u_t follows an AR(1) process with persistence ρ_u.

We can use the method of undetermined coefficients to solve:

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

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

The parameter Λ depends on the coefficients of the dynamic IS-equation and is needed to solve for the separate paths of inflation and output. 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 defined as the ratio of the average response of inflation and output over an increasing horizon h[8]. We can see that the Phillips multiplier implied by the model is constant:

$$\Psi = \frac{\kappa}{1-\beta\rho_u} \tag{12}$$

5.2 Structural estimation

Estimation. We estimate the parameters of the model by fitting the impulse response of prices and quantities after an idiosyncratic demand shock to the local projections estimated in section 4. In particular, we find parameters 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. We formulate the model in quarterly frequency and aggregate prices and quantities to annual frequency exactly as we do in the data (we use end-of-year prices and yearly totals for

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

(a) Impulse response of prices

(b) Impulse response of output

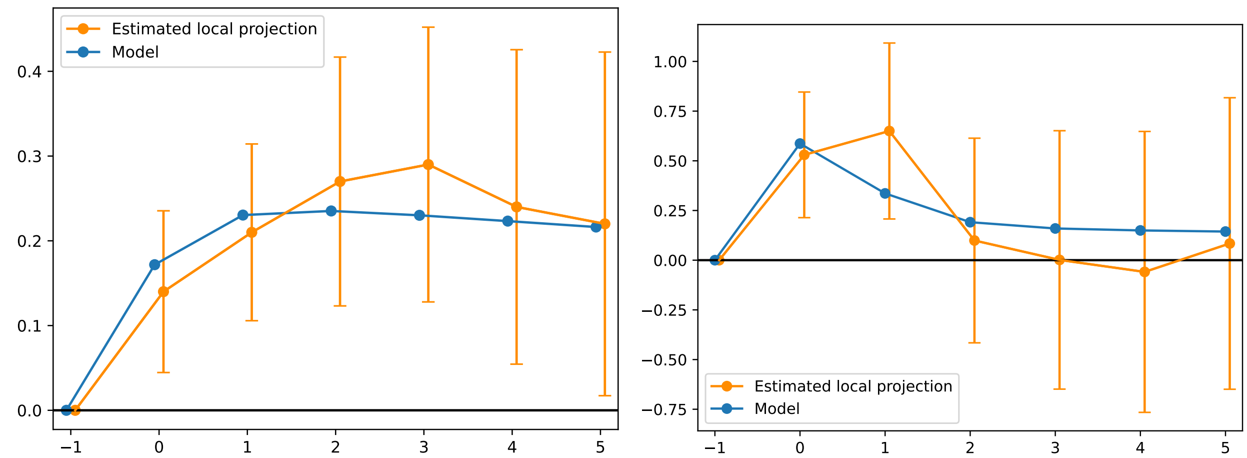

Figure 5: Impulse responses to an idiosyncratic demand shock in the data and our fitted model

quantities).

We fix the quarterly discount factor to $\beta = 0.99$, implying an annual interest rate of 0.04. We then estimate two sets of parameters. In our baseline estimation, we fix $\theta = 0.66$, the frequency of price adjustment in the Danish PPI microdata, and $\rho = 0.992$, which matches a persistence of 0.97 in our annual demand shock. We only estimate α and ε, the slope of the flexible price supply and demand curves, as well as a normalization of the initial value of the shock. In a robustness exercise, we estimate all parameters of the model except β unrestricted except for theoretical bounds.

Results at the firm level. Figure 5 shows the impulse responses of the fitted model compared to those estimated in the local projections. Our simple model with only two free parameters α and ε does remarkably well at fitting the estimated response of firm level prices and quantities. The response of prices reaches its maximum after 3 years in the model, compared to 4 years in the data, and the response of output reaches it's maximum in the first year, compared to the second year in the data, but overall these differences are small compared to the confidence intervals of our estimates.

