The Capacity Pressure Channel of the Phillips Curve

Christian Philip Hoeck
(University of Southern Denmark)

Tobias Renkin
(Danmarks Nationalbank)

March 24, 2026

Abstract

We estimate the response of domestic prices and total output of Danish manufacturing firms to persistent firm-level demand shocks that result from heterogeneity in firms' exposure to different export destinations. Our results suggest that supply curves at the firm level are steep—a demand shock that increases output by 1% raises prices by about 0.5% after three years. We use a New Keynesian model with firm-level demand shocks to show that the estimated response maps to the "capacity pressure" channel of the Phillips curve slope, which results from upward-sloping firm-level supply curves due to fixed factors of production. In the estimated model that fits firm-level responses, the capacity pressure channel contributes about 0.039 to the slope of the Phillips curve. This is larger than recent cross-sectional estimates of the total Phillips curve slope. Under an upper-bound assumption about the wage response to aggregate shocks, our estimated model suggests the capacity pressure channel accounts for two thirds of the overall Phillips curve slope.

Hoeck: University of Southern Denmark, Campusvej 55, 5230 Odense, Denmark, cpho@sam.sdu.dk;
Renkin: Danmarks Nationalbank, Langelinie Allé 47, 2100 København Ø, Denmark, tobias.renkin@gmail.com;

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

An earlier draft of this paper was circulated under the title "Demand Shocks and Prices: Micro Evidence and Macro Implications". We are grateful to Simone Lenzu, Jon Steinsson, Emi Nakamura, Jonathon Hazell, Andrea Cerrato, Federico Ravenna, and Giulia Gitti, as well as audiences at the NBER Summer Institute, the Theories and Methods in Macro (T2M) conference, the New Developments in Business Cycle Analysis conference at Norges Bank, the Understanding Inflation workshop at Collegio Carlo Alberto, the Banque de France and the Swiss National Bank for comments and suggestions. All remaining errors are our own.

1 Introduction

The Phillips curve is an aggregate supply curve that determines the relationship between aggregate prices and the output gap. Its slope determines 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 and 2020 has sparked a debate over a possible flattening of the Phillips curve. This debate has highlighted the considerable uncertainty about the slope of the Phillips curve due to weak identification in the traditional macroeconometric approach to estimation (Mavroeidis et al., 2014) and unaddressed simultaneity of aggregate supply and demand (McLeay and Tenreyro, 2019).

In this paper, we propose a new approach to estimate the "capacity pressure" channel of the Phillips curve slope as a lower bound to the overall slope. The "capacity pressure" channel follows from fixed factors of production that give a positive slope to supply curves at the firm level. Consequently, the strength of the aggregate capacity pressure channel is closely related to the response of firms' prices and output to firm-level demand shocks. We estimate the response of prices, output, and other firm-level variables to such shocks using micro data on Danish manufacturing firms, and use a New Keynesian model to map firm-level responses to the aggregate Phillips curve. Our analysis yields two key results. First, in our estimated model that is consistent with the firm-level responses to demand shocks, the capacity pressure channel adds about 0.039 to the slope of the Phillips curve. The estimate of the capacity pressure channel on its own is a lower bound for the slope of the Phillips curve, but already larger than recent cross-sectional estimates of the total Phillips curve slope. Our second main result is that this lower bound is useful—even when we calibrate an unrealistically steep wage Phillips curve, the capacity pressure channel in our estimated model amounts to two thirds of the overall price Phillips curve slope.

The main advantage of our focus on the capacity pressure channel over existing cross-sectional approaches is that it can be identified from firm-level data under weaker assumptions, while still providing a useful lower bound for the overall Phillips curve slope. Several recent contributions estimate the overall slope of the Phillips curve from cross-sectional data using U.S. state-level or sector-level shift-share demand shocks (Hazell et al., 2022, Gagliardone et al., 2025). To identify the overall slope of the Phillips curve, both contributions rely on the validity of their respective demand instruments *and* strong additional assumptions on factor mobility that are unlikely to hold in practice and would bias slope estimates toward zero if violated.[1] Our firm-level approach relaxes identification assumptions on the demand shocks we use, since firm-level estimation allows more granular fixed effects to absorb supply shocks. Moreover, it does not require additional assumptions beyond those inherent in a New Keynesian model and shared with other approaches in the literature. In addition, the strength of the capacity pressure channel is a policy-relevant parameter in itself. It measures the slope of the Phillips curve under a "neutral" wage policy that aligns real wages with productivity growth. In most European economies, aggregate wage dynamics are heavily influenced by collective bargaining partners, whose decisions might partially reflect policy choices. Separately identifying the capacity pressure channel thus enables a discussion of the slope of the Phillips curve and the inflation cost of aggregate demand management under different wage-setting policies.

Our empirical approach relies on the fact that the magnitude of the capacity pressure channel is determined by firm-level supply and demand parameters that can be identified from the response of firms' prices, output, and

[1] The approach of Hazell et al. (2022) requires perfect labor mobility between tradeable and non-tradeable sectors, and their slope estimates would be biased downward if there is less than perfect mobility. Gagliardone et al. (2025) require zero labor mobility between different manufacturing sectors and their results would likely be biased downward if there is some mobility.

material inputs to firm-level demand shocks. To estimate these responses, we rely on shift-share demand shocks to Danish manufacturing firms. We combine cross-sectional heterogeneity in the importance of different export destinations for Danish firms with fluctuations in aggregate imports in those destinations over time.[2] We estimate the dynamic response of prices, output, and material inputs to this demand shock via local projections, using time-sector fixed effects to control for aggregate supply shocks and inflation expectations. Our approach thus avoids the most challenging identification problems inherent in estimation of the Phillips curve slope from aggregate data. The central identification assumption to recover causal firm-level responses is that the variation in exposure to different export destinations is orthogonal to firm-level supply shocks (i.e., we rely on a shares-based shift-share identification strategy as in Goldsmith-Pinkham et al. (2020)).

The resulting reduced-form estimates are not informative about the slope of the Phillips curve on their own, since general equilibrium effects are absorbed in time-sector fixed effects (the "missing intercept"). To map our firm-level results to the aggregate, we augment a New Keynesian model with firm-level demand shocks that resemble the shocks in our data. We then estimate the model by matching firm-level impulse responses to those shocks in the model to the estimated local projections. This identifies the supply and demand parameters necessary to quantify the capacity pressure channel. Our main results are entirely driven by this matching procedure. In a model that fits the firm-level response, the slope of firms' supply curve is steeper than in common calibrations of the New Keynesian model, where firm-level supply curves are often completely flat or implicitly calibrated based on the labor share in aggregate income. In the estimated model, capacity pressure contributes 0.039 to the slope of the Phillips curve. This magnitude is robust in model extensions in which we add stronger strategic complementarity in price-setting (Kimball demand) and monopsony in labor markets.

Our work is related to the 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. Most of the classical literature on Phillips curve estimation focuses on this issue and uses macroeconomic time series with identification based on internal instruments and rational expectations assumptions. Mavroeidis et al. (2014) show that this approach is subject to severe weak identification problems. McLeay and Tenreyro (2019) discuss a second identification issue concerning simultaneity between aggregate supply and aggregate demand. If monetary policy is conducted 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 problem is to use deviations from monetary policy rules to identify the trade-off between output and inflation.[3] Barnichon and Mesters (2020) and Barnichon and Mesters (2021) follow this approach and obtain point estimates that suggest the Phillips curve has flattened after 1990, but still has a positive slope of 0.12 to 0.18. However, deviations of monetary policy from policy rules are infrequent and small, and the resulting shocks are weak instruments that yield imprecise estimates.

Given the substantial unresolved identification issues in the estimation of the Phillips curve from aggregate data, a recent literature has started to develop alternative approaches using panel data of regional aggregates or firm

[2]Similar export demand shocks have been used in different contexts e.g. in Hummels et al. (2014),Garin and Silvério (2024), Dhyne et al. (2025) or Boehm and Pandalai-Nayar (2022).

[3]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 does not 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).

level data. We contribute to this line of the literature. Our primary contribution is to use firm-level, rather than regional or sectoral demand shocks, that allow us to identify a lower bound to the Phillips curve slope under weak assumptions. The paper most closely related to ours is Gagliardone et al. (2025), who use firm-level cost shocks to estimate the marginal cost formulation of the Phillips curve. This pins down important Phillips curve parameters, but is not on its own informative about the relationship between inflation and *output*. To estimate this relationship, Gagliardone et al. construct industry-level demand shocks from aggregate monetary policy surprises interacted with the estimated sensitivity of industry output to those surprises. Compared to Gagliardone et al., our approach has a disadvantage and two advantages. The disadvantage is clear: in the best-case scenario, they identify the total slope of the Phillips curve, while our approach yields a lower bound. However, we think the advantages of our approach are important. First, the approach of Gagliardone et al. only recovers the slope of the Phillips curve under the assumption of no labor mobility between 4-digit industry boundaries and would be downward biased otherwise.[4] Second, their approach requires estimation of a large number of unknown first-stage sectoral output elasticities to monetary policy surprises—one for each 4-digit industry in their sample. This leads to a situation with many potentially weak instruments, in which the first stage of an IV estimator is overfitted, and the resulting IV estimates are biased toward the OLS estimate. Since OLS confounds demand and supply shocks, the OLS estimate is itself biased toward zero. Both of these concerns are potential explanations for why our slope estimate exceeds theirs despite being a lower bound.

Our paper is also closely related to the literature estimating Phillips curves using regional variation in prices and output or unemployment, most notably Hazell et al. (2022). Hazell et al. use unemployment and prices of non-tradeables at the state level to identify the slope of the U.S. Phillips curve. 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. The main advantage of our approach is that we can measure output directly at the firm level, while they use state-level unemployment as a proxy for slack in the non-tradeable sector. If labor mobility between sectors is limited, state-level unemployment could mismeasure slack for non-tradeable firms after a shock to tradeable sectors and lead to a downward bias in the slope of the estimated Phillips curve. Other recent papers in this literature include Fitzgerald et al. (2024), who show that the reduced-form relationship between local unemployment and inflation in the U.S. is relatively stable over the 1976–2018 period, and Cerrato and Gitti (2025), who use a similar identification strategy as Hazell et al. to show that the U.S. Phillips curve has substantially steepened during the COVID years.

Finally, our paper is related to work studying firm-level and sectoral responses to export demand shocks more generally. Dhyne et al. (2025) study the response of wages and employment of Belgian firms to a shift-share export demand shocks similar to ours, and find evidence consistent with monopsony power and an important overhead labor component. Our extended model of Section 4.1 is compatible with their findings. Boehm and Pandalai-Nayar (2022) study how the price response to demand shocks at the industry level depends on capacity utilization based on a model with "hard" capacity constraints embodied in a Leontief production function. They find evidence for an important role of initial capacity. Our model is based on a CES production technology which converges to the

[4]The aggregate Phillips curve reflects the pass-through of wage increases required to fulfill labor demand at higher output. In the aggregate, this is determined by the Frisch elasticity of labor supply. At the sectoral level, this depends on labor mobility across sectors. If workers are mobile, a small wage increase might be sufficient to fulfill higher labor demand after a sectoral shock, because workers can be poached from other sectors. Consequently, the elasticity of sectoral wages to sectoral output might be much smaller than the elasticity of aggregate wages to aggregate output.

Leontief case in the limit.

Our paper proceeds as follows. In Section 2, we introduce a New Keynesian model with firm-level demand shocks to illustrate the mapping between firm-level responses and the aggregate Phillips curve. In Section 3, we estimate the reduced-form response of prices, output and other outcomes to firm-level demand shocks. In Section 4, we estimate the structural supply and demand parameters of the model by matching impulse responses to firm-level demand shocks between the model and the data, and use these estimates to quantify the capacity pressure channel of the Phillips curve slope. In Section 5, we discuss the inflation response to monetary policy shocks implied by our estimates and compare it to others in the literature. We conclude in Section 6.

2 The capacity pressure channel in a New Keynesian model

We now introduce a New Keynesian model featuring firm-level demand heterogeneity. This model serves two purposes. First, we use it to illustrate the decomposition of the aggregate New Keynesian Phillips curve into a capacity pressure and a wage pressure channel. Second, we use the model to link the capacity pressure channel to the response of firms' prices, output, and material inputs to firm-level demand shocks. This forms the basis for our empirical approach.

Relative to a textbook model as in Gali (2008), we add three features to the model. First, we add firm-level demand shifters. These shifters do not affect the aggregate behavior of the model, but illustrate the link between the determinants of aggregate dynamics and the response to firm-level shocks. We will estimate key model parameters by matching the response to these shocks between the model and the data. Second, we include a material input in firms' production that is produced in a roundabout structure. This means our model is formulated in gross output rather than value-added terms, and maps cleanly to the data we will use, which covers firm-level output and output prices (rather than value added and value-added prices). Third, our model features a CES production function rather than a standard Cobb-Douglas specification. This introduces an additional parameter that governs the substitutability between fixed and flexible factors and allows for variation in the slope of firms' supply curves independently of factor shares.

Our model describes a closed economy. The empirical evidence we present below uses the *open* nature of the Danish economy for identification. We therefore interpret the closed economy model as a description of the Euro area, to which Denmark is closely connected through its long-standing fixed exchange rate regime[5]. A crucial assumption of our approach is then that Danish manufacturing firms are representative of Euro area firms. We think it is reasonable to extrapolate from Danish to Euro area manufacturing firms, even if they are not exactly the same in all dimensions. However, extrapolating from manufacturing to service-sector firms is a strong assumption. We share this strong assumption with our main references on cross-sectional Phillips curve estimation[6].

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

[6] Gagliardone et al. (2025) also model a closed economy and treat Belgian manufacturing firms as representative of this overall economy, and Hazell et al. (2022) model tradeables (manufacturing) and non-tradeables (services) separately but assume they have the same production functions.

2.1 Model setup

Final output. Gross final output in the model economy is a homogeneous good produced by perfectly competitive firms with flexible prices who use differentiated intermediate goods as inputs in a CES aggregator:

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

$Z_{i,t}$ is an idiosyncratic demand shifter for intermediate goods that follows an AR(1) process with persistence ρ and i.i.d. demand shocks $u_{i,t}$ in logarithms, i.e., $z_{i,t} = \rho z_{i,t-1} + u_{i,t}$. Final good producers choose intermediate inputs to minimize their expenditure $\int_i Y_{i,t} P_{i,t} di$ subject to (1). This results in a standard CES demand curve for intermediate goods $Y_{i,t}$:

$$Y_{i,t} = Y_t Z_{i,t} \left(\frac{P_{i,t}}{P_t} \right)^{-\sigma}. \tag{2}$$

Aggregate gross final output is divided into a value-added component Y_t^{gdp}, and a material component M_t that is used in the production of intermediates. Since final good producers are competitive and use intermediates as their only input, the intermediate CES price index P_t corresponds to the price of gross final output, value-added and materials. The roundabout structure allows us to explicitly account for second-round effects of higher input prices on inflation when aggregating to a Phillips curve.

Intermediate goods production. Intermediates are produced by monopolistically competitive firms who use a fixed factor K and a flexible factor for production. The flexible factor is a Cobb-Douglas composite of flexible labor L and materials M. Firms' overall production function takes a normalized CES form:

$$\frac{Y_{i,t}}{Y_{SS}} = \left(\nu \left(\frac{K_{i,t}}{K_{SS}} \right)^{\psi} + (1-\nu) \left(\left(\frac{L_{i,t}}{L_{SS}} \right)^{1-\alpha} \left(\frac{M_{i,t}}{M_{SS}} \right)^{\alpha} \right)^{\psi} \right)^{1/\psi} \tag{3}$$

In the long-run steady state, firms choose K optimally and operate with constant returns to scale. In the short run K is fixed, and firms operate with decreasing returns. We think of the fixed factor K as "production capacity" broadly defined, which includes capital inputs, but potentially also overhead labor.[7]

Short-run marginal cost varies with output depending on the importance of fixed factors in production ν and the degree ψ to which fixed and flexible inputs can be substituted. This is the source of the capacity pressure channel of the Phillips curve, and without fixed factors ($\nu = 0$) or with perfect substitutability ($\psi = 1$), the capacity pressure channel will reduce to zero. As $\psi \to -\infty$, the production function approaches the Leontief case with a "hard" capacity constraint as in Boehm and Pandalai-Nayar (2022).

[7] In practice, firms may adjust capacity subject to a costs as in Woodford (2005). Since our approach estimates the slope of firm-level supply curves from the observed response to demand shocks, we expect our quantification of the capacity pressure channel to be relatively insensitive to the specific structural source of incomplete short-term capacity adjustment.

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

$$\sum_{k=0}^{\infty}(\beta\theta)^k E_t\left(Y_{i,t+k}(1-\sigma)\left(P^*_{i,t}-\frac{\sigma}{\sigma-1}MC_{i,t+k}\right)\right)=0 \tag{4}$$

$Y_{i,t+k}$ and $MC_{i,t+k}$ depend on the chosen reset price, as well as the state of idiosyncratic demand and aggregate state variables in period $t+k$. The optimal reset price is a weighted average of hypothetical optimal flexible prices determined by the flexible-price supply curve $\sigma/(\sigma-1)MC_{i,t+k}$. We summarize the central supply-side parameters of our model using the slope δ of the log-linearized flexible-price supply curve:

$$\delta \equiv \frac{\partial \log \frac{\sigma}{\sigma-1}MC_{i,t+k}}{\partial \log Y_{i,t+k}} = \frac{(1-\psi)\nu}{1-\nu} \tag{5}$$

In our baseline model with CES aggregation, optimal flexible-price markups are constant and δ equals the output elasticity of marginal cost, but the characterization of the model in terms of δ also applies when optimal flexible-price markups are variable (see the model with Kimball demand in Section 4.1).

We log-linearize price-setting firms' FOC around a symmetric zero-inflation steady state in which the idiosyncratic demand shifters are equal to the mean. That means there is no heterogeneity in prices or quantities in the steady state. We use lower-case letters to denote log deviations of variables from their steady state value. The optimal reset price is a function of firms' discounted current and future nominal marginal cost:

$$p^*_{i,t} = (1-\beta\theta)\sum_{k=0}^{\infty}(\beta\theta)^k E_t\left(mc^R_{i,t+k}+p_{t+k}\right) \tag{6}$$

Under the assumption of competitive factor markets for flexible inputs, the deviation of marginal cost of firm i from its steady-state value is determined by the production function as $mc^R_{i,t+k} = mc^R_{t+k} + \delta(y_{i,t+k}-y_{t+k})$. Plugging in the linearized demand curve, future marginal cost at the current optimal reset price is equal to $mc^R_{i,t+k} = mc^R_{t+k} + \delta(z_{i,t+k} - \sigma(p^*_{i,t}-p_{t+k}))$. Together with $E(z_{i,t+k}|z_{i,t}) = \rho^k z_{i,t}$, the optimal reset price can be expressed as a function of the current realization of the firm-level demand shifter and current and expected aggregate states:

$$p^*_{i,t} - p_{t-1} = \frac{(1-\beta\theta)}{1-\beta\theta\rho}\frac{\delta}{1+\sigma\delta}z_{i,t} + \frac{(1-\beta\theta)}{1+\sigma\delta}\sum_{k=0}^{\infty}(\beta\theta)^k E_t\left(mc^R_{t+k}\right) + \sum_{k=0}^{\infty}(\beta\theta)^k E_t(\pi_{t+k}), \tag{7}$$

2.2 The response to firm-level shocks

We now illustrate how an average firm's observed prices and output respond to firm-level demand shocks. We denote as $\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})$ the average period $t+k$ price of firms with a period t history of idiosyncratic demand realizations $\mathcal{Z}_{i,t-1} = \{z_{i,\tau}\}_{\tau=-\infty}^{t-1}$ and current demand realization $z_{i,t}$, relative to the aggregate price level. Similarly, we denote as $\tilde{p}^*_{t+k}(z_{i,t})$ the average period $t+k$ reset price of firms with period t demand realization $z_{i,t}$, relative to the average reset price. While $\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})$ depends on the full history of demand realizations—because firms might not get a chance to reset their price—the reset price $\tilde{p}^*_{t+k}(z_{i,t})$ is not backward-looking and only depends on the period t realization of demand. Demand shocks received after period t are orthogonal to the history up to time t and average out. The observed period $t+k$ relative price is the average over firms' relative reset price and lagged relative price:

$$\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t}) = (1-\theta)\tilde{p}^*_{t+k}(z_{i,t}) + \theta\tilde{p}_{t+k-1}(\mathcal{Z}_{i,t-1}, z_{i,t}) \tag{8}$$

To derive the impulse response of relative prices to demand shocks, we differentiate (8) w.r.t. the demand shock $u_{i,t}$. The resulting recursion is:

$$\frac{\partial\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = (1-\theta)\frac{\partial\tilde{p}^*_{t+k}(z_{i,t})}{\partial u_{i,t}} + \theta\frac{\partial\tilde{p}_{t+k-1}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{9}$$

Since period t demand shocks are orthogonal to the history of demand realizations up to that point, we get as a boundary condition that $\partial\tilde{p}_{t-1}(\mathcal{Z}_{i,t-1}, z_{i,t})/\partial u_{i,t} = 0$. From equation (7), the derivative of the relative period $t+k$ reset price is:

$$\frac{\partial\tilde{p}^*_{t+k}(z_{i,t})}{\partial u_{i,t}} = \frac{1-\beta\theta}{1-\beta\theta\rho}\frac{\delta}{1+\sigma\delta}\rho^k \tag{10}$$

This allows us to solve the recursion (9). The resulting impulse response of prices to demand shocks is:

$$\frac{\partial\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \underbrace{\frac{\delta}{1+\sigma\delta}}_{\substack{\text{Flexible-price}\\\text{response}}} \cdot \underbrace{\frac{(1-\beta\theta)(1-\theta)}{1-\beta\theta\rho}}_{\substack{\text{Sticky price}\\\text{attenuation}}} \cdot \underbrace{\frac{\theta^{k+1}-\rho^{k+1}}{\theta-\rho}}_{\text{Dynamics}} \tag{11}$$

The magnitude of the price response depends on three terms, two static and one dynamic. The first term is given by the hypothetical instantaneous response to a demand shock under flexible prices. This term depends on the slopes of demand and the hypothetical flexible-price supply curves. The second term attenuates this response because shocks are transitory and prices are sticky. The last term determines the dynamics of the response and allows for three types of patterns. If idiosyncratic demand shocks are permanent, i.e., $\rho = 1$, then in response to a positive demand shock, the firm's relative price will converge to a permanently higher level. If demand shocks are transitory and $\rho > (1-\theta)$, the shock decays more slowly than prices adjust, and the response is hump-shaped—relative prices increase over several periods initially before slowly returning to zero. Finally, if $\rho < (1-\theta)$, demand shocks decay faster than prices adjust, and firms' relative price increases in the first period and then slowly returns to zero.

Given the response of prices, the response of output and flexible inputs follows from the demand curve and firms cost-minimizing input choices:

$$\frac{\partial\tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \rho^k - \sigma\frac{\partial\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{12}$$

$$\frac{\partial\tilde{m}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{\partial\tilde{l}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{1}{1-\nu}\frac{\partial\tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{13}$$

2.3 The capacity pressure channel

The demand shifter $z_{i,t}$ averages out across firms, and the model features the same aggregate dynamics as a version without demand heterogeneity. Around a zero inflation steady state, aggregate inflation is approximately equal to $\pi_t = (1-\theta)\left(\int_i p^*_{i,t}di - p_{t-1}\right)$. The marginal cost Phillips curve follows from combining this definition with equation (7):

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

The marginal cost coefficient equals $\lambda = (1-\beta\theta)(1-\theta)/(\theta(1+\sigma\delta))$. The deviation of aggregate real marginal cost from its steady-state value is equal to $mc^R_t = (1-\alpha)w^R_t + \delta y_t$, where w^R_t represents the deviation of the real wage and y_t the deviation of aggregate gross output from its steady-state value. The relationship between aggregate gross output and GDP around the steady state is given by:

$$y_t = \frac{(1-s_m)(1-\nu)}{1-\nu-s_m}y^{gdp}_t + \frac{s_m(1-\alpha)(1-\nu)}{1-\nu-s_m}w^R_t \tag{15}$$

$s_m = \alpha(1-\nu)/\mu$ denotes the steady-state share of materials in gross output. Compared to standard roundabout models, the presence of a fixed factor drives a wedge between gross output and GDP. Since gross output is produced with decreasing returns to scale, and some gross output is used as material, a given increase in GDP requires a larger corresponding increase in gross output. To derive the Phillips curve in terms of GDP, we define a simple reduced-form wage Phillips curve $w^R_t = \phi y^{gdp}_t$ that summarizes the GDP elasticity of real wages after an aggregate demand shock in a single parameter ϕ.[8] The Phillips curve in terms of deviations of GDP then takes its usual form:

$$\pi_t = \kappa y^{gdp}_t + \beta E_t(\pi_{t+1}). \tag{16}$$

The slope of the Phillips curve κ is the sum of two components reflecting capacity pressure and wage pressure:

$$\kappa = \underbrace{\overbrace{\frac{\delta}{1+\sigma\delta}}^{\text{Flexible-price response}} \overbrace{\frac{(1-\beta\theta)(1-\theta)}{\theta}}^{\text{Sticky price attenuation}} \overbrace{\frac{(1-s_m)(1-\nu)}{1-\nu-s_m}}^{\text{Roundabout amplification}}}_{\text{Capacity pressure}\equiv\kappa^y} + \underbrace{\frac{\phi(1-\alpha)}{1+\sigma\delta}\frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{(1-s_m)(1-\nu)-\psi\nu s_m}{1-\nu-s_m}}_{\text{Wage pressure}\equiv\kappa^w}. \tag{17}$$

The capacity pressure channel κ^y results from upward-sloping firm-level supply curves, which in turn results from production inputs that are fixed in the short run. The strength of the capacity pressure channel depends on three terms. First, the hypothetical instantaneous response of prices to a demand shock under flexible prices. If there are no fixed factors in production ($\nu = 0$), or fixed and flexible inputs are perfect substitutes ($\psi = 1$), supply curves are flat ($\delta = 0$) and the capacity pressure channel does not add to the slope. In the extreme case of Leontief production ($\psi \to -\infty$), this term is bounded by $1/\sigma$, as output does not respond to shifts in demand and prices move along the demand curve (Boehm and Pandalai-Nayar, 2022). The second term attenuates the flexible-price response due to price stickiness (under standard values for θ and β). Finally, roundabout production amplifies the capacity pressure

[8] In a model that features a specification of household decisions, ϕ is typically determined by the Frisch labor supply elasticity and households' intertemporal elasticity of substitution. For our purposes, going into more detail about the household side of the economy would not add additional insights, and we choose to summarize all relevant decisions by the single ϕ parameter.

channel, as increasing prices of intermediate inputs also subject to capacity pressure get passed into prices of final output. This roundabout amplification term is equal to 1 if there are no material inputs ($\alpha = 0$) or if there is no fixed factor ($\nu = 0$).

The second component κ^w reflects the "wage pressure" channel that results from co-movement of real wages and GDP after aggregate demand shocks. For a given sensitivity of real wages to GDP ϕ, the strength of this channel depends on the importance of labor in firms' flexible input $(1 - \alpha)$, which determines the relevance of wages for marginal cost. As with capacity pressure, the Calvo factor attenuates the response due to price stickiness, and the demand feedback factor $1/(1+\sigma\delta)$ dampens the channel as supply curves steepen—when price-adjusting firms raise prices in response to higher wages, demand decreases, partially offsetting the initial cost pressure. With roundabout production, there is additional amplification captured by the final multiplicative term, as higher real wages lead to a substitution from labor to materials, which raises marginal cost through capacity pressure on materials production.

There is a close connection between the response of prices to firm-level demand shocks and the capacity pressure channel of the Phillips curve—both depend on the hypothetical flexible-price response of prices to demand shocks $\delta/(1 + \sigma\delta)$. An important qualitative difference is that aggregate demand shocks generate a transitory increase in inflation and a permanent increase in the price level. Positive firm-level demand shocks generate a transitory increase in relative prices that fades away over time (unless the shocks themselves are permanent). Nevertheless, we can use estimates of firms' response to demand shocks to quantify the capacity pressure channel. In Section 3 we will estimate the firm-level response to demand shocks corresponding to equations (11) to (13) in the model. In Section 4, we map these reduced-form estimates to structural parameters and use our estimates to quantify the capacity pressure channel.

The response to firm-level shocks is not informative about the sensitivity of real wages to GDP ϕ and thus the wage pressure channel of the price Phillips curve. Our empirical approach below relies on estimating the dynamics of firm-level relative to aggregate outcomes after firm-level shocks. These relative changes are informative about the supply and demand parameters required to quantify the capacity pressure channel. But under our baseline assumption of competitive labor markets, the model does not feature variation in firm-level wages after demand shocks. We extend the model with monopsonistic labor markets in Section 4.1 to allow for wage effects at the firm level. Even then, wage effects are only informative about labor supply elasticities at the firm level (i.e., the strength of monopsony power) but not about the aggregate labor supply elasticity required to quantify the wage pressure channel. To quantify this channel requires demand shocks at the level of labor markets as in Hazell et al. (2022) and is outside the scope of this paper.

3 Estimating the firm-level response to demand shocks

In this section, we estimate the response of prices, output, intermediate purchases, and other outcomes of Danish manufacturing firms to firm-level demand shocks. Our identification strategy is based on demand shocks that arise from heterogeneity in firms' export exposure to fluctuations in import demand in different destination countries. Firms persistently export different products to different destination countries, and aggregate fluctuations in these countries are not perfectly synchronized. This leads to cross-sectional variation in firm-level demand that we exploit.

3.1 Data

Our empirical analysis is based on administrative and survey data covering production, sales, prices and input choices of Danish manufacturing firms. We combine 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 micro data we use is available at a quarterly or monthly frequency, trade data covering a large enough sample of countries over a sufficiently long period of time is only available at the annual level. Consequently, we conduct our empirical analysis at the annual frequency. Our sample covers the 1999–2021 period.

Producer price index microdata. Our main source of price data is the Danish Producer Price Index (PPI) survey. The Danish 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 CN codes. Firms mainly report domestic prices. Some firms also provide export prices, but the survey does not contain information on the export destination. The reported prices are transaction prices in Danish kroner, including temporary sales and discounts. The survey is designed to allow adjustments for quality changes and product substitutions.[9] The dataset is very balanced, with few gaps in price series. We perform quality adjustments and winsorize price changes at ± 1 log points in the monthly data. We then transform the dataset to annual frequency by keeping the price in the last month of each year.

Production and sales microdata. We use data on sales, production and exports of Danish manufacturing firms collected from two main sources. Data on total sales and production comes from a 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. Our main measure of firm-level output is total annual sales deflated with a firm-specific price index based on unit-value changes. We describe the construction of this index in the online Appendix B.1. Data on export and import values and quantities are based on administrative survey and customs data (UHDM). This data is collected for all exporters and importers above a small yearly volume cutoff and provides monthly export sales, import purchases and corresponding quantities at the level of 8-digit CN product categories. We use the export data to construct the shares in our shift-share demand shifter below.

Other firm-level data. We complement these datasets with basic firm information from annual balance sheets available in the Danish business register (FIRM) and the Danish accounting statistics (FIRE). We use total annual full-time equivalent employment, the annual wage sum and annual intermediate purchases from these datasets. The variables we use from these datasets are available for the universe of Danish firms. Finally, we use survey micro data on self-reported capacity utilization from the Danish Business Sentiment survey (Konjunkturbarometer). This dataset covers roughly 450 manufacturing firms.

Macro data on imports and exports. We use country-level data on imports and exports during the from the UN Comtrade database. Comtrade covers trade flows between a source and a destination country at the product level.

[9] 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.

Our baseline analysis uses flows at the 4-digit Harmonized System (HS) code level. We construct country-product level import growth rates that exclude imports from Denmark. These imports will serve as shifts in our shift-share demand shifter, and we leave out imports from Denmark in the construction to rule out a source of possible reverse causality.

Sample description. Our baseline sample of firms covers manufacturing firms that report domestic prices in the PPI or output in VARS. We restrict the sample to firms that have more than 20 employees for more than five consecutive years during the sample period. This excludes small firms that would otherwise move in and out of the sample as they cross the reporting threshold for VARS. We also require firms to have an export share of at least 1% of their total sales in the period before they are hit by a demand shock.

This leaves us with two samples over the 1999–2021 period. An output sample of about 2,400 firms who report output in VARS, and a price sample of about 800 firms who report prices in the PPI. In our baseline estimation, we use all firms who report output to estimate the effects of demand shocks on output and other firm-level outcomes, and the smaller sample of firms who report prices when estimating the effects of demand shocks on prices. In robustness checks, we restrict the estimation of firm-level outcomes to the smaller price sample, or use an alternative output price index that is available for all firms who report output in the larger output sample.

We provide descriptives for both samples in Appendix Table 2. The average firm in our output sample has 136 employees and sales of 52 million euro. Firms in the price sample are larger, with an average of 236 employees and sales of 100 million euro. Both are rather small by global standards, as Danish manufacturing is dominated by small and medium-sized enterprises. However, the samples 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 exports are about 60% of average sales in both the output and price sample. Firms in the output and price sample on average export products in 14 and 18 distinct 4-digit HS categories to 19 and 25 different export destinations, respectively. Finally, firms in the price sample on average report prices for 6 distinct products.

3.2 Estimation and Identification

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

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

$$\Delta z_{i,t} = \sum_{k \in K} \sum_{j \in J} \omega_{i,k,j,t-1} \Delta im_{k,j,t} \tag{18}$$

Note that the shares $\omega_{i,k,j,t-1}$ will generally not add up to one due to domestic sales. We calculate the total coverage of the demand shifter as $\Omega_{i,t-1} = \sum_{k \in K} \sum_{j \in J} \omega_{i,k,j,t-1}$ and recenter the shift-share demand shifter following Borusyak et al. (2022) by controlling for total coverage. In our baseline analysis, we winsorize $\Delta z_{i,t}$ at the 5th and 95th percentiles each year to deal with outliers. Online Appendix B.2 provides additional details on the construction of the demand shifters.

Properties of the demand shifters. Figure 1 shows properties of the demand shifters we construct. Panel (a) shows the mean and cross-sectional standard deviation of $\Delta z_{i,t}$. Since imports grow with output, our demand shifter is on average positive and varies with international business cycles. Our estimation only uses cross-sectional within-sector variation, as the mean will be absorbed in fixed effects. The within-sector cross-sectional standard deviation amounts to 4.5% on average over the sample period and increases to 10% during the Great Recession. A standard diagnostic for shift-share instruments is the Herfindahl index of exposure shares and its inverse, the effective number of shocks.[10] Borusyak et al. (2022) note that a Herfindahl index below 0.1 should make standard inference work reasonably well. We compute a Herfindahl index of exposure shares of 0.00006, corresponding to an effective number of about 15,000 shifts in the product and country dimensions used for estimation. The Herfindahl index considering only the country dimension is 0.004, corresponding to an effective number of about 230 shifts.

The demand shifter features very little autocorrelation in first differences. A regression of $\Delta z_{i,t}$ on its lag and time-sector fixed effects yields an insignificant coefficient of about 0.013. Consequently, we can reasonably treat $\Delta z_{i,t}$ as a demand *shock*. In panel (b) of Figure 1, we plot the results of a local projection of $\Delta z_{i,t}$ on its cumulative sum to show that shocks decay slowly over five years. The dynamics of the cumulative sum over five years are well-described by an AR(1) process with annual persistence 0.97. In our empirical estimation described below, we will focus on two types of specifications: one that treats the demand shifter as a shock directly, and one that controls for lagged values as is commonly done in the literature using local projections.

Reduced-form estimation. We estimate the response of prices, output and other outcomes to firm-level demand shocks using panel local projections following Jordà (2005). These local projections are the direct counterpart to equations (11) to (13) in the model. For output and other firm-level outcomes, we estimate local projections at the firm level in the output sample:

$$\Delta^h y_{i,t+h} = b^h \Delta z_{i,t} + a^h_{s(i),t} + \mathbf{c}^{h\prime} \mathbf{X}_{i,t} + v^y_{i,t+h} \tag{19}$$

[10]This is especially important when relying on a shock-based identification (see Borusyak et al. (2025)). We rely on a shares-based identification, but it is still informative for a shares-based perspective to know that the exposure shares are not too concentrated and a large set of products and export destinations is used for identification.

Figure 1: Properties of the demand shifter

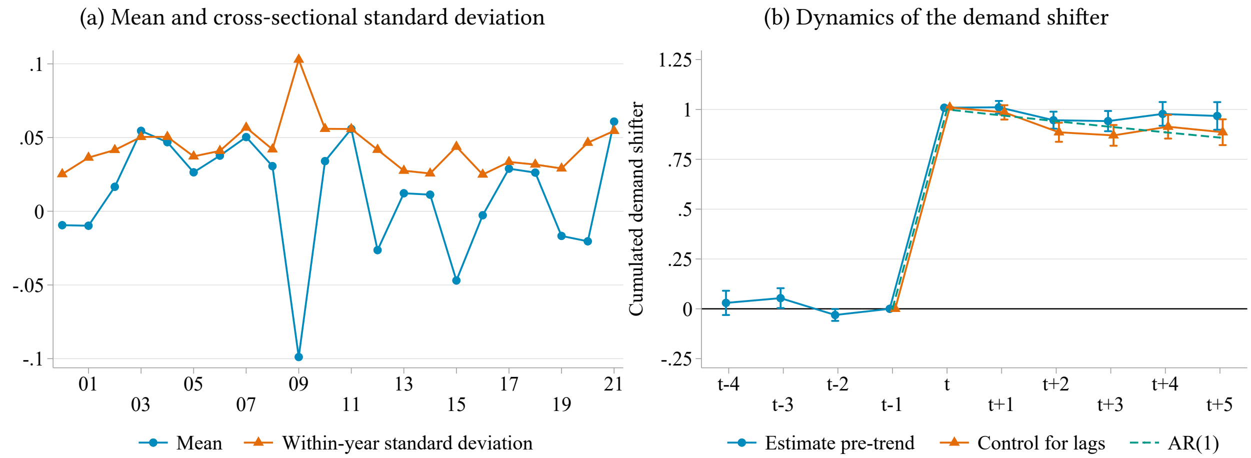

Notes: The figure plots properties of the demand shifter defined by equation (18). Panel (a) plots the mean and cross-sectional standard deviation of the demand shifter over time. Panel (b) plots a local projection of the demand shifter on its own cumulated sum, estimated using equations (19). Error bars denote 95% confidence intervals based on standard errors clustered at the firm level.

For prices, we estimate otherwise identical local projections with domestic prices at the firm-product level as the outcome in the price sample:

$$\Delta^h p_{i,j,t+h} = b^h \Delta z_{i,t} + a^h_{s(i),t} + \mathbf{c}^{h\prime}\mathbf{X}_{i,t} + v^p_{i,j,t+h} \tag{20}$$

We estimate local projections in the two different samples to maximize the precision of the estimated responses. To alleviate concerns that this might affect our results, we present robustness checks that estimate all firm-level responses in the smaller sample that includes only firms that appear in the PPI, and use the unit-value price index used to deflate sales to output as a price outcome in the larger sample. Our baseline estimates include 2-digit industry-time fixed effects $a_{s(i),t}$ that absorb supply shocks at the sector level—for example the development of real wages or the price of intermediate inputs—and inflation expectations. Moreover, we always control for the sum of lagged export shares interacted with industry-time dummies to recenter the shift-share instrument and account for possibly incomplete coverage as suggested in Borusyak et al. (2022). We use only domestic prices as our outcome variable in the baseline, since export prices might be directly affected by aggregate economic conditions in the destination in a way that cannot be absorbed by fixed effects (for example, through local distribution costs or pricing-to-market) and is not informative about firms' capacity pressure. To deal with outliers, we winsorize differences at each horizon Δ^h. We use the 5th and 95th percentiles as thresholds for prices and output, and the 2nd and 98th percentile for other firm-level outcomes that are less noisy.