Table 2 shows the estimated structural parameters. Our baseline estimates in column (1) are $\alpha = 0.85$ and $\varepsilon = 1.3$. Both values deviate from common macroeconomic calibrations or priors is estimated models. α and other output elasticities are usually set without considering their importance for the shape of firms supply curve, but rather to match the aggregate factor income distribution. Usually, $1 - \alpha$ is set to values around 2/3 to match the value of the labor share in output—our estimate of α is larger by a factor of 3 and implies substantially steeper supply curves. The elasticity of

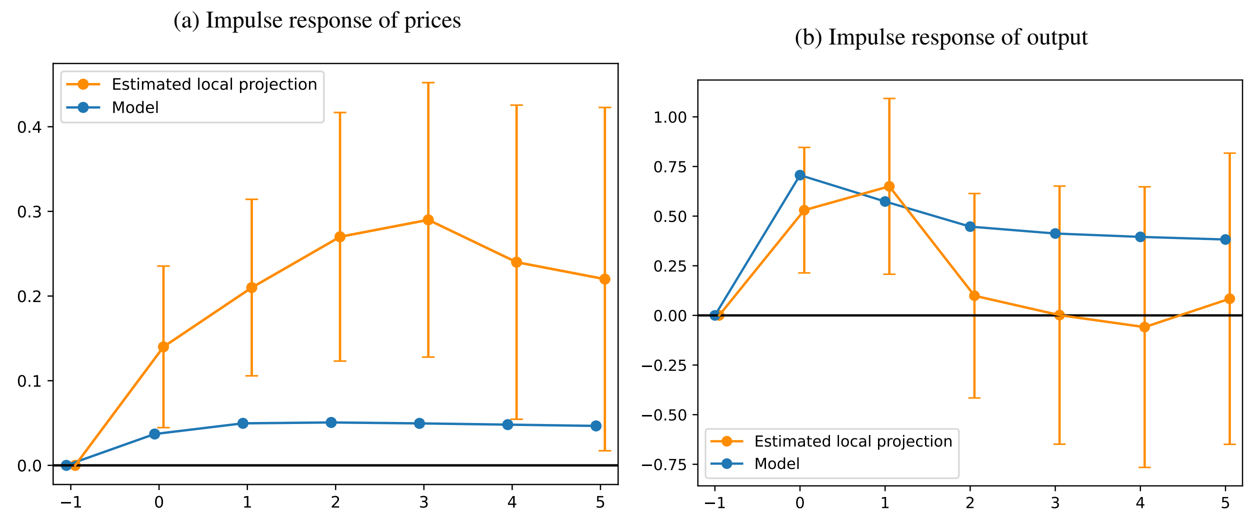

Figure 6: Counterfactual responses to an idiosyncratic demand shock with $\alpha = 0.33$ and $\varepsilon = 4$

demand ε is commonly set to values around 4 to match steady-state markups—our estimate suggests a flatter demand curve than is commonly assumed. In Figure 6 we show the response to firm level demand shocks using $\alpha = 0.33$ and $\varepsilon = 4$. The response of prices to a demand shock is much smaller, and the response of output much larger. At the firm level, the output-inflation trade-off is thus much weaker using a standard calibration than what we observe in the data.

Column (2) of Table 2 shows the unrestricted estimates of α, ε, ρ and θ. Our estimation chooses values of α and ε that are even higher and lower respectively than the estimate conditional on fixed values of ρ and θ. The estimated value of ρ is very close to the one we calibrate based on knowledge of the shock process (even though this is not targeted). The frequency of price adjustment θ is estimated to be 0.82, which is higher than in the data. For the slope of the supply and demand curves, the estimates are at the theoretical bounds we impose ($\alpha = 0.999$ and $\varepsilon = 1$). While the unrestricted model does naturally better in fitting the impulse response (see Figure 7 in the Appendix), we prefer our restricted estimate that uses information on the observed frequency of price adjustment.

Slope of the Phillips curve. Finally, Table 3 shows the implied slope of the Phillips curve κ and the implied Phillips multiplier Ψ. We separately calculate the components of the slope κ^P ("product market slack", see equation (8)) using our estimates of structural parameters, and κ^W ("labor market slack") using different assumptions on the response of real wages to an aggregate demand shock. We calculate the Phillips multiplier Ψ, using a persistence of the aggregate shock $\rho_u = 0.69$ taken from Jarocinski and Karadi (2020) to match the persistence of ECB monetary policy surprises.

Table 2: Structual parameter estimates

	(1) Restricted estimates	(2) Unrestricted estimates	(3) Gali (2008) calibration
α	0.858	0.999	0.330
ε	1.334	1.000	5.000
ρ	0.992	0.990	0.992
θ	0.660	0.815	0.660
MSE	4.115	1.925	44.333

Notes: The table plots the structural parameter estimates that matching the impulse response to idiosyncratic demand shocks in the model and the data. (1) restricts $\rho = 0.992$ and $\theta = 0.66$ to their values directly observed in the data. (2) leaves all parameters unrestricted within their theoretical bounds. (3) shows the calibration of Gali (2008) for comparison.