We estimate two baseline specifications. Our baseline specification "without controls" treats $\Delta z_{i,t}$ as a demand shock as is and includes only the interacted coverage of the shift-share demand shock as control. For this specification, we estimate placebo coefficients at negative horizons h that would detect possible violations of the parallel trends assumption, for example due to autocorrelation in $\Delta z_{i,t}$. In our baseline "with controls" we add further controls in the vector $X_{i,t}$ to preclude the possibility that the small autocorrelation in $\Delta z_{i,t}$ affects our results. These controls include lags of $\Delta z_{i,t}$ and lagged first differences of endogenous variables. The disadvantages of the latter approach are that we can no longer estimate placebo coefficients at negative horizons h, and that in a short panel, including lagged endogenous variables might introduce Nickell bias (see Anderson and Hsiao, 1982, Arellano and Bond, 1991).

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

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

IV estimation. The primary focus of our empirical evidence lies on the reduced-form regressions described above. An ex-ante interesting alternative to directly identify the flexible-price supply curve δ could be an IV specification that uses the demand shock as an instrument for output. This would recover δ in a scenario with flexible prices, but does not recover structural parameters in a sticky-price setting. We can use equations (6) and (8) of our model to derive firms' forward looking supply curve in terms of observed relative prices:

$$\tilde{p}_t(\mathcal{Z}_{i,t-1}, z_{i,t}) = \theta \tilde{p}_{t-1}(\mathcal{Z}_{i,t-1}, z_{i,t}) + (1-\theta)(1-\beta\theta)\delta \sum_{k=0}^{\infty} (\beta\theta)^k E_t[\tilde{y}_{t+k}(\tilde{p}_t^*(z_{i,t}))] \tag{21}$$

The current observed relative price depends on the lagged relative price and the discounted sum of expected future output at the *current relative reset price* $\tilde{p}_t^*$. However, in the data, we can at best observe realized output at the actual period $t+k$ relative price $\tilde{p}_{t+k}$. Even if forecast errors are random, the difference between those two prices depends on the current shock realization $u_{i,t}$ and unknown model parameters (this follows directly from equations (10) and (11) of the model). Consequently, if we were to construct and use the discounted sum of future output at realized prices, the IV exclusion restriction would be violated and we would not be able to infer δ.

We provide IV estimates of the elasticity of prices to a demand shock that increases period t output for completeness, but since this elasticity depends on the persistence of shocks and other parameters it should not be seen as a structural parameter.[11] In our IV specification, we regress price changes over horizon $t+h$ on firms' output growth $\Delta y_{i,t}$ in

[11] From equations (11) and (12), we can see that the IV coefficients correspond to $X_k/(1-\sigma X_0)$, where X_k is the reduced-form response of prices at horizon k.

period t, using $\Delta z_{i,t}$ as an instrument for $\Delta y_{i,t}$:

$$\Delta^h p_{i,j,t+h} = b^h \Delta y_{i,t} + a^h_{s(i),t} + \mathbf{c}^{h\prime} \mathbf{X}_{i,t} + v_{i,j,t+h} \tag{22}$$

In line with our reduced-form estimation in two samples, we implement the two-sample IV estimator of Inoue and Solon (2010) and estimate the first-stage regression of output on demand shocks at the firm level in the output sample, and the structural equation using prices at the product level in the price sample. We follow Pacini and Windmeijer (2016) to estimate clustered standard errors for the two-sample TSLS estimator. In our IV estimates, we control for lagged shocks, current and lagged values of the shift-share "supply shock" discussed above, and the coverage of the shift-share interacted with time-sector fixed effects.

3.3 Estimation results

Effects on output and capacity utilization. We present our main results starting with the effect of the demand shock on firms' output in Figure 2 panel (a). These results test whether the demand shock we construct is relevant and actually shifts demand. The two baseline specifications with and without controls yield very similar results. Output increases initially by about 0.4 log points for a unit increase in export-weighted demand, and then slowly declines. Consistent with the high persistence of the demand shock, output remains elevated by 0.27 log points five years after the shock. The placebo pre-shock coefficients are all close to zero and insignificant. Coefficient estimates, standard errors and summary statistics for both specifications are reported in Appendix Table 3.

For a smaller subset of firms, we observe self-reported capacity utilization. In panel (b), we show that the capacity utilization rate increases by about 0.14 to 0.20 percentage points for a unit increase in our demand shock. This is supportive evidence for our key assumption that some factors cannot be freely adjusted in the short run. In line with the idea that firms can adjust capacity over time, the effects on capacity utilization remain positive after five years but are insignificant and smaller than the effect on output.

Effects on prices. In Figure 3, we show the effects of demand shocks on prices. Panel (a) shows the reduced-form effect of demand shocks on domestic prices in the PPI. Prices increase by about 0.11 log points for a unit increase in the demand shifter on impact, and continue to rise up to a maximum increase of about 0.2 log points in the third year after the shock hits. Afterward, 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 the decline in prices after year three appears only in our specification with controls. Coefficient estimates, standard errors and other summary statistics for the effects on prices are reported in Table 5. Panel (b) of Figure 3 shows the response for the unit-value-based output price index that we use to deflate sales to output. This variable is a measure of prices for the whole production of a firm (not just domestic prices) and is available for all firms for which we compute output. The response of the output price index is very similar to the response of domestic PPI prices.

In panel (c) of Figure 3, we show the results for IV specifications that instrument the output response on impact with the demand shock. The estimates suggest that a demand shock that increases output by 10% on impact leads to a price increase of about 5% after three years. As in our reduced-form estimates, we find no significant pre-trends and the coefficient estimates from the two baseline specifications with and without controls are very similar. The

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

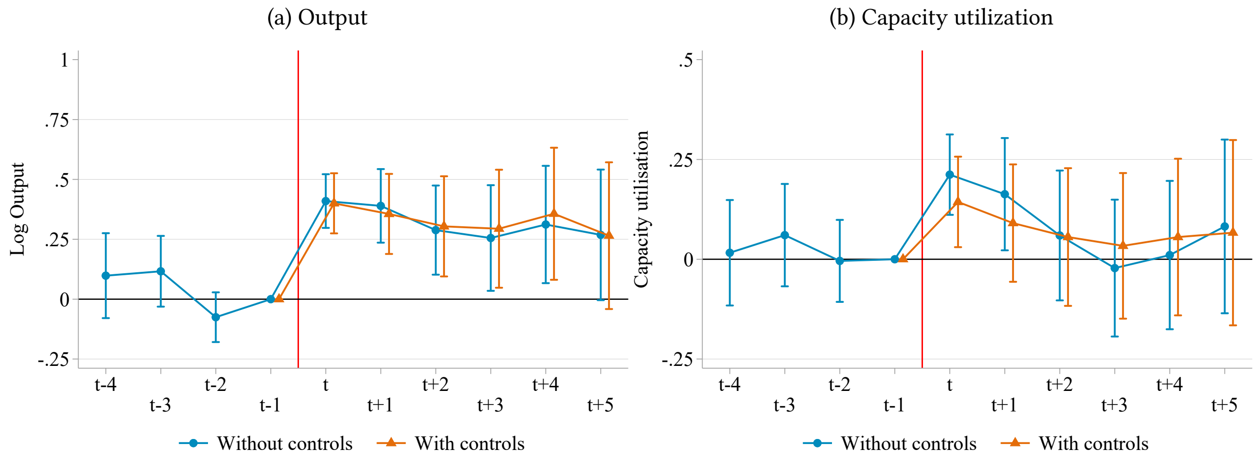

Notes: The figure plots local projections of the log output and the capacity utilization rate on the demand shock as described by equation (19). Specifications "without controls" only control for fixed effects and the coverage of the shift-share demand shock. Specifications "with controls" control in addition for the lagged demand shock, lagged first difference of the outcome in question, and current and lagged import-weighted "supply shocks". The error bars show 95% confidence intervals based on standard errors clustered at the firm level.

elasticity increases over the first four years after a shock and then starts to decline. Table 7 columns (1) and (2) present the coefficients for our two baseline specifications, as well as other summary statistics. The cluster-robust first-stage F statistic for our baseline specification with controls amounts to 39, suggesting our demand shocks are a strong instrument.

Effects on other outcomes. In Figure 4 and Appendix Table 9 we present the local projection results for several other firm outcomes. In panel (a), we show results on sales. These results are consistent in magnitude with the combined effect on prices and output. In panel (b), we show that intermediate input purchases increase, and that the response of intermediate purchases is larger than the response of output. That is consistent with decreasing returns to scale due to fixed factors. In panel (c), we show that employment also increases; however, the short-run elasticity of employment is lower than the elasticity of output at the same horizon. This is consistent with an important overhead employment component that cannot be adjusted in the short run. In panel (d), we show that firms' average wages increase slightly after firm-level shocks. This is consistent with some degree of monopsony in labor markets.

Robustness of baseline results. We estimate effects for several alternative specifications and sample restrictions to check the robustness of our main results on prices and output. First, we vary the regression specification by including firm fixed effects that control for trends in our—already differenced—local projection estimation. Second, we use 4-digit time-sector fixed effects instead of the 2-digit fixed effects in the baseline. Third, we estimate the local projections with time fixed effects only. Fourth, we winsorize prices and output at the 2nd/98th percentiles instead of the 5th/95th in our baseline. The reduced-form results for output and prices as outcomes are shown in Tables 3 and 5, and for the IV estimates in Table 7 in the Appendix. Across specifications, the results remain similar in terms of magnitude.

Figure 3: Effects on prices

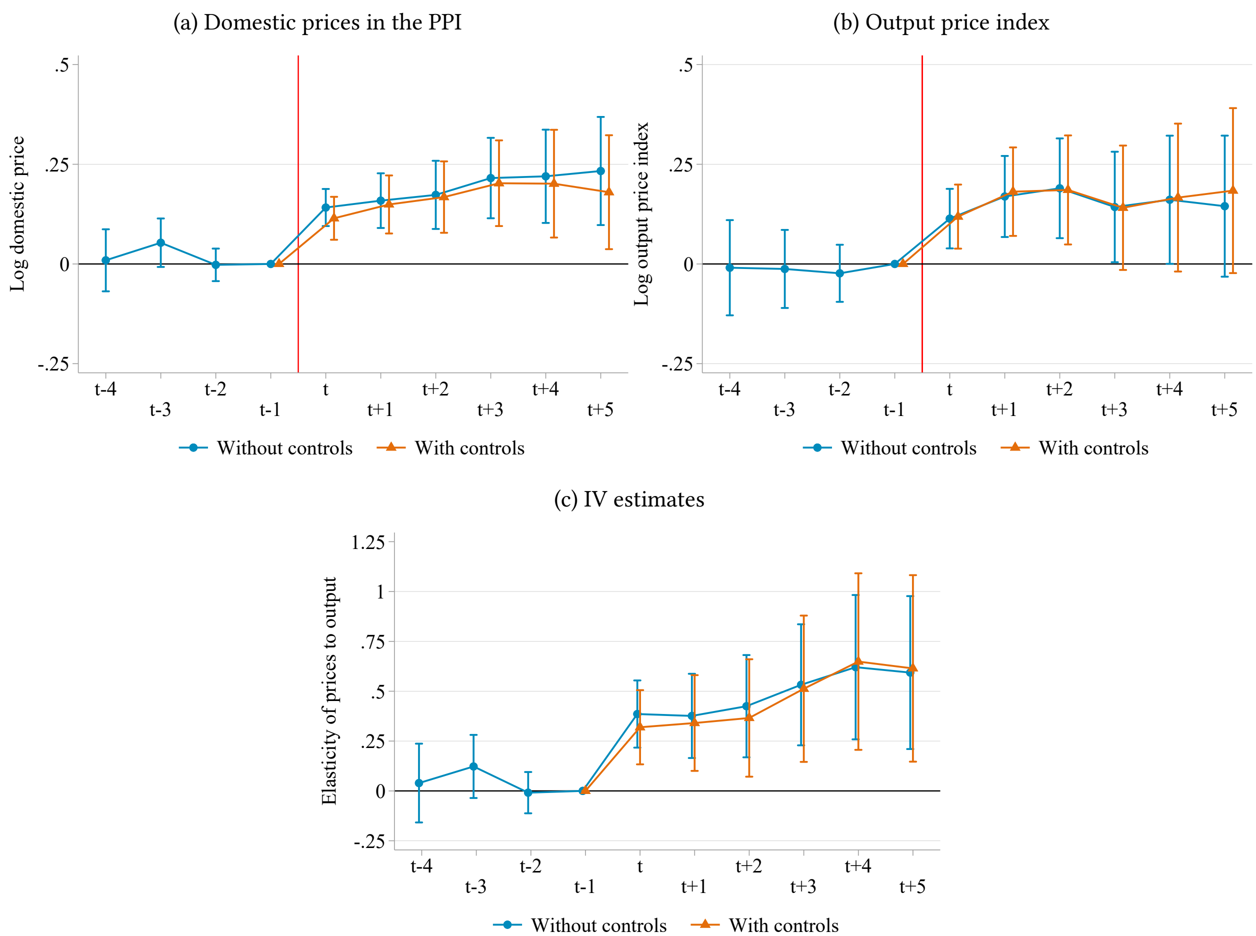

Notes: Panels (a) and (b) plot local projections of the log prices in the domestic PPI and the log output price index on the demand shock as described by equation (20). Panel (c) plots the IV-local projection of PPI prices on output growth, using the demand shock as an instrument as described by equation (22). Specifications "without controls" only control for fixed effects and the coverage of the shift-share demand shock. Specifications "with controls" also control for the lagged demand shock, lagged first difference of the outcome in question, and current and lagged import-weighted "supply shocks". The error bars show 95% confidence intervals based on standard errors clustered at the firm level.

We also show that our results are robust to variations of our sample restrictions. The reduced-form results for output and prices as outcomes are shown in Tables 4 and 6, and for the IV estimates in Table 8 in the Appendix. Our baseline restrictions are that firms have more than 20 employees for at least five years in a row during the sample period and are in a manufacturing sector. In addition, we restrict the sample of our local projections to firms that have an export share of at least 1% in the year prior to a given shock. We first restrict the sample used to estimate the effect on output to only firms that can be matched to the PPI micro data. The effect on output becomes slightly larger but remains comparable. Moreover, for both output and price outcomes, we estimate effects for a more balanced panel in which firms are required to have more than 20 employees for at least 10 years in a row in column (3), and restrict the panel to firms with an export share of at least 25% in column (4). We restrict the sample to data before 2020 to exclude the COVID period in column (5). The results for those three variations are very close to our baseline estimates. In

Figure 4: Effects on other outcomes

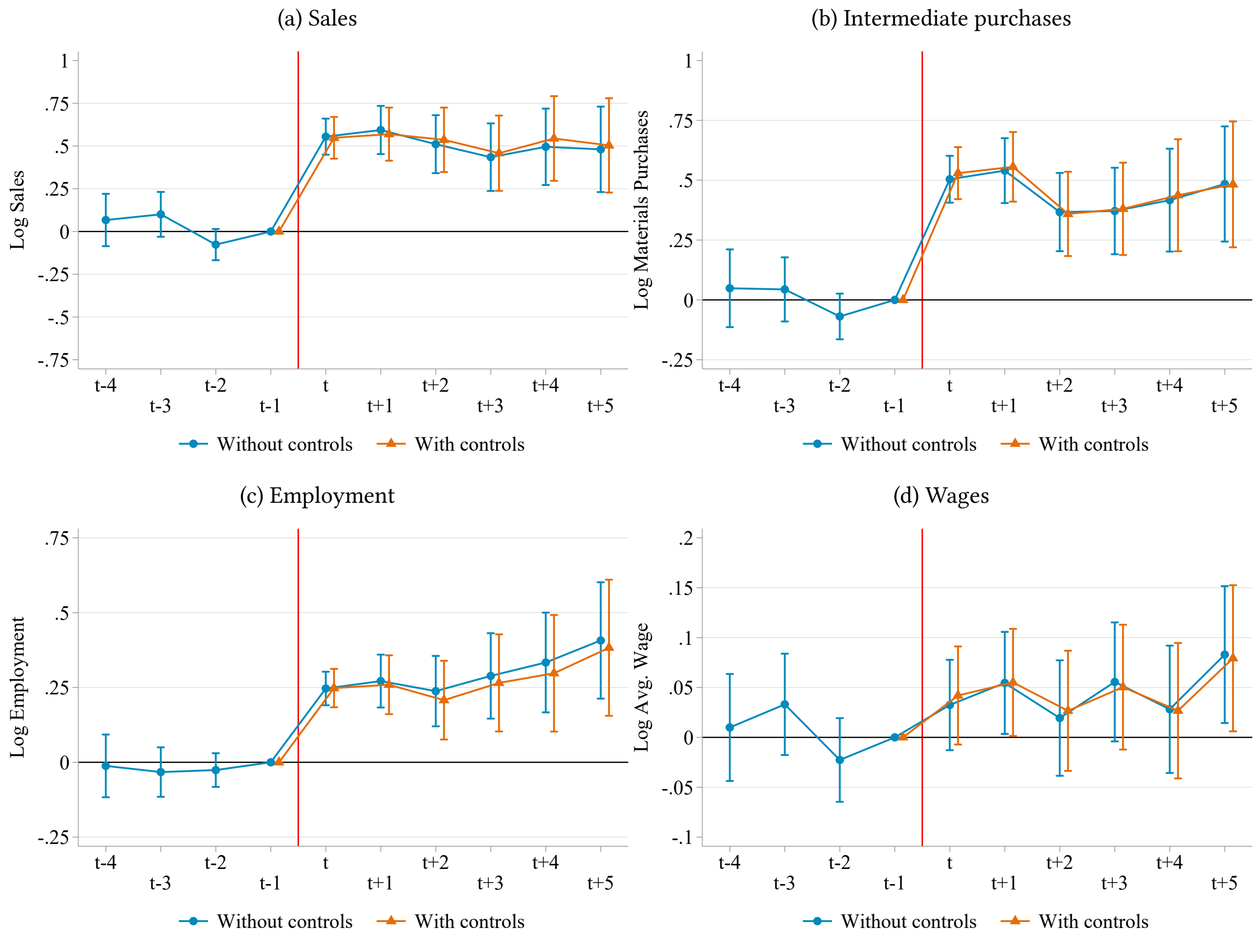

Notes: The figure plots local projections of the log sales, log intermediate purchases, log employment and log average wages on the demand shock as described by equation (19). Specifications "without controls" only control for fixed effects and the coverage of the shift-share demand shock. Specifications "with controls" control in addition for the lagged demand shock, lagged first difference of the outcome in question, and current and lagged import-weighted "supply shocks". The error bars show 95% confidence intervals based on standard errors clustered at the firm level.

columns (6) and (7), we construct demand shifters leaving out exports to Denmark's largest export partners Germany and Sweden. Again the results remain close to our baseline estimates.

4 Mapping the model to firm-level results

We now estimate key parameters of the model outlined in Section 2 by fitting the impulse responses estimated in Section 3 to their model counterparts. Our approach follows the impulse-response matching literature (Christiano et al., 2005), but instead of matching aggregate impulse responses, we match the firm-level impulse responses of prices, output, material inputs and sales to firm-level demand shocks. The estimated model thus produces aggregate behavior consistent with the firm-level response to shocks.

Setup. We formulate the model in quarterly frequency and aggregate impulse responses to annual frequency consistent with the aggregation of the data. In particular, we calculate responses for firms hit by a given shock in quarters one to four of year zero, and then aggregate the responses using the average log deviation in the last quarter of a year for prices, and yearly averages over quarterly log deviations for output, material inputs and sales. We calibrate the quarterly discount factor to $\beta = 0.99$, consistent with an annual interest rate of about 4%, and set the degree of price stickiness to $\theta = 0.65$, consistent with the frequency of price changes in the Danish PPI microdata. We set the quarterly persistence of idiosyncratic demand shocks to $\rho = 0.992$, which matches a persistence of about 0.97 in our annual demand shifter.

Impulse response matching. The key model parameters are the slope of firms' flexible-price supply curve δ, the demand elasticity σ, the fixed factor share ν and the share of materials in the flexible factor α. We estimate these parameters from matching the impulse response of prices, output, sales and material inputs to firm-level demand shocks, subject to the constraint that the aggregate materials share in gross output matches its empirical counterpart.

We make one adjustment to the model laid out in Section 2 and normalize the demand shifter with a scaling factor ζ—since the exposure shares of our shift-share demand shock do not add up to one (due to domestic sales) our shock has no natural scale. The normalization means that our model is identified from the relative magnitude of the empirical impulse responses, rather than their absolute magnitudes, and ζ can be understood similarly to the first-stage coefficient of an IV estimator. Firms' linearized demand is then equal to $y_{i,t} = \zeta z_{i,t} - \sigma(p_{i,t} - p_t)$ and the quarterly impulse responses are given by:

$$\frac{\partial \tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \zeta \cdot \frac{\delta}{1+\sigma\delta} \cdot \frac{(1-\beta\theta)(1-\theta)}{1-\beta\theta\rho} \cdot \frac{\theta^{k+1} - \rho^{k+1}}{\theta - \rho} \tag{23}$$

$$\frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \zeta \cdot \rho^k - \sigma \cdot \frac{\partial \tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{24}$$

$$\frac{\partial \tilde{py}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{\partial \tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} + \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{25}$$

$$\frac{\partial \tilde{m}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{1}{1-\nu} \cdot \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{26}$$

Without normalizing by ζ, we could directly back out δ and σ from the impulse response of prices and output at any horizon given the calibrated values of ρ, β and θ. With measurement error, we can still identify ζ, δ and σ from those two equations, but the minimal requirement is to observe the impulse response in t and $t+1$. Intuitively, it is the dynamics of the relative responses over time that allow us to separate the scale parameter ζ from the economically meaningful parameters σ and δ. The importance of the fixed factor ν is identified from the response of materials purchases relative to output, and allows us to decompose δ into its components. If the response of materials is larger than the response of output, this is an indication for the presence of fixed factors. The response of sales is not strictly necessary for identification, but improves the precision of the estimates. It adds additional information, since it is estimated in a larger sample of firms, and captures information about the average price charged during a year (rather than the price in the last quarter of a year). Finally, the share of materials in flexible inputs α does not affect the firm-level response to shocks, and is identified from the constraint that the aggregate materials share in gross output in the model $s_m = \alpha(1-\nu)(\sigma-1)/\sigma$ fits its empirical Euro area counterpart of 0.52.

We denote the stacked vector of estimated local projection coefficients as $\hat{J}$, the estimated covariance matrix of the impulse response coefficients[12] $\hat{V}$ and $J(\delta, \sigma, \zeta, \nu)$ the counterpart of model-implied impulse responses. To account for possible capacity adjustment at longer horizons, we match responses for the first three years after a shock hits. Our results are similar when we restrict to the first two years or extend to five years (Table 10), suggesting that capacity adjustment does not substantially affect our targeted moments over this horizon.[13] The estimates minimize the objective function:

$$\min_{\delta,\sigma,\zeta,\nu} (\hat{J} - J(\delta, \sigma, \zeta, \nu))' \hat{V}^{-1} (\hat{J} - J(\delta, \sigma, \zeta, \nu)) \quad \text{s.t.} \quad s_m = \alpha(1-\nu)(\sigma-1)/\sigma \tag{27}$$

As a baseline, we calculate asymptotic standard errors for the estimated parameters based on the covariance matrix $Var(\hat{\delta}, \hat{\sigma}, \hat{\zeta}, \hat{\nu}) = (G'\hat{V}^{-1}G)^{-1}$, where G is the Jacobian matrix of the model-implied impulse response vector at the minimum, and for $\hat{\alpha}$ and other derived parameters using the Delta method. We also compute confidence intervals for selected derived parameters (most importantly κ^y) using a parametric bootstrap that samples from the asymptotic distribution of our target moments. Unlike the asymptotic confidence intervals, the bootstrap respects theoretical bounds on the parameters (i.e., $\nu, \alpha \in [0,1]$ and $\sigma > 1$) and captures asymmetry in the distribution of derived parameters arising from nonlinear interactions between the underlying structural parameters.

Model fit and parameter estimates. Figure 5 shows the fitted model impulse responses of prices, output, material purchases and sales as well as the local projection estimates. Our simple model does well at fitting the relative magnitudes and dynamics of all four responses. Prices in the model increase over three years and decline slightly afterward, compared to a decline from year $t+4$ in the data. The gradual increase of prices reflects both price-stickiness and the staggered realization of the shock over the four quarters of year zero. The response of output peaks in the period the shock hits and declines afterward, tracking the response estimated in the data. The response of material inputs also peaks on impact in the model, while it peaks in the period after the shock hits in the data, but overall the two responses are close for the first three years after a shock. Figure 5 also shows the (non-targeted) response of flexible employment in the model relative to the response of overall employment in the data. Employment in the data initially responds less than flexible employment in the model. This is consistent with an important

[12] We estimate this covariance matrix by estimating a system of "stacked" local projections, see, e.g., Wooldridge (2003) p166–173.

[13] More generally, since the capacity pressure channel is disciplined by the firm-level responses we match, we expect that a richer model of capacity adjustment—as in Woodford (2005)—that fits the same responses would yield a similar estimate of κ^y.

overhead component of employment that is fixed in the short run (and would thus be part of the fixed factor of our model).

Under the null hypothesis that the model is correct, the value of the objective function at the minimum follows a χ^2 distribution with 12 degrees of freedom (we fit 16 moments with four parameters). Given a test statistic of 10.9, we do not reject the null that the model is correct (p-value: 0.54). However, this test has limited power given the modest number of overidentifying restrictions.

Figure 5: Impulse responses to a firm-level demand shock in the data and the estimated model

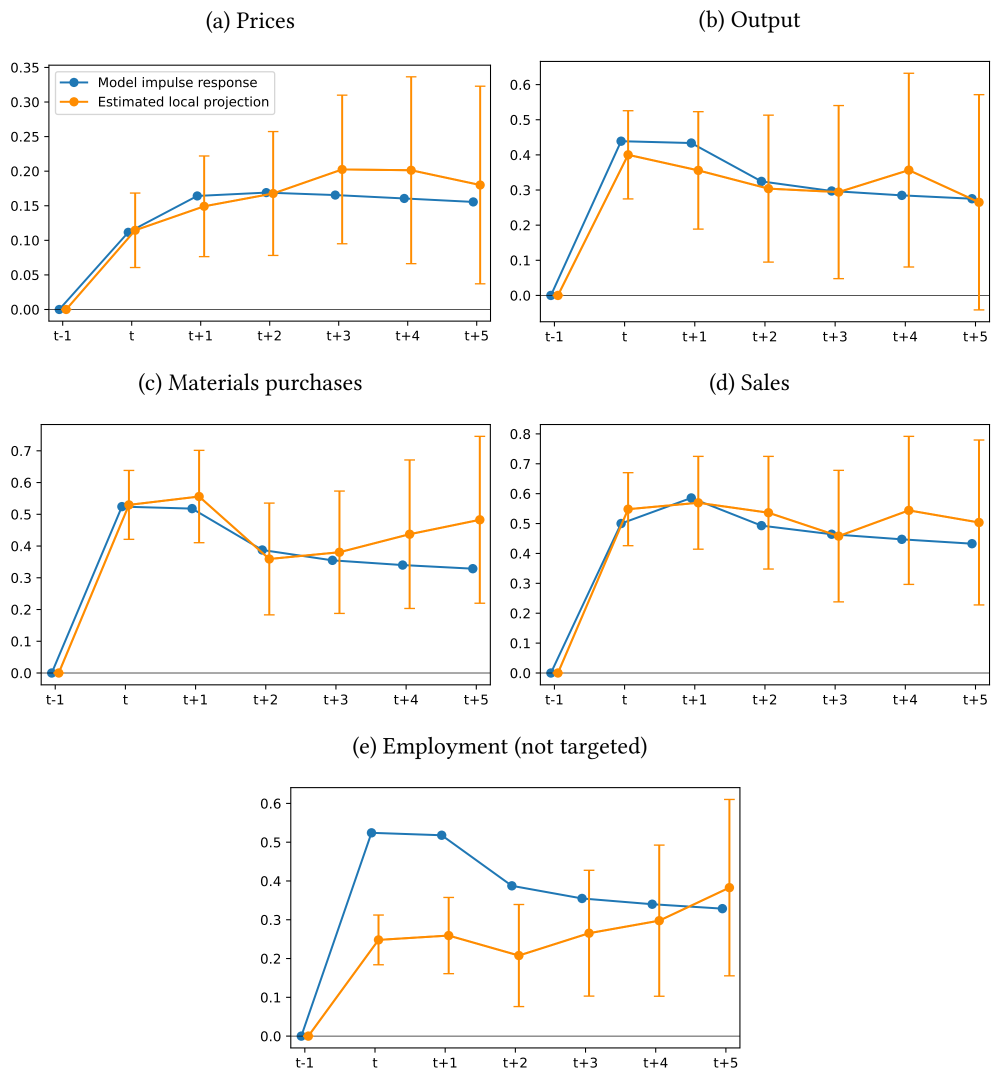

Notes: The figure plots the impulse responses to a firm-level demand shock in the estimated baseline model together with the corresponding local projections targeted in the estimation. The targeted local projections are our baseline estimates "with controls". The error bars show 95% confidence intervals based on standard errors clustered at the firm level.

Table 1 panel (I) shows the estimated model parameters. The estimated slope of the flexible-price supply curve δ is 0.57. The estimates imply a large role for fixed factors, since the empirical response of material inputs exceeds the response of output. The share parameter for the fixed factor ν is estimated to be 0.16. The estimates of δ and ν imply a CES substitution parameter between fixed and flexible factors ψ of -1.9. While the fixed factor share ν and the CES substitution parameter ψ are not precisely estimated, the overall slope of the flexible-price supply curve δ is significantly different from zero and we can reject that firms' supply curves are flat (which would be the case with $\nu = 0$ or $\psi = 1$). To match the empirical materials share in gross output with the estimated fixed factor share, the model requires materials to make up a large share of the flexible factor, and the estimate of α amounts to 0.79. In the data, the average share of firms' intermediate purchases in the sum of total payroll and intermediate purchases is about 0.69. Our results are thus consistent with about 60% of a firm's payroll going to flexible labor inputs and the remainder going to overhead labor. This is similar to results in Dhyne et al. (2025), who study the employment and wage responses of Belgian firms to similar trade demand shocks. The demand elasticity σ we estimate amounts to 4.5, consistent with a steady-state markup of about 28%. Finally, our estimate of the demand scaling parameter is 1.15.

The capacity pressure channel. We now discuss what these estimated values imply for the response of aggregate inflation to aggregate demand shocks. The slope of the Phillips curve in our model is given by equation (17):

$$\kappa = \underbrace{\frac{\delta}{1+\sigma\delta}\frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{(1-s_m)(1-\nu)}{1-\nu-s_m}}_{\text{Capacity pressure}\equiv\kappa^y} + \underbrace{\frac{\phi(1-\alpha)}{1+\sigma\delta}\frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{(1-s_m)(1-\nu)-\psi\nu s_m}{1-\nu-s_m}}_{\text{Wage pressure}\equiv\kappa^w}$$

Our empirical strategy pins down all parameters that determine the capacity pressure channel of the Phillips curve slope. Table 1 panel (III) shows the aggregate Phillips curve slopes implied by our estimation. Our estimate for the contribution of capacity pressure to the overall slope is 0.039 and statistically significantly different from zero. This magnitude is the product of the three components outlined in equation (17). The hypothetical flexible-price price response to a demand shock $\delta/(1+\sigma\delta)$ amounts to 0.16. This is attenuated by the Calvo factor $(1-\beta\theta)(1-\theta)/\theta$, which for our (standard) calibration choices amounts to 0.19. Finally, the second-round effects resulting from roundabout production $(1-s_m)(1-\nu)/(1-\nu-s_m)$ amplify the response by a factor of 1.27.

To determine the wage pressure component, we need to take a stance on the slope of the reduced-form wage Phillips curve ϕ, which governs the response of aggregate real wages to aggregate demand shocks. In a standard New Keynesian model, this parameter is determined by households' labor supply decisions. Our firm-level approach cannot identify ϕ—doing so would require variation in labor demand at the level of labor markets, not firms. To provide a sense of the importance of the capacity pressure channel relative to the overall slope of the Phillips curve, we deliberately choose an unrealistically high value of $\phi = 1$. This allows us to interpret the strength of the wage-pressure channel as an upper bound, and the relative importance of the capacity pressure channel as a lower bound in what follows.[14]

With our upper bound to the wage Phillips curve of $\phi = 1$, the overall slope of the Phillips curve including the wage pressure channel amounts to 0.059, and the capacity pressure channel accounts for two thirds of that slope. Why do the dynamics of real wages play a relatively small role for the Phillips curve in our estimated model? First, our

[14] The available evidence suggests real wages are not very cyclical conditional on monetary policy shocks. For example, the results of Christiano et al. (2005) suggest U.S. monetary policy shocks that increase output by 0.5% increase real wages by $\approx 0.1\%$, which implies $\phi \approx 0.2$. Cantore et al. (2021) find that real wages do not respond much to monetary policy shocks in several developed economies, i.e., $\phi \approx 0$.

Table 1: Model estimation results

(I) Estimated parameters	
Slope of flex-price supply δ	0.571
	(0.172)
Slope of demand σ	4.460
	(1.293)
Fixed input share ν	0.163
	(0.091)
Materials share in flexible input α	0.785
	(0.129)
Demand scaling parameter ζ	1.147
	(0.174)
Implied substitution parameter $\psi^{\dagger}$	-1.940
	(1.817)
Estimation fit	10.883
J-test p-value	0.539
(II) Calibrated parameters	
Discount factor β	0.990
Price stickiness θ	0.650
Firm-level demand persistence ρ	0.992
(III) Phillips curve slope	
Capacity pressure κ^y	0.039
	(0.014)
Bootstrapped 95% CI for κ^y	(0.020, 0.069)
Overall slope κ (w. $\phi = 1$)	0.059
Ratio κ^y/κ	0.66

Notes: The table shows the results of the impulse-response matching estimation described in Section 4. We target our baseline "with controls" local projection estimates. Panel (I) reports the parameter estimates and standard errors. Standard errors are calculated using the empirical covariance matrix together with the Jacobian of the model-implied impulse responses at the minimum. Panel (II) lists the values of calibrated parameters. Panel (III) shows the implied Phillips curve slopes implied by equation (17). The overall slope is calculated assuming a wage Phillips curve slope of $\phi = 1$. † The substitution parameter ψ is not estimated, but follows from δ and ν.

estimation suggests that the role for labor inputs in the flexible factor is limited due to our relatively high estimate of α and a large implied share of overhead labor. Consequently, wages play a smaller role for marginal cost than in a model where all labor is assumed to be flexible. Second, with an active capacity pressure channel ($\delta > 0$), GDP increases less and prices increase more after a positive aggregate demand shock. Since real wages move with GDP with elasticity ϕ, the wage pressure channel becomes less relevant overall.

Comparison to "standard" calibrations. We now compare firm-level and aggregate responses in the estimated model to two more "standard" calibrations of the New Keynesian model. First, we consider the textbook calibration in Gali (2008), augmented with firm-level demand shocks, but with a standard Cobb-Douglas production function and without material inputs. This corresponds to our model with $\psi = 0$ (Cobb-Douglas) and $\alpha = 0$ (no materials). The remaining parameters are calibrated to $\nu = 1/3$, $\sigma = 6$ and $\theta = 2/3$. The Galí model does not feature materials and is usually interpreted as a model in value-added terms, while our model is explicitly formulated in terms of gross output. To make firm-level responses between the two models comparable, we convert the responses of output and prices in our estimated model to their value-added counterparts. As we show in online Appendix C, value added and value-added prices in our model are given by:

$$\tilde{v}a_{i,t} = \frac{1}{1-s_m}\tilde{y}_{i,t} - \frac{s_m}{1-s_m}\tilde{m}_{i,t} \tag{28}$$

$$\tilde{p}^{va}_{i,t} = \frac{1}{1-s_m}\tilde{p}_{i,t} \tag{29}$$

s_m is the steady-state share of materials in gross output.

Figure 6 compares the firm-level responses of value added and the value-added price. In our estimated model the response of the value-added price is about three times as large, while the response of value added is slightly smaller, compared to the Galí calibration. The difference between the value-added price response is significant at all horizons based on the bootstrapped 95% confidence interval. The Phillips curve slopes between our estimated model and the Galí model are directly comparable since both are expressed in terms of GDP. In line with the smaller price response to demand shocks at the firm-level, the Galí model features a capacity pressure channel of 0.022, compared to our baseline estimate of 0.039. This falls just inside the 95% confidence interval for κ^y, with a two-sided p-value of 0.1.

Second, we compare our estimated model to a "standard" Cobb-Douglas calibration of our model with material inputs. We calibrate the share of materials in gross output to its empirical value of 0.52, and then choose the steady-state labor cost share $(1-\alpha)(1-\nu)$ to be twice as large as the fixed factor cost share ν as in the Galí calibration. The resulting parameter values are $\nu = 0.125$ and $\alpha = 0.713$. We set the remaining parameters $\sigma = 6$ and $\theta = 2/3$ as in the Galí calibration. Figure 7 shows that with this calibration, firms' prices respond substantially less, and output substantially more to demand shocks than in our estimated model. The differences are significant based on bootstrapped 95% confidence intervals. The aggregate capacity pressure channel for this calibration contributes 0.016 to the slope of the Phillips curve, which falls outside of the bootstrapped 95% confidence interval for our estimate of κ^y.

Robustness checks. We conduct several robustness checks in which we target different sets of empirical estimates. These are summarized in Appendix Table 10. Overall, we get similar parameter estimates and contributions of the capacity pressure channel to the Phillips curve for all targets we use. In column (2), we restrict the targeted moments to the first two local projection coefficients up to horizon $t+1$. This further alleviates concerns that the model with fixed factors might miss important capacity adjustments the longer the horizon. For completeness, we also show results for an extended set of targeted moments that includes all local projection coefficients up to $t+5$ in column (3). In both cases we get similar parameter estimates and capacity pressure contributions to the Phillips curve slope of 0.038 and 0.039, respectively. In column (4), we target our baseline estimates without controls, which also yields similar parameter estimates and a capacity pressure contribution of 0.035. In column (5), we use the output price index instead of domestic PPI prices as the empirical target for prices, which allows us to estimate the output, materials and price response based on the same larger sample of firms. This yields a capacity pressure contribution of 0.04. In column (6) we use estimates restricted to the sample of firms in the PPI for all outcomes, again to make

Figure 6: Firm-level responses in the estimated model vs. Gali (2008) calibration

(a) Value-added price (b) Value added

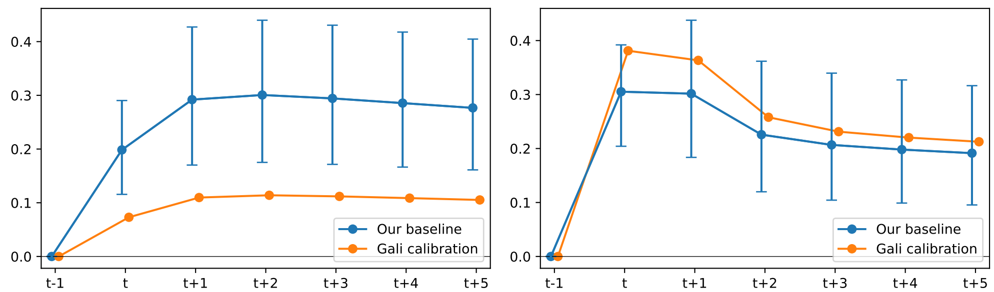

Figure 7: Firm-level responses in the estimated model vs. Cobb-Douglas calibration with materials

(a) Price (b) Output

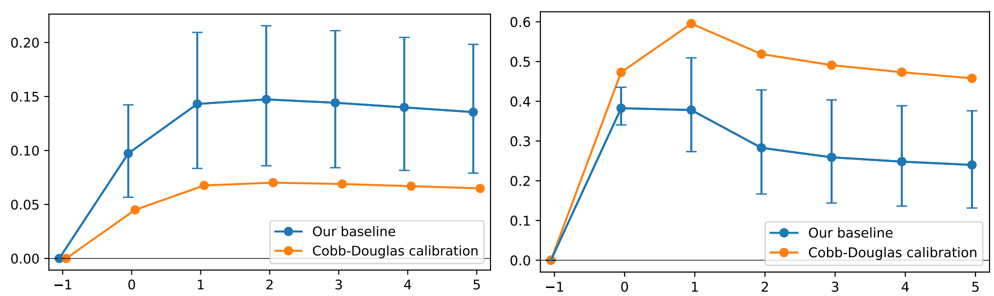

Notes: Figure 6 plots the impulse responses of value-added prices and value added to a firm-level demand shock in the estimated baseline model and in a model calibrated using parameter values from Gali (2008). Figure 7 plots the impulse responses of output prices and output to a firm-level demand shock in the estimated baseline model and in a model with material inputs and a "standard calibration" of a Cobb-Douglas production function. The error bars show bootstrapped 95% confidence intervals.

sure output, materials and price responses are estimated from the same (in this case smaller) sample. This yields a slightly smaller capacity pressure contribution of 0.032. All coefficient estimates we use as targets are summarized in Table 11.