Our estimates suggest that $\kappa^P = 0.119$. That means that downward sloping supply curves on their own lead to an (initial) increase in inflation by 0.119 in response to an aggregate demand shock that increases output by 1%, holding inflation expectations fixed. If we assume that real wages do not respond to an aggregate demand shock (for example because nominal wages are indexed to inflation), this is equal to the slope of the Phillips curve, and taking into account the effects that arise through changes in inflation expectations, the total medium-run increase in the price level corresponds to 0.376.

We consider two alternative scenarios for the behavior of real wages after an aggregate demand shock. First, 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) and use the average of the impulse responses. Danish real wages fall after a monetary policy surprise and the suggested output elasticity is -0.15. Combined with out micro level estimates, this suggests $\kappa^W = -0.003$ and the slope of the overall Phillips curve in this scenario is equal to 0.12. The slope of the Phillips curve is dominated by the "product market slack" channel. Second, we take the correlation 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 $1/3$. This results in $w^R = 2.5y$, i.e. real wages that move very strongly with output and $\kappa^W = 0.049$. The overall slope κ in this case amounts to 0.168 and the Phillips multiplier Ψ amounts to 0.531. While our asumptions on the behavior of real wages naturally matter for the slope of the Phillips curve, all scenarios we consider result in what would qualify as "steep" Phillips curves relative to other recent estimates.

Discussion of alternative estimates We compare the slope of the Phillips curve implied by our estimated to several values in the recent literature. 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 slope of $\kappa = 0.128$ and a Phillips multiplier of $\Psi = 0.402$. That is very similar to the Phillips curve implied by our estimates, but the source of the slope is different: our results imply a steep Phillips curve results from steep firm level supply curves, but in the Gali calibration this results from a strong co-movement of real wages and output. This is at odds with the behavior of real wages at least in Danish data. We also compare our results to recent macro estimates of the Phillips curve slope in Barnichon and Mesters (2020) and the Phillips multiplier in Barnichon and Mesters (2021).[9] In both papers the trade-off between output and inflation is relatively steep, albeit imprecisely estimated. The magnitude of the Phillips curve slope implied by our estimates is comparable to these macro estimates.

Our paper estimates steeper Phillips curves than two recent papers using regional and firm-level panel data. Gagliardone et al. (2023) provide estimates for two Phillips curve slopes—their main estimates are for the marginal cost Phillips curve. These estimates suggest a relatively steep marginal cost Phillips curve with a slope of 0.053, but are difficult to compare with our estimates without an estimated elasticity of marginal cost to output (which is a part of the supply curves we estimate). They provide direct estimates of an output Phillips curve from a firm-level regression of prices on output with monetary policy surprises interacted with firm dummies as instruments. This is similar in spirit with our estimation of firm supply curves, but produces a much lower slope estimate of 0.03 and implies an output Phillips curve slope coefficient of about 0.017. We think a possible explanation of the discrepancy between their estimate and ours lies in the weakness of the instruments they use. It is well understood that monetary policy surprises are a weak instrument in aggregate regressions. Gagliardone et al. (2023) interact this instrument with firm dummies, which makes the number estimated first-stage parameters equal to the number of firms in the data. This turns an estimation with one weak instrument into a situation with many (likely also) weak instruments and leads to a plethora of econometric problems. Most imporantly, TSLS with many and especially with many weak instruments is inconsistent and estimates are severely biased toward OLS (see Hansen et al., 2008, Mikusheva and Sun, 2024, for reviews), which in itself would be biased downward (if one wants to estimate a supply curve) because of simultaneity of demand and supply. We view the main contribution of our work relative to theirs in our focus on the output Phillips curve, and use of a stronger cross-sectional demand shock for reliable estimation of firms' supply curves and subsequently the aggregate Phillips curve.

Hazell et al. (2022) estimate Phillips curves using state-level data on prices of non-tradeables. They use variation in exposure to aggregate shocks to tradeable industries as a source of variation in local aggregate demand. They

[9] We refer to their estimates of US Phillips curves and multipliers for the post 1990 period using HFI shocks.