In Appendix Figures 12 to 14 we show the sensitivity of our baseline results to alternative values of the calibrated parameters β, ρ and θ. We vary parameter values one-by-one from our baseline and re-estimate the model for each parameter combination. Varying the discount factor β has a negligible impact on the slope of the Phillips curve or the fit of the model. Varying ρ also yields similar Phillips curve slopes, but lower values of ρ decrease the fit of the firm-level impulse responses. Finally, higher values of θ decrease the slope of the Phillips curve we estimate. The fit of the firm-level impulse responses improves slightly for small decreases in θ. Since we calibrate θ directly to the frequency of price changes we observe in the PPI microdata, we are confident in our baseline calibration.

Finally, in Table 12 column (2), we show estimation results for the model with a Cobb-Douglas production function (i.e. with the restriction $\psi = 0$). Compared to our baseline, this produces a larger capacity pressure channel with magnitude 0.058. However, the model fit is substantially worse and some parameter estimates appear implausible. To fit the price response to firm-level shocks with a unit elasticity of substitution between fixed and flexible factors, the model estimates a fixed factor share of 0.31, about double our baseline value. Consequently, the model requires an implausibly high share of materials in the flexible factor to match the materials share in gross output ($\alpha = 0.99$). This would imply that almost all labor is fixed in the short run and is inconsistent with the observed employment response to demand shocks. For this reason, we prefer our baseline with CES production, even if the parameter ψ is not precisely estimated.

4.1 Model extensions

To further demonstrate the robustness of our results, we extend the basic model of Section 2 in two dimensions. First, we discuss a version of the model with Kimball demand that features variable desired markups as in Gagliardone et al. (2025). Second, we discuss a version of the model with monopsonistic labor markets to match the response of wages to firm-level demand shocks estimated in Section 3.3. We provide a brief discussion of both extensions and results below, and delegate a more detailed description of the extended model to the online Appendix C.

Kimball demand. In the model with Kimball demand, gross final output is produced using a Kimball aggregator defined by $\int_i Z_{i,t} \Upsilon \left(\frac{Y_{i,t}}{Z_{i,t} Y_t} \right) di = 1$. Intermediate producers now face a demand curve with variable elasticity. This affects the desired pass-through of marginal cost into prices, and leads to markups that vary with the state of firm-level demand. The rate of desired cost pass-through is determined by the relative-price elasticity of markups Γ, which is constant around the steady state and depends on the parameters of the Kimball aggregator Υ. The slope of the flexible-price supply curve δ in this setting is given by:

$$\delta = \frac{(1-\psi)\nu}{(1-\nu)(1+\Gamma)} \tag{30}$$

$\Gamma = 0$ corresponds to the baseline model with CES demand. We calibrate $\Gamma = 1.33$ following Gagliardone et al. (2025) and Amiti et al. (2019). The response of firms' prices, output, sales and materials purchases to firm-level demand shocks is still given by the same equations (23) to (26) as in our baseline model with CES demand—but with an underlying definition of δ corresponding to equation (30). We thus estimate the same parameter values independent of the value of Γ (see Table 12, column (3)) and for any calibration of Γ, the estimated model features the exact same response of prices, output, sales and material inputs.

The calibration of Γ determines the changes in markups and marginal cost underlying the response of prices. Figure 9 illustrates the response of the three variables to a firm-level shock in the baseline model with CES demand and the extension with Kimball demand. In the CES case, markups fall slightly on impact due to the sluggish adjustment of prices, but return to their original level after three years. With Kimball demand, desired pass-through of marginal cost is incomplete. To match the price response in the data, the model thus generates a stronger increase in marginal cost (through a lower value of ψ) and a larger, persistent drop in markups (reflecting incomplete pass-through). Markups remain below and marginal cost above their initial level as long as output and prices are elevated.

Figure 9: Response of prices, marginal cost and markups to firm-level demand shocks with Kimball and CES demand

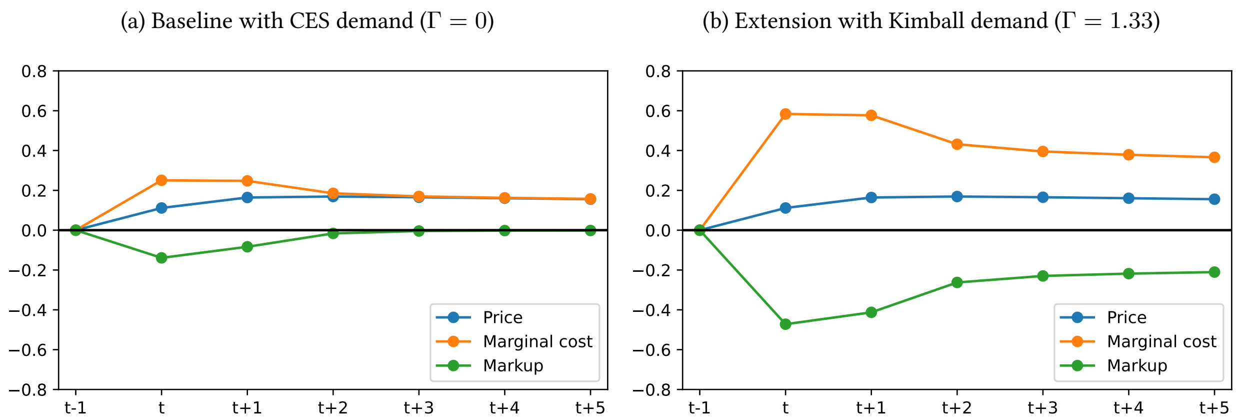

Notes: The Figure plots the evolution of log prices, log marginal cost and log markups in our estimated baseline model, compared with the estimated model with Kimball demand ($\Gamma = 1.33$).

In the aggregate, the Phillips curve of the model with Kimball demand is given by:

$$\kappa = \underbrace{\frac{\delta}{1+\sigma\delta}\frac{(1-\beta\theta)(1-\theta)}{\theta}\cdot\frac{(1-s_m)(1-\nu)}{1-\nu-s_m}}_{=\text{Capacity pressure } \kappa^y} + \underbrace{\frac{\phi(1-\alpha)}{1+\sigma\delta}\frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{(1-s_m)(1-\nu)-\psi\nu s_m}{1-\nu-s_m}\frac{1}{1+\Gamma}}_{\text{Wage pressure}\equiv\kappa^w} \tag{31}$$

The expression for the capacity pressure channel is identical to the baseline model with CES demand, since incomplete pass-through of the marginal cost increase associated with higher output is already reflected in the slope of the flexible-price supply curve δ. This is not the case for the marginal cost increase associated with higher wages, and the wage-pressure channel is attenuated by the pass-through rate $1/(1+\Gamma)$. Consequently, the capacity pressure channel in the model with Kimball demand has the same magnitude as in our baseline, while the wage pressure channel is smaller (see Table 12, column (3)). For our baseline calibration of $\Gamma = 1.33$, the share of the capacity pressure channel in the overall Phillips curve slope thus increases to 74%. In Appendix Figure 15 we show the results of estimating the model for a wide range of calibrated values for Γ—the magnitude of the capacity pressure channel is constant, but its share of the overall Phillips curve increases as we increase Γ.

Monopsony. In Section 3 we show that firms' wages increase slightly in response to firm-level demand shocks. This is inconsistent with our baseline assumption of competitive labor markets and suggests firms have some monopsony power. We extend the model to include this feature and explore its quantitative implications. We assume a CES labor supply curve $L_{i,t} = L_t \left(W_{i,t}/W_t\right)^{1/\eta}$. The degree of labor market power is determined by firms' labor supply elasticity $1/\eta$ and the case of competitive labor markets corresponds to $\eta \to 0$. In this setting, the slope of firms' flexible-price supply curve is given by:

$$\delta = \frac{(1-\psi)\nu}{1-\nu} + \frac{\eta(1-\alpha)}{(1-\nu)(1+\alpha\eta)} \tag{32}$$

The first term reflects capacity pressure, while the second reflects the marginal cost impact of the wage increase necessary to attract additional workers. After aggregate demand shocks, all firms increase wages and only the first term is relevant for the capacity pressure channel. However, matching the response to firm-level shocks identifies the overall δ. To address this issue, we estimate the model with monopsony and identify η by matching the impulse response of wages to the estimated wage response. Figure 10 shows that the extended model matches the response of firm-level wages to demand shocks. The corresponding parameter estimates are reported in Table 12, column (4). We estimate η to be 0.1, which means that to increase employment by 10% after a firm-level demand shock a firm needs to raise wages by 1%. This is lower than estimates in Dhyne et al. (2025), because we estimate a smaller reduced-form response of wages to demand shocks. The values of the other estimated parameters remain similar to the baseline case.

Figure 10: Response of wages to firm-level shock

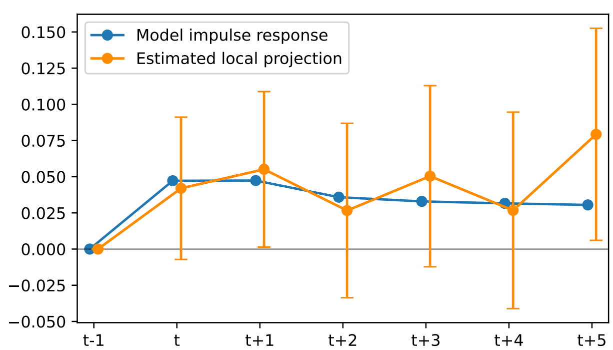

Notes: The figure plots the impulse response of wages to a firm-level demand shock in the estimated model with monopsony together with the corresponding local projection targeted in the estimation. The targeted local projections are our baseline estimates "with controls". The error bars show 95% confidence intervals based on standard errors clustered at the firm level.

We can then calculate the correct contribution of the capacity pressure channel to the Phillips curve slope. In the model with monopsony, the slope of the Phillips curve is given by:

$$\kappa = \underbrace{\frac{\delta}{1+\sigma\delta}\frac{\chi^a}{\chi^i}\frac{(1-\beta\theta)(1-\theta)}{\theta}\cdot\frac{(1-s_m)(1-\nu)}{1-\nu-s_m}}_{=\text{Capacity pressure } \kappa^y} + \underbrace{\frac{\phi(1-\alpha)}{1+\sigma\delta}\frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{(1-s_m)(1-\nu)-\psi\nu s_m}{1-\nu-s_m}}_{\text{Wage pressure}\equiv\kappa^w} \tag{33}$$

In this expression, χ^a/χ^i corrects for the monopsony term captured in δ. The term depends on η and other model parameters (see online Appendix C). As shown in Table 12 column (4), the impact on the magnitude of the capacity pressure channel is marginal. The reason for this small quantitative effect is that the labor supply elasticity remains relatively close to the competitive limit, and that the impact of wages on marginal cost is limited due to the importance of materials in our baseline model.

Full model with both extensions. Table 12 also shows estimation results for the full model of online Appendix C with both Kimball demand and monopsony in labor markets. As suggested by the results for both extensions separately, the magnitude of the implied capacity pressure channel remains almost identical to our baseline model, while its share in the overall slope of the Phillips curve increases from two thirds in the baseline to 74%.

5 The response of prices to monetary policy shocks

In the previous section we have estimated the magnitude of the capacity pressure channel of the Phillips curve slope. This allows us to quantify how monetary policy or other aggregate demand shocks affect aggregate prices through that channel. The slope of the Phillips curve determines the response of inflation to an aggregate demand shock at fixed inflation expectations. We are interested in the overall response including the effect coming from adjustments in expectations. To derive an expression for this overall effect, we need to complement the Phillips curve with a dynamic IS equation:

$$\pi_t = \kappa y_t^{gdp} + \beta E_t \pi_{t+1} \tag{34}$$

$$y_t^{gdp} = \lambda_y E_t y_{t+1}^{gdp} + \lambda_\pi E_t \pi_{t+1} + u_t. \tag{35}$$

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

$$\pi_{t+k} = \frac{\kappa}{1-\beta\rho_u} \Lambda u_{t+k} \tag{36}$$

$$y_{t+k}^{gdp} = \Lambda u_{t+k} \tag{37}$$

Λ depends on the coefficients of the dynamic IS equation, but cancels out in the relative response. The response of period $t+k$ inflation to an aggregate demand shock that increases log output by one in period t is thus given by $\pi_{t+k} = \kappa/(1-\beta\rho_u)\rho_u^k$, and the effect on the price level is given by the cumulative sum of inflation rates:

$$\frac{\partial p_{t+k}}{\partial u_t} \frac{\partial u_t}{\partial y_t^{gdp}} = \frac{\kappa}{1-\beta\rho_u} \frac{1-\rho_u^{k+1}}{1-\rho_u} \tag{38}$$

We discuss the response to a monetary policy shock with persistence $\rho_u = 0.69$ and compare our results to alternative cross-sectional estimates of the Phillips curve slope in Figure 11. Hazell et al. (2022) estimate a Phillips curve slope of $\kappa = 0.008$ using U.S. data on prices of non-tradeables.[15] Gagliardone et al. (2025) use two specifications to estimate a Phillips curve slope based on micro data on Belgian manufacturing firms. The first is based directly on the response of prices to the demand shock they construct, and is more directly comparable to our approach. This yields an estimate of $\kappa = 0.021$. The second is based on the response of their measure of marginal cost, and yields an estimate of $\kappa = 0.006$.

Given these slope estimates, a monetary policy shock that increases output by 1% in period t increases the aggregate price level after three years by 0.08% based on the slope estimate of Hazell et al. and 0.06% and 0.21% based on the slope estimates of Gagliardone et al. Our estimates imply that the price level increases by 0.39% through capacity pressure alone. Even though this is still a lower bound, our estimate for the cumulative inflation cost is about double the higher estimate of Gagliardone et al. and about five times that of Hazell et al. and the lower estimate of Gagliardone et al. Based on bootstrapped 95% confidence intervals for κ^y, we can reject that the inflation cost implied by our estimate is equal to that implied by the point estimates of Hazell et al. (2022) and the lower of the two Gagliardone et al. point estimates. The effect implied by the higher point estimates of Gagliardone et al. (2025) falls just within our 95% confidence interval, but could be rejected with a two-sided p-value of about 0.08.

Figure 11: Response of the price level to a monetary policy shock

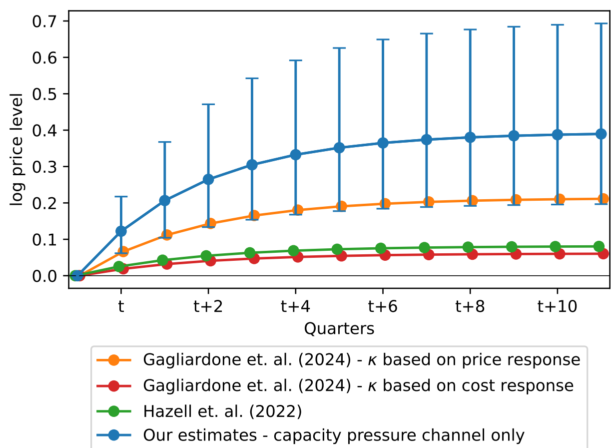

Notes: The figure compares the implied cumulative inflation response to a monetary policy shock that increases output by 1% in period 0. The results of Hazell et al. (2022) are converted to an output Phillips curve using the covariance between the U.S. unemployment and output gaps.

6 Conclusion

In this paper, we identify the "capacity pressure" channel of the Phillips curve slope in a New Keynesian model from the response of prices, output and material inputs to firm-level demand shocks. Our results suggest that the capacity pressure channel contributes about 0.039 to the slope of the Phillips curve. Even though this represents a lower bound, it already implies a substantially stronger response of prices to monetary policy or other aggregate demand shocks than previous cross-sectional estimates of the total Phillips curve slope. Our results also suggest that this

[15] The estimates in Hazell et al. (2022) are for the unemployment gap Phillips curve, which we convert to an output-based slope using the covariance between the U.S. unemployment gap and output gap.

lower bound is useful—even with an upper-bound assumption on the elasticity of real wages to aggregate demand shocks, the capacity pressure channel makes up about two thirds of the Phillips curve slope.

We see several promising directions for further research building on our approach. We estimate average supply and demand parameters that are the same for all firms in the economy and constant over time. There is potential for substantial heterogeneity in these parameters, but more data—perhaps for a country with a larger firm population than Denmark—would be needed to identify heterogeneity with useful precision. We think that identifying such heterogeneity could provide both improved estimates of the response of prices to aggregate demand shocks, and possibly allow for more targeted policy that minimizes the inflationary impact of demand stimulus.

References

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

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

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

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

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

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, April 2022, *112* (4), 1147–1193.

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

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

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

_ , _ , and _ , "A Practical Guide to Shift-Share Instruments," *Journal of Economic Perspectives*, February 2025, *39* (1), 181–204.

Cantore, Cristiano, Filippo Ferroni, and Miguel León-Ledesma, "The Missing Link: Monetary Policy and The Labor Share," *Journal of the European Economic Association*, June 2021, *19* (3), 1592–1620.

Cerrato, Andrea and Giulia Gitti, "The Return of the Phillips Curve: Evidence from US Cities," *Mimeo*, June 2025.

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

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot, "Firm Responses and Wage Effects of Foreign Demand Shocks with Fixed Labor Costs and Monopsony," *American Economic Review*, December 2025, *115* (12), 4328–4368.

Fitzgerald, Terry, Callum Jones, Mariano Kulish, and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation?," *American Economic Journal: Macroeconomics*, October 2024, *16* (4), 114–142.

Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens, "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," *American Economic Review*, November 2025, *115* (11), 3941–3974.

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

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.

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.

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

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

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

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.

Woodford, Michael, "Firm-Specific Capital and the New Keynesian Phillips Curve," *International Journal of Central Banking*, 2005, *1* (2).

Wooldridge, J M, *Econometric Analysis of Cross Section and Panel Data*, 2 ed., Cambridge, Massachusetts: The MIT Press, 2003.

Appendix

A Supplementary Figures and Tables

Table 2: Descriptive statistics of firms in the estimation sample

	(I) Output sample				(II) Price sample			
	Mean	Median	10th pctl.	90th pctl.	Mean	Median	10th pctl.	90th pctl.
Sales (Mio EUR)	51.76	11.80	3.47	74.31	100.51	27.31	6.89	141.18
Employment (FTE)	136.54	50.00	19.00	242.97	236.85	98.00	30.00	426.00
Assets (Mio EUR)	51.93	7.75	2.11	59.06	105.77	18.94	4.32	127.75
Goods exports (Mio EUR)	21.96	4.45	0.25	35.75	42.24	12.34	1.55	73.52
Exports (Mio EUR)	30.47	5.13	0.46	43.09	59.83	13.80	1.76	88.41
Imports (Mio EUR)	14.15	1.92	0.10	21.36	28.07	5.95	0.51	42.03
Export destinations	18.58	15.00	2.00	41.00	25.02	22.00	5.00	50.00
Exported products (4-digit HS)	13.46	8.00	2.00	30.00	17.69	11.00	3.00	39.00
Share of payroll in payroll+purchases	0.31	0.31	0.15	0.46	0.27	0.26	0.13	0.42
Prices reported in the PPI					5.69	4.00	2.00	11.00
Share in MFG employment	0.60				0.45			
Firms	2,444				824			
N	28,748				12,297			

Notes: The table reports summary statistics for the pooled 1999–2021 period. Nominal amounts are deflated to 2020 prices using the Danish CPI and then converted to EUR using the fixed DKK/EUR exchange rate. Panel (I) "Output sample" covers all manufacturing firms that fulfill the sampling requirements outlined in Section 3.1. Panel (II) "Price sample" describes all manufacturing firms that fulfill the sampling requirements outlined in Section 3.1 and can also be matched to the PPI micro data. Note that due to missing lagged controls and singletons within fixed effect cells, the number of distinct firms and observations is larger than the sample for estimated regressions.