Table 3: Macro parameters

	κ^P	κ^W	κ	Ψ^P	Ψ
Based on our estimates					
with $\phi = 0.00$	0.119	0.0	0.119	0.376	0.376
with $\phi = 2.50$	0.119	0.049	0.168	0.376	0.531
with $\phi = -0.15$	0.119	-0.003	0.116	0.376	0.366
Comparison to other estimates					
Gali (2008) calibration	0.021	0.106	0.128	0.067	0.402
Barnichon and Mesters (2020)			0.12		0.061
Barnichon and Mesters (2021)			0.181		0.181
Hazell et. al. (2022)			0.008		0.013
Gagliardone et. al. (2024) - Output Phillips curve			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^W for different sets of parameters. Row 1 uses our baseline estimates for α, ε and $\theta = 0.66$ and assumes that real wages do not vary with aggregate output, i.e. $\phi = 0$. Row 2 uses the elasticity of real wages to the output corresponding to the Gali (2008) calibration. Row 3 uses the elasticity of real wages to the output estimated from aggregate Danish data. Row 4 gives results for the Gali (2008) calibration. Row 5 shows the estimates of Hazell et al. (2022) converted from their unemployment gap Phillips curves to an output-based Phillips curve using the covariance of the quarterly unemployment and output gaps in the US. Row 5 gives the values estimated for the Euro area in the semi-structural ECB-Base model.

estimate[10] a slope of 0.005 and a Phillips multiplier of 0.035. This is much flatter than implied by our estimates. There are several possible explanations for this discrepancy. First there is the possibility that their estimates are biased downward. This would happen if there is correlation between local supply shocks and their instrument. For example, it could be that states that are highly exposed to booming industries receive inflows of workers and capital and therefore both aggregate supply and demand are affected. An alternative explanation is that prices of non-tradeable goods and services in their price index simply react less to aggregate demand than those of manufacturing firms. This is supported by evidence in Aruoba and Drechsel (2024), who find that consumer prices of services generally respond very little to monetary policy surprises, while consumer prices of non-durables but especially of durable goods respond faster and more strongly. In this case the true aggregate Phillips curve would be an average between their estimate and ours.

[10] 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.

6 Conclusion

Our paper makes several contributions to our understanding of 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 demand shocks. Our findings reveal that firm-level supply curves are substantially steeper than commonly assumed in macroeconomic models—a demand shock that increases output by 1% leads to a 0.3% 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. Our fitted model suggests that the slope of the Phillips curve is primarily determined by the aggregation of firm-level supply curves rather than labor market dynamics. Our preferred estimate for the slope of the output Phillips curve is 0.119, implying that a moderately persistent monetary policy shock that increases output by 1% results in a 0.37% increase in the price level in the medium run.

Our findings contribute to the ongoing debate about the flatness of the Phillips curve. While some recent studies using regional variation find flatter Phillips curves, our results based on firm-level evidence suggest a more substantial trad,e-off between inflation and output. The difference may arise from our focus on manufacturing firms' producer prices rather than consumer prices of services, or from the strength of our identification strategy that allows us to control for all aggregate supply variation.

Our work highlights the possibilities of using microdata and credible identification strategies to understand aggregate price dynamics. The steep firm-level supply curves we document suggest that individual firms face significant constraints in adjusting their output in response to demand, leading to price adjustments that aggregate into meaningful inflation dynamics. This has important implications for monetary policy, suggesting that central banks face a steeper short-run trade-off between stabilizing output and inflation than indicated by some recent studies.

References

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.

Aruoba, S. Borağan and Thomas Drechsel, "The Long and Variable Lags of Monetary Policy: Evidence from Disaggregated Price Indices," *Journal of Monetary Economics*, July 2024, p. 103635.

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.

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.

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," Technical Report w31382, National Bureau of Economic Research, Cambridge, MA June 2023.

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.

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 Supplementary 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	0.53***	0.59***	0.49***	0.54***	0.46**	0.75***
	(0.12)	(0.11)	(0.13)	(0.12)	(0.20)	(0.10)
t+1	0.66***	0.65***	0.50***	0.58***	0.91***	0.82***
	(0.16)	(0.15)	(0.19)	(0.16)	(0.30)	(0.14)
t+3	0.075	0.27	-0.22	-0.0092	0.31	0.14
	(0.24)	(0.24)	(0.32)	(0.24)	(0.46)	(0.20)
t+5	-0.11	0.18	-0.66	-0.034	-0.0056	0.16
	(0.30)	(0.30)	(0.44)	(0.24)	(0.57)	(0.26)
firms	717	762	706	677	671	720
N	7,642	9,449	7,336	7,602	6,796	7,670
F	9.098	28.791	5.354	13.539	4.773	15.260

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	0.53***	0.56***	0.59***	0.57***	0.55***	0.50***
	(0.12)	(0.11)	(0.099)	(0.13)	(0.13)	(0.13)
t+1	0.66***	0.54***	0.62***	0.59***	0.59***	0.60***
	(0.16)	(0.15)	(0.14)	(0.17)	(0.18)	(0.19)
t+3	0.075	0.34	0.45**	0.15	0.20	0.20
	(0.24)	(0.22)	(0.22)	(0.26)	(0.27)	(0.27)
t+5	-0.11	0.13	0.26	0.020	0.085	0.085
	(0.30)	(0.27)	(0.26)	(0.31)	(0.32)	(0.32)
firms	717	791	1,817	629	556	545
N	7,642	8,662	14,740	6,945	5,742	5,105
F	9.098	9.804	9.577	6.749	6.301	5.746

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	0.66***	0.73***	0.60***	0.69***	0.41***	0.87***
	(0.074)	(0.067)	(0.078)	(0.074)	(0.11)	(0.067)
t+1	0.73***	0.79***	0.63***	0.70***	0.51***	0.93***
	(0.095)	(0.087)	(0.10)	(0.094)	(0.16)	(0.089)
t+3	0.43***	0.54***	0.35***	0.30***	0.28	0.58***
	(0.11)	(0.11)	(0.13)	(0.098)	(0.18)	(0.10)
t+5	0.22	0.40***	0.11	0.17	0.040	0.49***
	(0.16)	(0.15)	(0.18)	(0.12)	(0.27)	(0.14)
firms	820	854	817	782	775	820
N	9,516	11,740	9,336	9,478	8,616	9,527
F	21.203	117.802	20.019	21.943	5.633	38.704

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	0.66***	0.75***	0.61***	0.63***	0.64***	0.60***
	(0.074)	(0.069)	(0.055)	(0.074)	(0.074)	(0.078)
t+1	0.73***	0.75***	0.68***	0.71***	0.72***	0.71***
	(0.095)	(0.088)	(0.070)	(0.096)	(0.099)	(0.10)
t+3	0.43***	0.52***	0.54***	0.42***	0.46***	0.46***
	(0.11)	(0.10)	(0.093)	(0.12)	(0.12)	(0.12)
t+5	0.22	0.31**	0.49***	0.30*	0.36**	0.36**
	(0.16)	(0.14)	(0.12)	(0.17)	(0.17)	(0.17)
firms	820	911	2,269	711	624	615
N	9,516	10,928	19,865	8,643	7,066	6,317
F	21.203	26.839	33.004	21.038	21.320	18.271

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 utilisation — 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***	0.19**	0.25***	0.14	0.23***
	(0.063)	(0.055)	(0.075)	(0.067)	(0.18)	(0.054)
t+1	0.22***	0.24***	0.14	0.25***	0.30	0.22***
	(0.080)	(0.078)	(0.10)	(0.082)	(0.25)	(0.060)
t+3	-0.00053	-0.065	-0.0033	0.088	-0.058	-0.0066
	(0.093)	(0.091)	(0.12)	(0.089)	(0.29)	(0.088)
t+5	0.046	0.085	0.11	0.19	-0.22	0.15
	(0.11)	(0.11)	(0.15)	(0.12)	(0.45)	(0.092)
firms	408	448	386	362	341	411
N	3,190	3,896	2,962	3,141	2,393	3,250
F	27.334	17.924	8.075	30.637	7.825	31.745

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 utilisation — 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***	0.20***	0.18**	0.17**	0.14*
	(0.063)	(0.057)	(0.057)	(0.069)	(0.073)	(0.077)
t+1	0.22***	0.19***	0.13*	0.16*	0.12	0.15
	(0.080)	(0.069)	(0.077)	(0.090)	(0.096)	(0.10)
t+3	-0.00053	0.046	0.091	-0.025	-0.033	-0.033
	(0.093)	(0.086)	(0.087)	(0.10)	(0.10)	(0.10)
t+5	0.046	0.031	0.041	-0.024	-0.078	-0.078
	(0.11)	(0.11)	(0.12)	(0.13)	(0.13)	(0.13)
firms	408	448	650	371	328	302
N	3,190	3,579	4,128	2,931	2,385	2,073
F	27.334	24.218	32.744	25.218	22.431	20.009

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	0.15***	0.20***	0.15***	0.13**	0.13**	0.13**
	(0.049)	(0.043)	(0.055)	(0.053)	(0.051)	(0.063)
t+1	0.22***	0.24***	0.25***	0.17***	0.24***	0.095
	(0.053)	(0.059)	(0.061)	(0.055)	(0.084)	(0.059)
t+3	0.29***	0.33***	0.36***	0.14*	0.30***	0.20**
	(0.083)	(0.074)	(0.090)	(0.074)	(0.11)	(0.084)
t+5	0.23**	0.37***	0.35**	0.037	0.088	-0.019
	(0.10)	(0.098)	(0.15)	(0.082)	(0.15)	(0.13)
firms	589	715	583	581	581	589
N	16,998	25,277	16,907	16,990	16,861	17,006
F	6.822	20.516	1.922	7.483	4.668	5.918