Table 3: Reduced-form effect on output — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Firm FE	(4) 4d-nace FE	(5) No sector FE	(6) Winsor 2/98
t	0.40***	0.41***	0.40***	0.29***	0.59***	0.60***
	(0.064)	(0.057)	(0.096)	(0.11)	(0.079)	(0.067)
t+1	0.36***	0.39***	0.25**	0.19	0.54***	0.54***
	(0.085)	(0.078)	(0.11)	(0.14)	(0.10)	(0.089)
t+3	0.29**	0.26**	0.34**	0.28	0.48***	0.38***
	(0.13)	(0.11)	(0.17)	(0.21)	(0.14)	(0.12)
t+5	0.27*	0.27*	0.23	0.31	0.46***	0.41***
	(0.16)	(0.14)	(0.15)	(0.24)	(0.16)	(0.15)
Firms	2,016	2,316	1,780	1,965	2,016	2,016
N	75,162	158,115	74,039	71,218	75,262	75,262
F	39.051	51.306	17.565	6.774	55.990	79.665

Notes: The table reports local projections of log output on the demand shock as described by equation (19). Standard errors in parentheses are clustered at the firm level. $*$ p$<$0.1 $**$ p$<$0.05 $***$ p$<$0.01. The sample is the "Output sample", but the number of observations varies with the availability of controls and fixed effects. The number of distinct firms and observations refer to the union of samples across horizons. The F-statistic tests the null hypothesis that all coefficients jointly equal zero. The included specifications vary included controls and fixed effects. (1) Baseline specification with controls. (2) Baseline specification without controls. (3) Adds a firm FE that controls for trends. (4) Replaces 2d-sector $\times$ time FE with 4d-sector $\times$ time FE. (5) Replaces 2d-sector $\times$ time FE with time FE. (6) Winsorizes outcomes at the 2nd/98th percentile.

Table 4: Reduced-form effect on output — different sample restrictions

	(1) Baseline	(2) Only PPI	(3) More balanced	(4) Higher exports	(5) Pre 2020	(6) Exclude Germany	(7) Exclude Sweden
t	0.40***	0.41***	0.43***	0.41***	0.39***	0.38***	0.44***
	(0.064)	(0.083)	(0.096)	(0.095)	(0.095)	(0.088)	(0.094)
t+1	0.36***	0.46***	0.36***	0.39***	0.39***	0.30***	0.35***
	(0.085)	(0.11)	(0.12)	(0.12)	(0.12)	(0.11)	(0.11)
t+3	0.29**	0.34**	0.41**	0.39**	0.41**	0.41**	0.41**
	(0.13)	(0.15)	(0.18)	(0.18)	(0.17)	(0.16)	(0.16)
t+5	0.27*	0.22	0.24	0.33*	0.30	0.34*	0.31*
	(0.16)	(0.19)	(0.20)	(0.20)	(0.19)	(0.18)	(0.19)
Firms	2,016	756	1,488	1,525	1,972	2,010	2,011
N	75,162	38,526	67,873	54,627	72,426	74,930	75,030
F	39.051	24.898	20.491	18.809	16.885	18.454	21.780

Notes: The table reports local projections of log output on the demand shock as described by equation (19). Standard errors in parentheses are clustered at the firm level. $*$ p$<$0.1 $**$ p$<$0.05 $***$ p$<$0.01. The base sample is the "Output sample", but the number of observations varies with the additional imposed restrictions. The number of distinct firms and observations refer to the union of samples across horizons. The F-statistic tests the null hypothesis that all coefficients jointly equal zero. The table reports results for the baseline specification with controls across different sample restrictions. (1) Baseline sample. (2) Restricts to firms that appear in the PPI sample. (3) Adds restriction of $\geq$10 uninterrupted years of firm activity. (4) Adds restriction of export share $>$0.25 in the year before a given shock. (5) Restricts sample to the pre-2020 period. (6) Excludes exports to Germany and (7) excludes exports to Sweden when constructing demand shocks.

Table 5: Reduced-form effect on prices — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Firm FE	(4) 4d-nace FE	(5) No sector FE	(6) Winsor 2/98
t	0.11^{***}	0.14^{***}	0.12^{***}	0.10^{***}	0.14^{***}	0.14^{***}
	(0.027)	(0.024)	(0.030)	(0.037)	(0.030)	(0.037)
t+1	0.15^{***}	0.16^{***}	0.14^{***}	0.10^{*}	0.13^{***}	0.16^{***}
	(0.037)	(0.035)	(0.039)	(0.058)	(0.034)	(0.044)
t+3	0.20^{***}	0.22^{***}	0.15^{***}	0.18^{**}	0.23^{***}	0.23^{***}
	(0.055)	(0.051)	(0.053)	(0.082)	(0.051)	(0.064)
t+5	0.18^{**}	0.23^{***}	0.16^{***}	0.069	0.13^{**}	0.19^{**}
	(0.073)	(0.069)	(0.061)	(0.11)	(0.056)	(0.082)
Firms	647	771	636	641	647	647
N	92,465	193,100	92,409	91,556	92,477	92,465
F	17.379	35.610	16.208	7.751	21.466	14.247

Notes: The table reports local projections of log prices on the demand shock as described by equation (20). Standard errors in parentheses are clustered at the firm level. $*$ $p<0.1$ $**$ $p<0.05$ $***$ $p<0.01$. The sample is the "Price sample", but the number of observations varies with the availability of controls and fixed effects. The number of distinct firms and observations refer to the union of samples across horizons. The F-statistic tests the null hypothesis that all coefficients jointly equal zero. The included specifications vary included controls and fixed effects. (1) Baseline specification with controls. (2) Baseline specification without controls. (3) Adds a firm FE that controls for trends. (4) Replaces 2d-sector $\times$ time FE with 4d-sector $\times$ time FE. (5) Replaces 2d-sector $\times$ time FE with time FE. (6) Winsorizes outcomes at the 2nd/98th percentile.

Table 6: Reduced-form effect on prices — sample restrictions

	(1) Baseline	(2) More balanced	(3) Higher exports	(4) Pre 2020	(5) Exclude Germany	(6) Exclude Sweden
t	0.11^{***}	0.11^{***}	0.12^{***}	0.11^{***}	0.086^{***}	0.12^{***}
	(0.027)	(0.028)	(0.029)	(0.028)	(0.026)	(0.027)
t+1	0.15^{***}	0.13^{***}	0.15^{***}	0.16^{***}	0.11^{***}	0.15^{***}
	(0.037)	(0.036)	(0.039)	(0.039)	(0.035)	(0.037)
t+3	0.20^{***}	0.19^{***}	0.19^{***}	0.20^{***}	0.17^{***}	0.20^{***}
	(0.055)	(0.055)	(0.057)	(0.055)	(0.053)	(0.055)
t+5	0.18^{**}	0.18^{**}	0.23^{***}	0.18^{**}	0.16^{**}	0.19^{**}
	(0.073)	(0.075)	(0.076)	(0.073)	(0.071)	(0.074)
Firms	647	601	535	628	647	647
N	92,465	90,956	68,548	88,487	92,243	92,338
F	17.379	14.871	16.897	15.173	11.065	20.218

Notes: The table reports local projections of log prices on the demand shock as described by equation (20). Standard errors in parentheses are clustered at the firm level. $*$ $p<0.1$ $**$ $p<0.05$ $***$ $p<0.01$. The sample is the "Price sample", but the number of observations varies with the additional imposed restrictions. The number of distinct firms and observations refer to the union of samples across horizons. The F-statistic tests the null hypothesis that all coefficients jointly equal zero. The table reports results for the baseline specification with controls across different sample restrictions. (1) Baseline sample. (2) Adds restriction of ≥ 10 uninterrupted years of firm activity. (3) Adds restriction of export share >0.25 in the year before a given shock. (4) Restricts sample to the pre-2020 period. (5) Excludes exports to Germany and (6) excludes exports to Sweden when constructing demand shocks.

Table 7: IV effect on prices — specifications and fixed effects

	(1) Baseline	(2) No controls	(3) Firm FE	(4) 4d-nace FE	(5) No sector FE	(6) Winsor 2/98
t	0.32***	0.39***	0.32***	0.37*	0.29***	0.37***
	(0.095)	(0.086)	(0.11)	(0.19)	(0.072)	(0.12)
t+1	0.34***	0.38***	0.29**	0.12	0.21***	0.34**
	(0.12)	(0.11)	(0.13)	(0.27)	(0.075)	(0.14)
t+3	0.51***	0.53***	0.36**	0.80	0.38***	0.57***
	(0.19)	(0.15)	(0.18)	(0.49)	(0.11)	(0.21)
t+5	0.61**	0.59***	0.48**	0.84	0.25**	0.68**
	(0.24)	(0.20)	(0.22)	(0.62)	(0.12)	(0.27)
2nd stage observations	18,796	25,943	18,786	18,661	18,799	18,796
2nd stage Firms	587	707	577	582	588	587
1st stage observations	16,652	22,133	16,423	15,894	16,662	16,652
1st stage Firms	2,032	2,317	1,803	1,982	2,032	2,032
1st stage F statistic	38.947	51.681	32.577	10.002	94.204	38.947

Table 8: IV effect on prices — sample restrictions

	(1) Baseline	(2) Only PPI	(3) More balanced	(4) Higher exports	(5) Pre 2020	(6) Exclude Germany	(7) Exclude Sweden
t	0.32***	0.29***	0.27***	0.32***	0.30***	0.29***	0.38***
	(0.095)	(0.096)	(0.087)	(0.099)	(0.099)	(0.10)	(0.11)
t+1	0.34***	0.31***	0.27**	0.34***	0.37***	0.31**	0.41***
	(0.12)	(0.12)	(0.11)	(0.13)	(0.14)	(0.13)	(0.14)
t+3	0.51***	0.47**	0.49***	0.45**	0.55***	0.49**	0.60***
	(0.19)	(0.18)	(0.18)	(0.19)	(0.20)	(0.20)	(0.22)
t+5	0.61**	0.57**	0.57**	0.72***	0.66**	0.65**	0.70***
	(0.24)	(0.23)	(0.24)	(0.26)	(0.26)	(0.27)	(0.27)
2nd stage observations	18,796	18,796	18,392	13,991	16,444	18,737	18,766
2nd stage Firms	587	587	545	489	574	586	587
1st stage observations	16,652	8,115	14,531	12,013	14,795	14,750	14,768
1st stage Firms	2,032	764	1,500	1,540	1,990	1,984	1,985
1st stage F statistic	38.947	23.225	37.707	34.830	31.389	30.366	30.527

Notes: Tables 7 and 8 report IV local projections of the log price on log output growth as described by equation (22). Output growth is instrumented with the demand shock. Standard errors in parentheses are clustered at the firm level. $*$ p$<$0.1 $**$ p$<$0.05 $***$ p$<$0.01. We use the two-sample TSLS estimator of Inoue and Solon (2010): the first-stage regression is estimated in the output sample and the structural equation in the price sample. The number of distinct firms and observations in the structural equation refers to the union of samples across horizons. We report the cluster-robust first-stage F statistic. Table 7 reports results across specifications that vary the baseline controls and fixed effects structure: (1) Baseline specification with controls and (2) without controls. (3) Adds a firm FE that controls for trends. (4) Replaces 2d-sector $\times$ time FE with 4d-sector $\times$ time FE. (5) Replaces 2d-sector $\times$ time FE with time FE. (6) Winsorizes outcomes at the 2nd/98th percentile. Table 8 reports results for the baseline specification with controls across different sample restrictions: (1) Baseline sample. (2) Restricts first-stage sample to firms in the PPI. (3) Adds restriction of $\geq$10 uninterrupted years of firm activity. (4) Adds restriction of export share $>$0.25 in the year before a given shock. (5) Restricts sample to the pre-2020 period. (6) Excludes exports to Germany and (7) excludes exports to Sweden when constructing demand shocks.

Table 9: Reduced-form effect on other firm-level outcomes

	(1) Capacity Utilization	(2) Output price index	(3) Materials	(4) Average Wage	(5) Sales	(6) Employment
t	0.14**	0.12***	0.53***	0.042*	0.55***	0.25***
	(0.058)	(0.041)	(0.055)	(0.025)	(0.062)	(0.033)
t+1	0.091	0.18***	0.56***	0.055**	0.57***	0.26***
	(0.075)	(0.057)	(0.074)	(0.027)	(0.079)	(0.050)
t+3	0.034	0.14*	0.38***	0.050	0.46***	0.27***
	(0.093)	(0.080)	(0.098)	(0.032)	(0.11)	(0.083)
t+5	0.067	0.18*	0.48***	0.079**	0.50***	0.38***
	(0.12)	(0.11)	(0.13)	(0.037)	(0.14)	(0.12)
Firms	629	2,026	2,109	2,110	2,016	2,110
N	17,895	76,488	99,991	99,908	75,162	100,023
F	6.164	8.400	91.470	2.800	77.363	57.488

Notes: The table reports local projections of several outcomes on the demand shock as described by equation (19). Standard errors in parentheses are clustered at the firm level. $*$ p$<$0.1 $**$ p$<$0.05 $***$ p$<$0.01. The sample is the "Output sample", but the number of observations varies with the availability of controls and fixed effects. The number of distinct firms and observations refer to the union of samples across horizons. The F-statistic tests the null hypothesis that all coefficients jointly equal zero. All specifications differ only in the outcome variable and use the same controls and fixed effects structure as the baseline specification with controls. (1) Reported capacity utilization. (2) Firm-level price index based on unit values. (3) Log intermediate purchases. (4) Log average wage calculated as total salaries divided by the number of full-time equivalent employees. (5) Log sales. (6) Log full-time equivalent employees.

Table 10: Alternative targets for structural estimation

	(1) Baseline	(2) Estimates up to t+1	(3) Estimates up to t+5	(4) Estimates w/o controls	(5) Output price index	(6) Estimates in PPI sample
(I) Estimated parameters						
Slope of flex-price supply δ	0.571	0.544	0.580	0.577	0.554	0.510
	(0.172)	(0.174)	(0.174)	(0.155)	(0.245)	(0.168)
Slope of demand σ	4.460	4.375	4.457	4.312	4.562	5.134
	(1.293)	(1.413)	(1.282)	(1.067)	(1.810)	(1.616)
Fixed input share ν	0.163	0.156	0.169	0.087	0.185	0.131
	(0.091)	(0.093)	(0.090)	(0.085)	(0.094)	(0.105)
Materials share in flexible input α	0.785	0.784	0.791	0.727	0.802	0.729
	(0.129)	(0.138)	(0.130)	(0.101)	(0.160)	(0.119)
Demand scaling parameter ζ	1.147	1.127	1.135	1.196	1.122	1.306
	(0.174)	(0.185)	(0.172)	(0.153)	(0.172)	(0.216)
Implied substitution parameter $\psi^{\dagger}$	-1.940	-1.936	-1.851	-5.050	-1.438	-2.393
	(1.817)	(1.933)	(1.701)	(6.125)	(1.360)	(2.976)
Estimation fit	10.883	5.386	14.809	13.556	7.092	7.155
J-test p-value	0.539	0.250	0.787	0.330	0.527	0.847
(II) Calibrated parameters						
Discount factor β	0.990	0.990	0.990	0.990	0.990	0.990
Price stickiness θ	0.650	0.650	0.650	0.650	0.650	0.650
Firm-level demand persistence ρ	0.992	0.992	0.992	0.992	0.992	0.992
(III) Phillips curve slope						
Capacity pressure κ^y	0.039	0.038	0.039	0.035	0.040	0.032
	(0.014)	(0.015)	(0.014)	(0.010)	(0.020)	(0.012)
Bootstrapped 95% CI for κ^y	(0.020, 0.069)	(0.019, 0.067)	(0.021, 0.070)	(0.021, 0.060)	(0.012, 0.071)	(0.017, 0.067)
Overall slope κ (w. $\phi = 1$)	0.059	0.059	0.059	0.060	0.058	0.056
Ratio κ^y/κ	0.66	0.65	0.67	0.59	0.68	0.58

Notes: The table shows model estimation results for different empirical targets used for impulse response matching. The estimation is described in Section 4. Column (1) restates our baseline estimates that target our baseline specification "with controls". Column (2) targets estimates from our baseline specification up to horizon $t + 1$. Column (3) targets estimates from our baseline specification up to horizon $t + 5$. Column (4) targets our baseline specification "without controls". Column (5) uses the estimates for the output price index instead of domestic PPI prices as target for the price outcome (using the "with controls" specification). Column (6) uses estimates that are only based on the sample of firms that can be matched to PPI data targets (using the "with controls" specification). † The substitution parameter ψ is not estimated, but follows from δ and ν.

Table 11: Targets used in model estimation

	(I) Baseline				(II) No Controls			
	(1) Log Output	(2) Log Price	(3) Log Materials	(4) Log Sales	(5) Log Output	(6) Log Price	(7) Log Materials	(8) Log Sales
t	0.40***	0.11***	0.53***	0.55***	0.41***	0.14***	0.50***	0.55***
	(0.064)	(0.027)	(0.055)	(0.062)	(0.057)	(0.024)	(0.050)	(0.054)
t+1	0.36***	0.15***	0.56***	0.57***	0.39***	0.16***	0.54***	0.59***
	(0.085)	(0.037)	(0.074)	(0.079)	(0.078)	(0.035)	(0.069)	(0.072)
t+2	0.30***	0.17***	0.36***	0.54***	0.29***	0.17***	0.37***	0.51***
	(0.11)	(0.046)	(0.090)	(0.096)	(0.095)	(0.044)	(0.083)	(0.086)
t+3	0.29**	0.20***	0.38***	0.46***	0.26**	0.22***	0.37***	0.43***
	(0.13)	(0.055)	(0.098)	(0.11)	(0.11)	(0.051)	(0.092)	(0.10)
t+4	0.36**	0.20***	0.44***	0.54***	0.31**	0.22***	0.42***	0.50***
	(0.14)	(0.069)	(0.12)	(0.13)	(0.12)	(0.060)	(0.11)	(0.11)
t+5	0.27*	0.18**	0.48***	0.50***	0.27*	0.23***	0.48***	0.48***
	(0.16)	(0.073)	(0.13)	(0.14)	(0.14)	(0.069)	(0.12)	(0.13)
Firms	2,016	647	2,109	2,016	2,316	771	2,353	2,316
N	75,162	92,465	99,991	75,162	158,115	193,100	187,827	158,115

	(III) Output Price Index				(IV) Only PPI			
	(9) Log Output	(10) Log Price	(11) Log Materials	(12) Log Sales	(13) Log Output	(14) Log Price	(15) Log Materials	(16) Log Sales
t	0.40***	0.12***	0.53***	0.55***	0.41***	0.11***	0.56***	0.57***
	(0.064)	(0.041)	(0.055)	(0.062)	(0.083)	(0.027)	(0.078)	(0.076)
t+1	0.36***	0.18***	0.56***	0.57***	0.46***	0.15***	0.57***	0.67***
	(0.085)	(0.057)	(0.074)	(0.079)	(0.11)	(0.037)	(0.098)	(0.10)
t+2	0.30***	0.19***	0.36***	0.54***	0.34***	0.17***	0.36***	0.55***
	(0.11)	(0.070)	(0.090)	(0.096)	(0.13)	(0.046)	(0.12)	(0.13)
t+3	0.29**	0.14*	0.38***	0.46***	0.34**	0.20***	0.40***	0.53***
	(0.13)	(0.080)	(0.098)	(0.11)	(0.15)	(0.055)	(0.14)	(0.14)
t+4	0.36**	0.17*	0.44***	0.54***	0.36**	0.20***	0.39**	0.62***
	(0.14)	(0.095)	(0.12)	(0.13)	(0.17)	(0.069)	(0.18)	(0.16)
t+5	0.27*	0.18*	0.48***	0.50***	0.22	0.18**	0.46**	0.47**
	(0.16)	(0.11)	(0.13)	(0.14)	(0.19)	(0.073)	(0.20)	(0.19)
Firms	2,016	2,026	2,109	2,016	756	647	768	756
N	75,162	76,488	99,991	75,162	38,526	92,465	48,143	38,526

Notes: The table summarizes the four sets of local projection estimates used for impulse response matching. Those estimates are partially shown in other tables and figures and restated here. Panel (I) shows the set of targets used in our baseline impulse response matching based on the baseline local projection specification with controls. Panel (II) to (IV) shows the targets used in robustness checks summarized in Table 10. Panel (II) shows the targets based on our local projection specification without controls. Panel (III) shows the targets with firms' output price index instead of domestic PPI prices used as a target for prices. Panel (IV) shows the targets when we restrict the sample to firms in the PPI. The number of distinct firms and observations refer to the union of all samples across local projection horizons.