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	0.15***	0.15***	0.15***	0.13**	0.15**
	(0.049)	(0.045)	(0.055)	(0.056)	(0.066)
t+1	0.22***	0.23***	0.22***	0.20***	0.21***
	(0.053)	(0.049)	(0.060)	(0.059)	(0.059)
t+3	0.29***	0.32***	0.34***	0.27***	0.27***
	(0.083)	(0.079)	(0.091)	(0.088)	(0.088)
t+5	0.23**	0.26***	0.34***	0.31***	0.31***
	(0.10)	(0.092)	(0.11)	(0.11)	(0.11)
firms	589	649	536	462	443
N	16,998	19,408	15,988	12,884	11,112
F	6.822	4.275	5.815	5.487	7.654

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*	0.29**	0.21*	0.23	0.25**
	(0.15)	(0.12)	(0.12)	(0.19)	(0.12)
t+1	0.28*	0.28**	0.19	0.27	0.17*
	(0.15)	(0.13)	(0.12)	(0.22)	(0.098)
t+3	0.42**	0.40**	0.25*	0.49	0.32**
	(0.21)	(0.17)	(0.14)	(0.35)	(0.14)
t+5	0.45**	0.43**	0.18	0.42	0.24**
	(0.22)	(0.19)	(0.14)	(0.32)	(0.12)
firms	751	808	736	746	751
N	17,305	20,597	17,293	17,224	17,307
rkf	8.537	31.463	6.261	2.526	10.893

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*	0.22	0.19*	0.20*	0.19*	0.23
	(0.15)	(0.13)	(0.099)	(0.11)	(0.11)	(0.14)
t+1	0.28*	0.22*	0.19*	0.18*	0.18	0.21
	(0.15)	(0.14)	(0.11)	(0.11)	(0.11)	(0.13)
t+3	0.42**	0.34*	0.29**	0.29*	0.26*	0.28*
	(0.21)	(0.19)	(0.14)	(0.15)	(0.15)	(0.17)
t+5	0.45**	0.35*	0.32*	0.33**	0.33**	0.36*
	(0.22)	(0.18)	(0.17)	(0.17)	(0.17)	(0.20)
firms	751	843	740	657	578	558
N	17,305	19,498	16,667	16,530	14,149	14,149
rkf	8.537	7.092	8.193	6.541	6.243	5.515

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.

(a) Impulse response of prices

(b) Impulse response of output

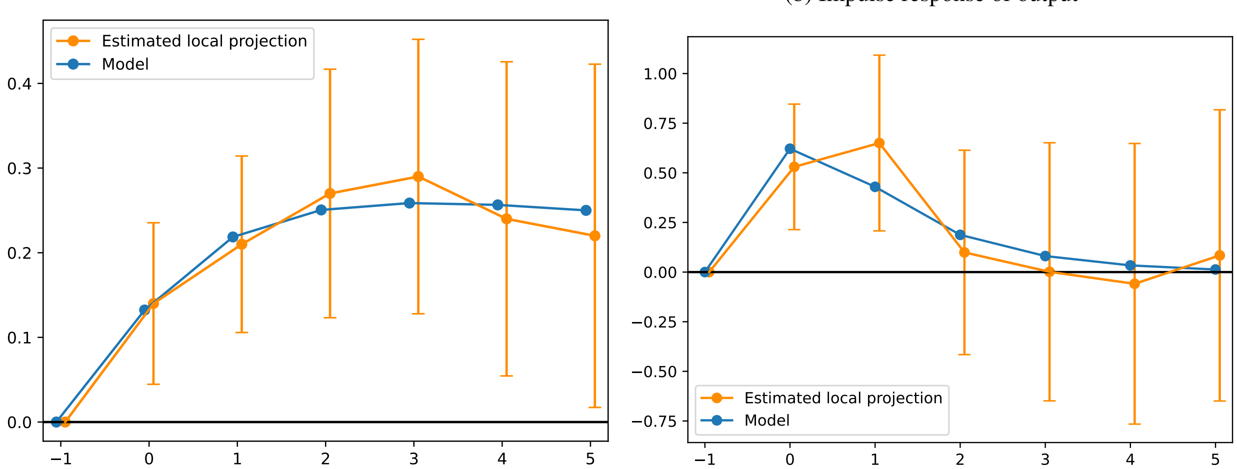

Figure 7: Impulse responses to an idiosyncratic demand shock in the data and our unrestricted fitted model