Table 12: Model variations and extensions

	(1) Baseline	(2) restricted to Cobb-Douglas	(3) with Kimball demand	(4) with Monopsony	(5) w. Kimball demand & Monopsony
(I) Estimated parameters					
Slope of flex-price supply δ	0.571	0.441	0.571	0.564	0.564
	(0.172)	(0.130)	(0.172)	(0.164)	(0.164)
Slope of demand σ	4.460	3.835	4.460	4.263	4.263
	(1.293)	(1.365)	(1.293)	(1.242)	(1.242)
Fixed input share ν	0.163	0.306	0.163	0.153	0.153
	(0.091)	(0.063)	(0.091)	(0.103)	(0.103)
Materials share in flexible input α	0.785	0.994	0.785	0.787	0.787
	(0.129)	(0.188)	(0.129)	(0.144)	(0.144)
Demand scaling parameter ζ	1.147	0.928	1.147	1.125	1.125
	(0.174)	(0.140)	(0.174)	(0.172)	(0.172)
Firms' labor supply elasticity η	—	—	—	0.099	0.099
				(0.052)	(0.052)
Implied substitution parameter $\psi^{\dagger}$	-1.940	—	-5.851	-1.984	-6.122
	(1.817)		(4.234)	(2.107)	(5.125)
Estimation fit	10.883	16.708	10.883	12.261	12.261
J-test p-value	0.539	0.213	0.539	0.659	0.659
(II) Calibrated parameters					
Discount factor β	0.990	0.990	0.990	0.990	0.990
Price stickiness θ	0.650	0.650	0.650	0.650	0.650
Firm-level demand persistence ρ	0.992	0.992	0.992	0.992	0.992
Markup elasticity Γ	—	—	1.330	—	1.330
(III) Phillips curve slope					
Capacity pressure κ^y	0.039	0.058	0.039	0.038	0.038
	(0.014)	(0.030)	(0.014)	(0.016)	(0.015)
Bootstrapped 95% CI for κ^y	(0.020, 0.069)	(0.020, 0.060)	(0.020, 0.070)	(0.019, 0.070)	(0.020, 0.070)
Overall slope κ (w. $\phi = 1$)	0.059	0.059	0.052	0.058	0.052
Ratio κ^y/κ	0.66	0.99	0.74	0.65	0.74

Notes: The table shows model estimation results for variations and extended versions of the model. The extended model is described in online Appendix C. The estimation is described in Section 4. The empirical target is our baseline specification "with controls". Column (1) restates our baseline estimates. Column (2) reports estimates for the baseline model restricted to a Cobb-Douglas production function. Column (3) estimates for the model with Kimball demand and $\Gamma = 1.33$. Column (4) reports estimates with monopsony, using the response of firms' average wage to demand shocks reported in Table 9 as an additional target. Column (5) reports estimates for the model with both extensions. † The substitution parameter ψ is not estimated, but follows from δ and ν.

Figure 12: Sensitivity of estimates to varying β

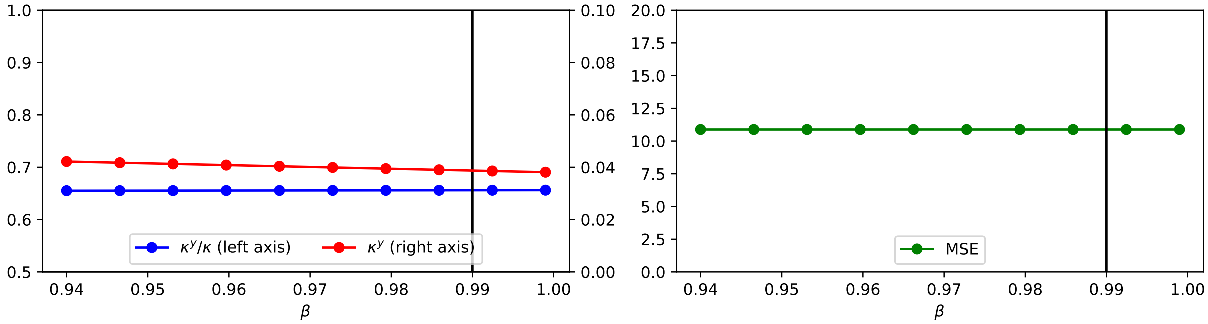

Figure 13: Sensitivity of estimates to varying ρ

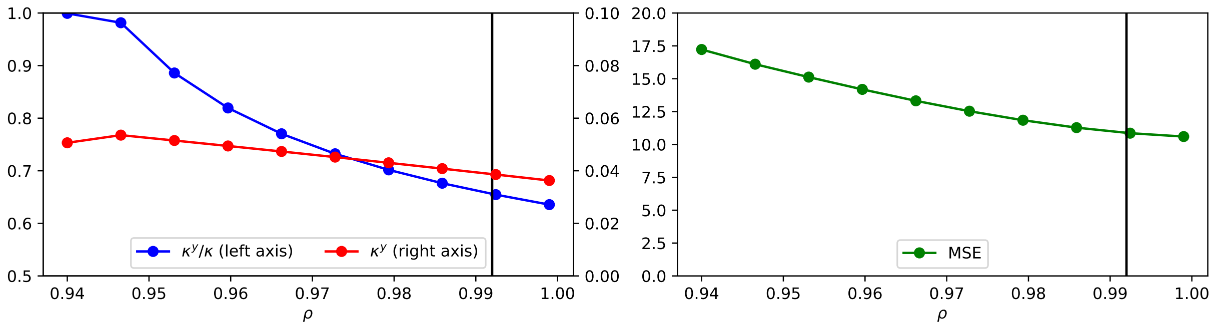

Figure 14: Sensitivity of estimates to varying θ

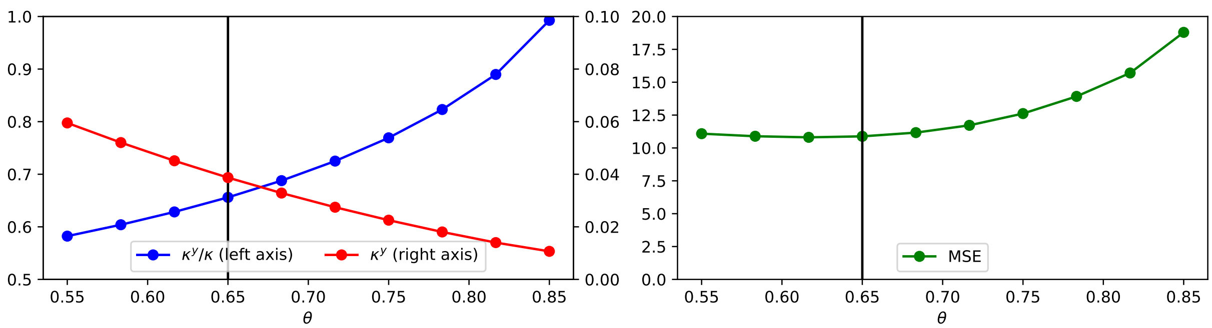

Figure 15: Sensitivity of estimates to varying Γ

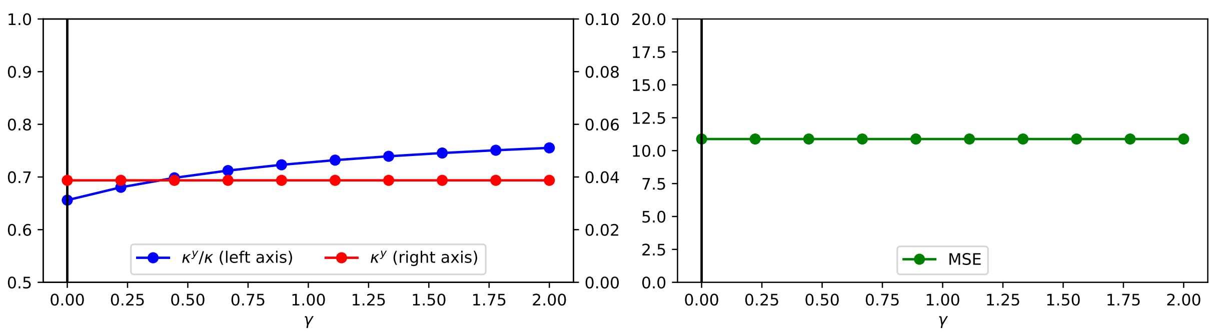

Notes: Figures 12, 13, 14 and 15 show the estimated magnitude of the capacity pressure channel and its share in the overall Phillips curve with different calibrated values of β, ρ, θ and Γ. The baseline calibration is $\beta = 0.99$, $\rho = 0.992$, $\theta = 0.65$ and $\Gamma = 0$ ($\Gamma = 1.33$ in the Kimball demand extension). We vary each parameter keeping the others fixed to their baseline value. We estimate the model parameters as described in Section 4. The right figures show the minimum value of the estimation objective function. The black bars illustrate the baseline parameter values.

Online Appendix

B Data

B.1 Construction of firm-level output indices

We measure firm-level real output by deflating nominal sales with a firm-specific price index based on unit-value changes. The main complication is that product definitions change over time: the 8-digit Combined Nomenclature (CN8) codes used to classify products in the VARS dataset may be relabeled, split, or merged across years. We address this using the official correspondence tables between CN schemas.

Data. VARS contains the micro data underlying the Danish version of PRODCOM. It provides sales values $v_{i,p,t}$ and quantities $q_{i,p,t}$ at the firm-product-quarter level, where products are defined at the CN8 level.

Mapping groups. To construct comparable unit values over time, we use the official correspondence between CN schemas to define year-to-year *mapping groups* $\mathcal{R}_{t-1,t}$. Each mapping group $r \in \mathcal{R}_{t-1,t}$ pairs a set of CN8 codes $\mathcal{P}_r^{t-1}$ valid in year $t-1$ with a corresponding set $\mathcal{P}_r^t$ valid in year t that represents the same underlying products. When a code is simply relabeled, both sets are singletons. When codes are split or merged, the sets contain multiple elements. While the aggregate correspondence tables involve many complex many-to-many mappings, at the firm level most mappings reduce to one-to-one correspondences since individual firms produce only a limited set of products.

Unit-value changes. For each firm i and mapping group $r \in \mathcal{R}_{t-1,t}$, we aggregate sales values and quantities across all CN8 codes in the mapping group to obtain group-level sales $V_{i,r,t} = \sum_{p \in \mathcal{P}_r^t} v_{i,p,t}$ and quantities $Q_{i,r,t} = \sum_{p \in \mathcal{P}_r^t} q_{i,p,t}$, and analogously for period $t-1$. The unit value is $UV_{i,r,t} = V_{i,r,t}/Q_{i,r,t}$. We restrict attention to mapping groups that have positive sales in both periods—groups where a product is introduced or discontinued do not yield a comparable unit-value change. We compute unit-value changes $\Delta \ln UV_{i,r,t} = \ln UV_{i,r,t} - \ln UV_{i,r,t-1}$ and winsorize them between -1 and 1.

Firm-level price index. We construct firm-level price changes as a value-weighted average of unit-value changes across mapping groups:

$$\Delta \ln P_{i,t} = \sum_{r \in \mathcal{R}_{t-1,t}} \omega_{i,r,t} \, \Delta \ln UV_{i,r,t}$$

where the weights $\omega_{i,r,t} = V_{i,r,t} / \sum_{r'} V_{i,r',t}$ are the value shares of each mapping group in the firm's total sales of continuing products. The firm-level price index is obtained by compounding these price changes from the first year t_0 a firm is observed in an uninterrupted sequence of observations: $P_{i,t} = \prod_{s=t_0+1}^{t} \exp(\Delta \ln P_{i,s})$.

Real output. Firm-level real output is total nominal sales deflated by the firm-specific price index:

$$Y_{i,t} = \frac{\sum_p v_{i,p,t}}{P_{i,t}}$$

B.2 Construction of shift-share demand shifters

The exports-based demand shifters consist of two sub-components: country-product-level imports for the shifters, and firm-product-level exports for the shares.

The country-product shifters are based on the UN Comtrade database, which contains an annual panel of the US Dollar value of countries' aggregate goods imports at the source-country-product level, using 6-digit Harmonized System (HS6) product codes. We aggregate these imports to the 4-digit level (HS4). To ensure that changes in a country's imports of a specific product are not driven by shocks to Danish firms, we exclude imports originating from Denmark from each country-product import cell. We additionally drop import series for countries with gaps in coverage or fewer than 20 years of observations over the 1996–2021 period. In practice, this only affects a negligible share of Danish manufacturing exports. We then deflate to constant US Dollars using the US CPI and calculate real annual import growth rates $\Delta im_{k,j,t}$ for country k, product j, and year t.

The firm-product-level export shares are based on the UHDM registry maintained by Statistics Denmark, a monthly panel containing all Danish exports at the firm-product level using 8-digit Combined Nomenclature (CN8) product codes. We first aggregate to the yearly and HS4 product level. 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). We then divide each firm's country-product export value by its total sales from VARS rather than total exports:

$$\omega_{i,k,j,t} = \frac{\text{EXPORTS}_{i,k,j,t}}{\text{SALES}_{i,t}}$$

This ensures that each market's weight reflects its overall importance to the firm's business, rather than merely its export exposure. When constructing the shift-share demand shifter, we use lagged shares when aggregating, as indicated in equation (18), reproduced here for convenience:

$$\Delta z_{i,t} = \sum_{k \in K} \sum_{j \in J} \omega_{i,k,j,t-1} \, \Delta im_{k,j,t}$$

C Full model with extensions

This appendix presents the complete specification of our New Keynesian model with firm-level demand shocks. The model nests the baseline version presented in Section 2 (for the special case $\Gamma = 0$ and $\eta = 0$) and incorporates extensions with Kimball demand ($\Gamma > 0$) and monopsony power in labor markets ($\eta > 0$). We follow the notation of the main text throughout.

C.1 Model environment

Final goods production. Gross final output Y_t is produced by perfectly competitive firms using a Kimball aggregator that combines differentiated intermediate goods:

$$\int_i Z_{i,t} \Upsilon\left(\frac{Y_{i,t}}{Z_{i,t} Y_t}\right) di = 1 \tag{39}$$

Following Klenow and Willis (2016), we specify the aggregator function Υ such that cost minimization by final goods producers yields the demand curve:

$$Y_{i,t} = Y_t Z_{i,t} \Psi\left(\frac{P_{i,t}}{P_t D_t}\right) \tag{40}$$

where D_t is a price dispersion index and $\Psi(\cdot)$ is a decreasing function. The demand elasticity faced by firm i and the relative-price elasticity of the desired markup are:

$$\varepsilon(x_{i,t}) = -\frac{\Psi'(x_{i,t})}{\Psi(x_{i,t})} x_{i,t}, \qquad \Gamma(x_{i,t}) = -\frac{d \log \mu(x_{i,t})}{d \log x_{i,t}} \tag{41}$$

where $x_{i,t} = P_{i,t}/(P_t D_t)$. With CES aggregation, demand is isoelastic with $\varepsilon = \sigma$ and $\Gamma = 0$. With the Kimball specification of Υ, both ε and Γ vary with the firm's relative price. Around the symmetric zero-inflation steady state ($x_{i,t} = 1$ for all i), both reduce to constants:

$$\varepsilon = \sigma, \qquad \Gamma = \frac{\tau}{\sigma - 1} \tag{42}$$

where $\tau \geq 0$ is the curvature parameter of the Kimball aggregator ($\tau = 0$ recovers CES). We calibrate $\Gamma = 1.33$ following Gagliardone et al. (2025) and Amiti et al. (2019). Since all our results require only σ and Γ evaluated at steady state, we work with these two parameters directly and suppress the functional forms of Υ and Ψ from here on. The log-linearized demand curve around the steady state is:

$$y_{i,t} - y_t = z_{i,t} - \sigma(p_{i,t} - p_t) \tag{43}$$

where lowercase letters denote log deviations from steady-state values.

The variable $Z_{i,t}$ is an idiosyncratic demand shifter that follows an AR(1) process in logs:

$$z_{i,t} = \rho z_{i,t-1} + u_{i,t} \tag{44}$$

where $u_{i,t}$ is i.i.d. with mean zero.

Roundabout production structure. Gross final output can be decomposed into value added (GDP) and material inputs used in intermediate production:

$$Y_t = Y_t^{gdp} + M_t \tag{45}$$

Since final good producers are competitive and use intermediates as their only input, the CES price index P_t is the price of gross final output, value added, and materials.

Intermediate goods production. Intermediate goods producers use a fixed factor $K_{i,t}$, labor $L_{i,t}$, and materials $M_{i,t}$ to produce output. The production function takes a normalized CES form in the fixed factor and a flexible input composite:

$$\frac{Y_{i,t}}{Y_{SS}} = \left(\nu \left(\frac{K_{i,t}}{K_{SS}}\right)^{\psi} + (1-\nu) X_{i,t}^{\psi}\right)^{1/\psi} \tag{46}$$

where the flexible composite $X_{i,t}$ is Cobb-Douglas in labor and materials:

$$X_{i,t} = \left(\frac{L_{i,t}}{L_{SS}}\right)^{1-\alpha} \left(\frac{M_{i,t}}{M_{SS}}\right)^{\alpha} \tag{47}$$

We assume $K_{i,t} = K_{SS}$ (the fixed factor is fixed at its steady-state value). The parameter $\nu \in [0,1]$ determines the share of the fixed factor, $\alpha \in [0,1]$ is the materials share within the flexible composite, and $\psi \leq 1$ governs the elasticity of substitution between fixed and flexible inputs ($\psi \to 0$: Cobb-Douglas; $\psi \to -\infty$: Leontief).

Labor markets with monopsony. Firms face an upward-sloping labor supply curve:

$$L_{i,t} = L_t \left(\frac{W_{i,t}}{W_t}\right)^{1/\eta} \tag{48}$$

The parameter $\eta \geq 0$ governs the degree of monopsony power. The limiting case $\eta \to 0$ corresponds to competitive labor markets with perfectly elastic labor supply. This functional form can be microfounded using a framework with differentiated jobs and workers with extreme-value preferences over employers as in Berger et al. (2022).

Price setting. Intermediate producers set prices subject to Calvo frictions: each period, a firm resets its price with probability $1-\theta$.

C.2 Firm optimization

We now characterize intermediate goods producers' optimal input choices and pricing decisions. All derivations are for the general model with both monopsony ($\eta \geq 0$) and Kimball demand ($\Gamma \geq 0$). The baseline model corresponds to $\eta = 0$ and $\Gamma = 0$.

Cost minimization. With the labor supply curve (48), the effective cost of hiring $L_{i,t}$ workers is $W_{i,t}L_{i,t} = W_t L_t^{-\eta} L_{i,t}^{1+\eta}$, where the firm internalizes that expanding employment raises the wage it must pay. Firms choose labor and materials to minimize total variable cost:

$$\min_{L_{i,t},M_{i,t}} W_t L_t^{-\eta} L_{i,t}^{1+\eta} + P_t M_{i,t} \tag{49}$$

subject to the production function (46)–(47) with $K_{i,t} = K_{SS}$.

Since the flexible composite is Cobb-Douglas, the first-order conditions for $L_{i,t}$ and $M_{i,t}$ yield a constant expenditure ratio:

$$\frac{(1+\eta)W_{i,t}L_{i,t}}{P_t M_{i,t}} = \frac{1-\alpha}{\alpha} \tag{50}$$

The factor $(1+\eta)$ reflects the marginal cost of labor inclusive of the monopsony wage effect: hiring one additional unit of labor costs the firm $W_{i,t}(1+\eta)$ rather than $W_{i,t}$, because expanding employment raises the wage on all inframarginal workers.