B Derivations of dynamic price-setting model

Our setup is equal to the production side of a textbook New Keynesian model. Firms maximize their future discounted profit when they have an opportunity to reset their price to a new price P_t^*. This yields the first order condition:

$$\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t\left(Q_{t+k}(P_t^*)(P_t^*-\mu Q_{t+k}(P_t^*)^{\frac{\alpha}{1-\alpha}}W_{t+k})\right)=0, \tag{13}$$

where μ is the flexible price markup and $Q_{t+k}(P_t^*)^{\frac{\alpha}{1-\alpha}}W_{t+k}$ is nominal marginal cost in period $t+k$ at price P_t^*. We log-linearize this equation around a point where all shocks are equal to their mean. We get:

$$p_t^*=(1-\theta\beta)\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t mc_{t+k|t} \tag{14}$$

Aggregate dynamics We first recap the basic aggregate dynamics in this setting. Using the definition of demand and the cost function, we can express future marginal cost for the average firm who last reset prices at time t as $mc_{t+k|t}=mc_{t+k}-\frac{\alpha\varepsilon}{1-\alpha}(p_t^*-\overline{p}_{t+k})$. We can plug this expression into (14):

$$\begin{aligned}
p_t^* &= \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon}\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t(mc_{t+k}+\frac{\alpha\varepsilon}{1-\alpha}\overline{p}_{t+k})\\
&= \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon}\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t(mc^R_{t+k}+\frac{1-\alpha+\alpha\varepsilon}{1-\alpha}\overline{p}_{t+k})\\
&= \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon}\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t mc^R_{t+k}+(1-\theta\beta)\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t\overline{p}_{t+k}
\end{aligned}$$

We subtract the price level in the base period:

$$\begin{aligned}
p_t^*-\overline{p}_{t-1} &= \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon}\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t mc^R_{t+k}+\sum_{k=0}^{\infty}(\beta\theta)^k\mathbb{E}_t\pi_{t+k}\\
p_t^*-\overline{p}_{t-1} &= \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon}mc^R_t+\pi_t+\beta\theta\mathbb{E}_t(p^*_{t+1}-\overline{p}_t)
\end{aligned}$$

Finally, we combine this with the definition of inflation as $\pi_t = (1-\theta)(p_t^* - \overline{p}_{t-1})$ to get the New Keynesian marginal cost Phillips curve:

$$\pi_t = \frac{(1-\alpha)(1-\theta\beta)(1-\theta)}{1-\alpha+\alpha\varepsilon} mc_t^R + (1-\theta)\pi_t + (1-\theta)\theta\beta\mathbb{E}_t(p_{t+1}^* - \overline{p}_t)$$
$$\pi_t = \frac{(1-\alpha)(1-\theta\beta)(1-\theta)}{\theta(1-\alpha+\alpha\varepsilon)} mc_t^R + \beta\mathbb{E}_t(\pi_{t+1})$$

Firm level dynamics Our only departure from the production side of a textbook New Keynesian model is the addition of firm level demand shocks. We derive the impulse response of the average price of a firm hit by a given demand shock relative to the aggregate price level. Heterogeneity in the price paths of different firms hit by the same shock arises due to the stochastic nature of price adjustment. We average over firms that update their prices at different times.

We can express marginal cost of a firm that last reset its prices at time t, conditional on demand realization z_t as $mc_{t+k|t,z_z} = mc_{t+k} + \frac{\alpha}{1-\alpha}(z_{t+k} - \varepsilon(p_t^* - \overline{p}_{t+k}))$. Consequently, the reset price conditional on z_t is equal to:

$$p_{t|z_t}^* = \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon} \sum_{k=0}^{\infty} (\beta\theta)^k \mathbb{E}_t(mc_{t+k} + \frac{\alpha}{1-\alpha} z_{t+k} + \frac{\alpha\varepsilon}{1-\alpha}\overline{p}_{t+k}) \tag{15}$$

$$\tag{16}$$

Which we next express relative to the unconditional average reset price:

$$p_{t|z_t}^* - p_t^* = \frac{(1-\alpha)(1-\theta\beta)}{1-\alpha+\alpha\varepsilon} \sum_{k=0}^{\infty} (\beta\theta)^k \frac{\alpha}{1-\alpha} \mathbb{E}_t z_{t+k} \tag{17}$$

$$p_{t|z_t}^* - p_t^* = z_t \frac{\alpha(1-\theta\beta)}{(1-\alpha+\alpha\varepsilon)(1-\beta\theta\rho)} \tag{18}$$