Marginal cost in terms of output. Using (50) and the CES production function, the marginal cost of firm i can be expressed as:

$$MC_{i,t} = \frac{1}{1-\nu} \frac{W_{i,t}^{1-\alpha} P_t^{\alpha} (1+\eta)^{1-\alpha}}{(1-\alpha)^{1-\alpha}\alpha^{\alpha}} \left(\frac{X_{i,t}}{X_{SS}}\right)^{1-\psi} \left(\frac{Y_{i,t}}{Y_{SS}}\right)^{\psi-1} \tag{51}$$

This expression has a direct economic interpretation. The first fraction is the unit cost of the flexible composite (accounting for monopsony). The remaining terms capture the CES production structure: as the firm expands output relative to its fixed factor, the ratio $X_{i,t}/Y_{i,t}$ increases (the flexible composite must do more of the work), raising marginal cost.

Optimal pricing. When resetting its price, a firm chooses $P_{i,t}^*$ to maximize expected discounted profits:

$$\max_{P_{i,t}^*} \sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(P_{i,t}^* Y_{i,t+k} - C(Y_{i,t+k}, W_{t+k}, P_{t+k})\right) \tag{52}$$

subject to the demand curve (40), where $C(\cdot)$ denotes the cost function derived from the cost minimization problem (49) and $Y_{i,t+k}$ is the demand the firm faces at price $P_{i,t}^*$ in period $t+k$. Differentiating with respect to $P_{i,t}^*$ and using $MC_{i,t+k} \equiv \partial C/\partial Y_{i,t+k}$, the first-order condition is:

$$\sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(Y_{i,t+k}(1-\varepsilon_{i,t+k}) \left(P_{i,t}^* - \frac{\varepsilon_{i,t+k}}{\varepsilon_{i,t+k}-1} MC_{i,t+k} \right) \right) = 0 \tag{53}$$

where $\varepsilon_{i,t+k}$ is the demand elasticity faced by firm i in period $t+k$, evaluated at the firm's relative price. The ratio $\varepsilon_{i,t+k}/(\varepsilon_{i,t+k}-1)$ is the firm's desired markup. With CES demand ($\Gamma = 0$), $\varepsilon_{i,t+k} = \sigma$ is constant and the desired markup equals $\sigma/(\sigma-1)$. With Kimball demand ($\Gamma > 0$), $\varepsilon_{i,t+k}$ varies with the firm's relative price, so the desired markup falls when the firm charges a high relative price.

C.3 Log-linearization

We now log-linearize the model around the symmetric zero-inflation steady state. We first derive linearized expressions for firm-level marginal cost and the optimal reset price, then aggregate to the Phillips curve in Section C.4, and finally characterize firm-level impulse responses in Section C.5.

Firm-level marginal cost. We log-linearize the production function (46) with $K_{i,t} = K_{SS}$:

$$y_{i,t} = (1-\nu)x_{i,t} \tag{54}$$

where we use the normalization that $K_{i,t}/K_{SS} = 1$. Log-linearizing the Cobb-Douglas composite (47):

$$x_{i,t} = (1-\alpha)l_{i,t} + \alpha m_{i,t} \tag{55}$$

Log-linearizing the labor supply curve (48):

$$w_{i,t} - w_t = \eta(l_{i,t} - l_t) \tag{56}$$

Log-linearizing the expenditure ratio (50):

$$w_{i,t} + l_{i,t} = m_{i,t} + p_t \tag{57}$$

Note that the η terms cancel in the log-linearized expenditure ratio since $W_{i,t}(1+\eta)$ is already accounted for in steady state. Combining equations (57) and (56), $m_{i,t} - l_{i,t} = w_t - p_t + \eta(l_{i,t} - l_t)$. In deviations from the aggregate: $m_{i,t} - m_t - (l_{i,t} - l_t) = \eta(l_{i,t} - l_t)$, so:

$$m_{i,t} - m_t = (1+\eta)(l_{i,t} - l_t) \tag{58}$$

This says that when a firm expands, materials increase more than employment (by a factor $1+\eta$), because the rising wage under monopsony makes the firm substitute toward materials. Substituting (58) into the production function (54)–(55), we can express all input deviations in terms of output deviations:

$$l_{i,t} - l_t = \frac{1}{(1-\nu)(1+\alpha\eta)}(y_{i,t} - y_t) \tag{59}$$

$$m_{i,t} - m_t = \frac{1+\eta}{(1-\nu)(1+\alpha\eta)}(y_{i,t} - y_t) \tag{60}$$

$$w_{i,t} - w_t = \frac{\eta}{(1-\nu)(1+\alpha\eta)}(y_{i,t} - y_t) \tag{61}$$

To derive (59), substitute (58) into the linearized composite $x_{i,t} - x_t = (1-\alpha)(l_{i,t} - l_t) + \alpha(m_{i,t} - m_t) = [(1-\alpha) + \alpha(1+\eta)](l_{i,t} - l_t) = (1+\alpha\eta)(l_{i,t} - l_t)$, and use $y_{i,t} - y_t = (1-\nu)(x_{i,t} - x_t)$. Equations (60) and (61) follow from (58) and (56).

Log-linearizing marginal cost (51) and using the input-output relationships above, we obtain:

$$mc_{i,t}^R = mc_t^R + \chi^i (y_{i,t} - y_t) \tag{62}$$

where the elasticity of firm-level real marginal cost to firm-level output (relative to aggregate) is:

$$\chi^i = \underbrace{\frac{(1-\psi)\nu}{1-\nu}}_{\text{Capacity pressure}} + \underbrace{\frac{(1-\alpha)\eta}{(1-\nu)(1+\alpha\eta)}}_{\text{Monopsony}} \tag{63}$$

The first term reflects the direct effect of diminishing returns: as the firm expands output against its fixed factor, the marginal product of the flexible composite falls. The second term captures the monopsony channel: expanding output requires more labor, which raises the firm-specific wage, increasing marginal cost. In the baseline model with competitive labor markets ($\eta = 0$), only the first term remains.

Aggregate marginal cost. Having characterized firm-level marginal cost, we now derive its aggregate counterpart. After an aggregate demand shock, all firms expand symmetrically and wages rise uniformly. The monopsony term therefore drops out at the aggregate level:

$$mc_t^R = (1-\alpha) w_t^R + \chi^a y_t \tag{64}$$

where:

$$\chi^a = \frac{(1-\psi)\nu}{1-\nu} \tag{65}$$

The term $(1-\alpha) w_t^R$ captures the effect of aggregate real wages on marginal cost. The term $\chi^a y_t$ captures the capacity pressure effect at the aggregate level—this is the same as the first term of χ^i because it reflects only the fixed-factor channel.

Optimal reset price. We log-linearize the Calvo pricing first-order condition (53) around the symmetric zero-inflation steady state. With Kimball demand, the log-linearized FOC takes the form:

$$p_{i,t}^* - p_{t-1} = (1-\beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k \frac{1}{1+\Gamma} E_t\left(mc_{i,t+k}^R\right) + \sum_{k=0}^{\infty} (\beta\theta)^k E_t(\pi_{t+k}) \tag{66}$$

Relative to the baseline CES case, the factor $1/(1+\Gamma)$ attenuates the pass-through of marginal cost due to strategic complementarities in price-setting.

We now substitute the expression for firm-level marginal cost. Using the log-linearized demand curve (43), firm-level output deviations from aggregate are $y_{i,t+k} - y_{t+k} = z_{i,t+k} - \sigma(p_{i,t+k} - p_{t+k})$. Firms that set their price in period t and have not yet re-optimized by period $t+k$ have $p_{i,t+k} = p_{i,t}^*$. Substituting into (62):

$$mc_{i,t+k}^R = mc_{t+k}^R + \chi^i \left(z_{i,t+k} - \sigma(p_{i,t}^* - p_{t+k})\right) \tag{67}$$

Plugging into (66) and collecting the terms involving $p^*_{i,t}$ on the left-hand side:

$$p^*_{i,t} - p_{t-1} = \frac{1-\beta\theta}{1+\sigma\delta} \sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(\frac{1}{1+\Gamma} mc^R_{t+k} + \delta z_{i,t+k} \right) + \sum_{k=0}^{\infty} (\beta\theta)^k E_t(\pi_{t+k}) \tag{68}$$

where we define the slope of the flexible-price supply curve as:

$$\delta \equiv \frac{\chi^i}{1+\Gamma} \tag{69}$$

In the baseline model ($\Gamma = 0$, $\eta = 0$), this reduces to $\delta = (1-\psi)\nu/(1-\nu)$. With Kimball demand, δ is smaller for given χ^i because incomplete pass-through dampens the price response. With monopsony, χ^i is larger because the wage channel adds to marginal cost sensitivity.

Using $E_t[z_{i,t+k}] = \rho^k z_{i,t}$ and evaluating the geometric sum:

$$p^*_{i,t} - p_{t-1} = \frac{1-\beta\theta}{1-\beta\theta\rho} \frac{\delta}{1+\sigma\delta} z_{i,t} + \frac{1-\beta\theta}{1+\sigma\delta} \sum_{k=0}^{\infty} (\beta\theta)^k E_t \left(\frac{1}{1+\Gamma} mc^R_{t+k} \right) + \sum_{k=0}^{\infty} (\beta\theta)^k E_t(\pi_{t+k}) \tag{70}$$

This is the general reset price equation corresponding to equation (7) in the main text. The first term captures the effect of the firm-specific demand state. The second and third terms depend on aggregate conditions and are common to all firms resetting in period t.

C.4 Phillips curve

We now aggregate the firm-level pricing decisions to derive the Phillips curve.

Marginal cost Phillips curve. Since $z_{i,t}$ has mean zero across firms, the aggregate reset price follows from integrating (70) over firms:

$$p^*_t - p_{t-1} = \frac{1-\beta\theta}{(1+\sigma\delta)(1+\Gamma)} mc^R_t + \beta\theta E_t(p^*_{t+1} - p_t) + \pi_t \tag{71}$$

Substituting the Calvo aggregation formula $\pi_t = (1-\theta)(p^*_t - p_{t-1})$ and rearranging:

$$\pi_t = \frac{(1-\beta\theta)(1-\theta)}{\theta(1+\sigma\delta)(1+\Gamma)} mc^R_t + \beta E_t(\pi_{t+1}) \tag{72}$$

This is the marginal cost Phillips curve. In the CES baseline ($\Gamma = 0$), the coefficient on marginal cost is $\lambda = (1-\beta\theta)(1-\theta)/(\theta(1+\sigma\delta))$, as in the main text.

Phillips curve in terms of gross output. Substituting the aggregate marginal cost expression (64):

$$\pi_t = \frac{(1-\beta\theta)(1-\theta)}{\theta} \left(\frac{\chi^a}{(1+\sigma\delta)(1+\Gamma)} y_t + \frac{(1-\alpha)}{(1+\sigma\delta)(1+\Gamma)} w^R_t \right) + \beta E_t(\pi_{t+1}) \tag{73}$$

Using $\delta = \chi^i/(1+\Gamma)$ and the definition of χ^a, the coefficient on y_t can be rewritten as:

$$\frac{\chi^a}{(1+\sigma\delta)(1+\Gamma)} = \frac{\delta}{1+\sigma\delta}\frac{\chi^a}{\chi^i} \tag{74}$$

This decomposition clarifies that the output coefficient in the Phillips curve equals the flexible-price supply curve response $\delta/(1+\sigma\delta)$ scaled by χ^a/χ^i. In the baseline without monopsony ($\eta = 0$), $\chi^a = \chi^i$ and this ratio equals 1. With monopsony, $\chi^a < \chi^i$ because the firm-level monopsony channel does not operate at the aggregate level. The Phillips curve in terms of gross output is therefore:

$$\pi_t = \frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{\delta}{1+\sigma\delta}\frac{\chi^a}{\chi^i}y_t + \frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{(1-\alpha)}{(1+\sigma\delta)(1+\Gamma)}w_t^R + \beta E_t(\pi_{t+1}) \tag{75}$$

From gross output to GDP. To express the Phillips curve in terms of GDP y_t^{gdp}, we need the relationship between gross output and GDP. Log-linearizing $Y_t = Y_t^{gdp} + M_t$ around steady state:

$$y_t = (1-s_m)y_t^{gdp} + s_m m_t \tag{76}$$

where $s_m = M_{SS}/Y_{SS} = \alpha(1-\nu)/\mu$ is the steady-state materials share in gross output and $\mu = \sigma/(\sigma-1)$ is the steady-state markup. To eliminate m_t, we use the linearized expenditure ratio (57) and production function (54)–(55) in the aggregate (where all firms expand symmetrically) to obtain $m_t = y_t/(1-\nu) + (1-\alpha)w_t^R$. Substituting into (76) and solving for y_t:

$$y_t = \frac{(1-s_m)(1-\nu)}{1-\nu-s_m}y_t^{gdp} + \frac{s_m(1-\alpha)(1-\nu)}{1-\nu-s_m}w_t^R \tag{77}$$

The coefficient on y_t^{gdp} exceeds 1 whenever $s_m > 0$, because a given increase in GDP requires a more-than-proportional increase in gross output to also produce the additional material inputs.

Phillips curve in terms of GDP. Substituting (77) into the Phillips curve and using $w_t^R = \phi y_t^{gdp}$, where ϕ is the reduced-form elasticity of real wages with respect to GDP:

$$\pi_t = (\kappa^y + \kappa^w)y_t^{gdp} + \beta E_t(\pi_{t+1}) \tag{78}$$

where:

$$\kappa^y = \frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{\delta}{1+\sigma\delta}\frac{\chi^a}{\chi^i}\frac{(1-s_m)(1-\nu)}{1-\nu-s_m} \tag{79}$$

$$\kappa^w = \frac{(1-\beta\theta)(1-\theta)}{\theta}\frac{\phi(1-\alpha)}{(1+\sigma\delta)(1+\Gamma)}\frac{(1-s_m)(1-\nu)-\psi\nu s_m}{1-\nu-s_m} \tag{80}$$

κ^y is the capacity pressure channel and κ^w is the wage pressure channel. The total Phillips curve slope is $\kappa = \kappa^y + \kappa^w$. To derive (80), collect the terms multiplying ϕy_t^{gdp} from both the direct wage term in the Phillips curve and the indirect effect through y_t, and use $\chi^a = (1-\psi)\nu/(1-\nu)$ to simplify.

Special cases. In the CES baseline ($\Gamma = 0$, $\eta = 0$): $\chi^a = \chi^i$, so $\chi^a/\chi^i = 1$ and $\delta = (1-\psi)\nu/(1-\nu)$. The expressions reduce to those in equation (17) of the main text.

With Kimball demand only ($\Gamma > 0, \eta = 0$): $\chi^a/\chi^i = 1$ still holds, so κ^y is unchanged—the incomplete pass-through is already captured in δ. The wage channel is attenuated by $1/(1+\Gamma)$.

With monopsony only ($\Gamma = 0$, $\eta > 0$): $\chi^a/\chi^i < 1$ because χ^i includes the monopsony term that does not operate at the aggregate level.

C.5 Firm-level impulse responses

We now derive the impulse response of firm-level variables to idiosyncratic demand shocks. These expressions form the basis for our impulse response matching estimation in Section 4. Following the main text, we denote as $\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})$ the average period $t+k$ price of firms with a period t history of idiosyncratic demand realizations $\mathcal{Z}_{i,t-1} = \{z_{i,\tau}\}_{\tau=-\infty}^{t-1}$ and current demand realization $z_{i,t}$, relative to the aggregate price level. The reset price $\tilde{p}^*_{t+k}(z_{i,t})$ depends only on the current realization. We use analogous notation for other variables.

Price response. From the reset price equation (70), the relative reset price of a firm hit by shock $z_{i,t}$ is:

$$\tilde{p}^*_{t+k}(z_{i,t}) = \frac{1-\beta\theta}{1-\beta\theta\rho}\frac{\delta}{1+\sigma\delta}\rho^k z_{i,t} \tag{81}$$

This uses $E_t[z_{i,t+k}] = \rho^k z_{i,t}$ and the fact that the aggregate terms cancel when taking deviations from the average reset price.

The observed relative price in period $t+k$ (averaging over firms that have and have not reset since period t) satisfies the recursion:

$$\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t}) = (1-\theta)\,\tilde{p}^*_{t+k}(z_{i,t}) + \theta\,\tilde{p}_{t+k-1}(\mathcal{Z}_{i,t-1}, z_{i,t}) \tag{82}$$

Since period t demand shocks are orthogonal to the history of demand realizations up to that point, the boundary condition is $\partial\tilde{p}_{t-1}(\mathcal{Z}_{i,t-1}, z_{i,t})/\partial u_{i,t} = 0$.

Differentiating (82) with respect to $u_{i,t}$ and solving the recursion:

$$\frac{\partial\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{(1-\beta\theta)(1-\theta)}{1-\beta\theta\rho}\frac{\delta}{1+\sigma\delta}\cdot\frac{\theta^{k+1}-\rho^{k+1}}{\theta-\rho} \tag{83}$$

To verify, one can substitute this expression into the differentiated recursion and confirm that it satisfies both the recursion and the boundary condition.[16]

Output and input responses. From the log-linearized demand curve (43), the output response is:

$$\frac{\partial\tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \rho^k - \sigma\frac{\partial\tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{84}$$

[16] Denoting the coefficient $(1-\theta)\frac{1-\beta\theta}{1-\beta\theta\rho}\frac{\delta}{1+\sigma\delta}$ as A and trying the solution $A(\theta^{k+1}-\rho^{k+1})/(\theta-\rho)$, one verifies by substitution that the recursion holds and the boundary condition is satisfied since $(\theta^1-\rho^1)/(\theta-\rho) = 1$.

From equations (59)–(61), the input responses are proportional to the output response:

$$\frac{\partial \tilde{w}^R_{i,t+k}}{\partial u_{i,t}} = \frac{\eta}{(1-\nu)(1+\alpha\eta)} \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{85}$$

$$\frac{\partial \tilde{l}_{i,t+k}}{\partial u_{i,t}} = \frac{1}{(1-\nu)(1+\alpha\eta)} \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{86}$$

$$\frac{\partial \tilde{m}_{i,t+k}}{\partial u_{i,t}} = \frac{1+\eta}{(1-\nu)(1+\alpha\eta)} \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{87}$$

In the baseline without monopsony ($\eta = 0$), the input responses collapse to $\frac{1}{1-\nu} \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}}$, as stated in equation (13) of the main text. The ratio of the materials response to the employment response is $1+\eta$: under monopsony, firms substitute toward materials (whose price is unaffected by the firm's hiring decisions) and away from labor.

Value-added responses. Since firm i has measure zero, the aggregate price level is unchanged by an idiosyncratic shock. We can therefore convert the gross output responses to value-added equivalents, which allows us to compare our roundabout production model to a standard model without intermediate inputs as in Section 4. Define real value added as $VA_i = Y_i - M_i$ and the value-added price deflator P^{VA}_i via $P^{VA}_i \cdot VA_i = P_i Y_i - PM_i$. Log-linearizing around the symmetric steady state where $VA_{SS} = Y_{SS}(1 - s_m)$:

$$\tilde{va}_{i,t+k} = \frac{1}{1-s_m}\tilde{y}_{i,t+k} - \frac{s_m}{1-s_m}\tilde{m}_{i,t+k} \tag{88}$$

$$\tilde{p}^{VA}_{i,t+k} = \frac{1}{1-s_m}\tilde{p}_{i,t+k} \tag{89}$$

The value-added price conversion (89) follows because the firm's price change applies to gross revenue, while materials costs are unchanged (the aggregate price is fixed for an idiosyncratic shock). The entire nominal effect therefore accrues to value added, which is a smaller base than gross output, amplifying the price response by $1/(1-s_m)$. The impulse responses in value-added terms are:

$$\frac{\partial \tilde{p}^{VA}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{1}{1-s_m} \frac{\partial \tilde{p}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{90}$$

$$\frac{\partial \tilde{va}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} = \frac{1}{1-s_m} \frac{\partial \tilde{y}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} - \frac{s_m}{1-s_m} \frac{\partial \tilde{m}_{t+k}(\mathcal{Z}_{i,t-1}, z_{i,t})}{\partial u_{i,t}} \tag{91}$$