We are interested in the average k period ahead relative price level after an initial shock to z_t. We can use the definition of the price level as $p_t = (1-\theta)p_t^* + \theta p_{t-1}$ and the equivalent definition for a the price level conditional on an initial shock and iterate forward:

$$\overline{p}_t|z_t - \overline{p}_t = (1-\theta)(p_t^*|z_t - p_t^*) + \theta(\overline{p}_{t-1}|z_t - \overline{p}_{t-1}) \tag{19}$$

$$\overline{p}_{t+1}|z_t - \overline{p}_{t+1} = (1-\theta)(p_{t+1}^*|z_t - p_{t+1}^*) + \theta(\overline{p}_t|z_t - \overline{p}_t) \tag{20}$$

$$\overline{p}_{t+2}|z_t - \overline{p}_{t+2} = (1-\theta)(p_{t+2}^*|z_t - p_{t+2}^*) + \theta(\overline{p}_{t+1}|z_t - \overline{p}_{t+1}) \tag{21}$$

$$\ldots = \ldots \tag{22}$$

We can use that the initial shock is i.i.d. and hence uncorrelated with the initial price level, and $(\overline{p}_{t-1}|z_t - \overline{p}_{t-1}) = 0$. We get that

$$\overline{p}_{t+k}|z_t - \overline{p}_{t+k} = (1-\theta)\sum_{j=0}^{k}\theta^{k-j}(p_{t+j}^*|z_t - p_{t+j}^*) \tag{23}$$

$$= \frac{\alpha(1-\theta\beta)(1-\theta)}{(1-\alpha+\alpha\varepsilon)(1-\beta\theta\rho)}\sum_{j=0}^{k}\theta^{k-j}\rho^j z_t \tag{24}$$

Finally, calculating the sum of the finite geometric series, we get that:

$$\overline{p}_{t+k}|z_t - \overline{p}_{t+k} = z_t\frac{\alpha(1-\theta\beta)(1-\theta)}{(1-\alpha+\alpha\varepsilon)(1-\beta\theta\rho)}\frac{\theta^{k+1}-\rho^{k+1}}{\theta-\rho} \tag{25}$$

C 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:

$$\Delta^h y_t = \beta^h u_t + \sum_{k=1}^{4} u_{t-k} + \sum_{k=1}^{4} \Delta y_{t-k} \tag{26}$$

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.

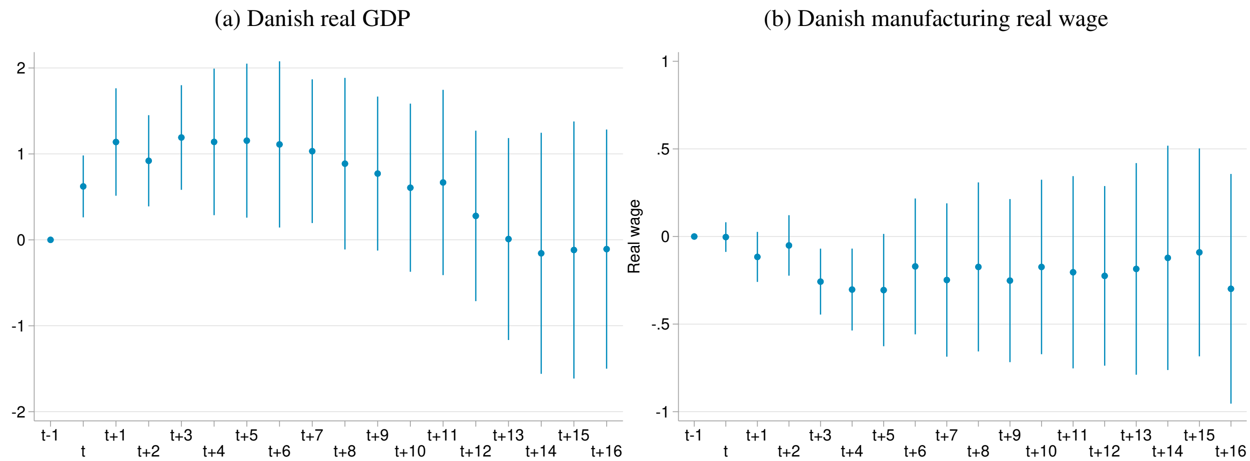

Figure 8: Effects of an ECB monetary policy surprise