Anatomy of the Phillips Curve: Micro Evidence and Macro Implications

Luca Gagliardone* Mark Gertler[†] Simone Lenzu[‡] Joris Tielens[§]

July, 2023

Abstract

We use a unique high-frequency micro-dataset to estimate the slope of the primitive form of the New Keynesian Phillips curve, which features marginal cost as the relevant real activity variable. Our dataset encompasses product-level prices, costs, and output within the Belgian manufacturing sector over twenty years, recorded at a quarterly frequency. Leveraging the richness of the data, we adopt a "bottom-up" approach that identifies the Phillips curve's slope by estimating the primitive parameters that govern firms' pricing behavior, including the degrees of price rigidity and strategic complementarities in price setting. Our estimates indicate a high slope for the marginal cost-based Phillips curve, which contrasts with the low estimates of the conventional unemployment or output gap-based formulations in the literature. We reconcile the difference demonstrating that although the pass-through of marginal cost into inflation is substantial, the elasticity of marginal cost in relation to the output gap is low. Furthermore, through an examination of the transmission of oil shocks, we illustrate how the marginal cost-based Phillips curve offers a transparent means of capturing the impact of supply shocks on inflation.

JEL: E30.

* New York University. Email: luca.gagliardone@nyu.edu; † New York University and NBER. Email: mark.gertler@nyu.edu; ‡ New York University, Stern School of Business. Email: slenzu@stern.nyu.edu; § National Bank of Belgium and KU Leuven. Email: joris.tielens@nbb.be. We thank Fernando Alvarez, Xavier Gabaix, Jonathon Hazell, Man Chon Iao, Virgiliu Midrigan, Alessio Moro, Carolin Pflueger, Ludwig Straub, Gianluca Violante, Olivier Wang, and seminar participants at various institutions for helpful comments and conversations. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium, the Eurosystem, or any other institution with which the authors are affiliated.

1 Introduction

Understanding the drivers of inflation continues to be an important though unresolved matter in macroeconomics. At the heart of this inquiry lies the challenge of estimating the slope of the Phillips curve, which measures how inflation reacts to fluctuations in real economic activity. To illustrate the issue, let us consider the New Keynesian version of the Phillips curve (NKPC), which is now the textbook formulation in the literature. Let π_t denote inflation and \tilde{y}_t the output gap, the percentage difference between real output and its natural level. Then (what we will refer to as) the *conventional form* of the NKPC is given by:

$$\pi_t = \kappa \ \widetilde{y}_t + \beta \ \mathbb{E}_t \ \{\pi_{t+1}\} + u_t,\tag{1}$$

where u_t is a random disturbance capturing, e.g., cost shocks, and β is a subjective discount factor, typically a parameter close to unity. The NKPC asserts that inflation depends positively on both \tilde{y}_t , which can be interpreted as a measure of excess demand, and on expected future inflation. The main object of interest is κ , the slope coefficient on the output gap.

There are two interrelated sets of issues involved in uncovering κ . The first set revolves around the econometric identification of this parameter. First, as emphasized by McLeay and Tenreyro (2020), the output gap is an endogenous object. If the central bank acts to adjust \tilde{y}_t to stabilize π_t in response to positive cost shocks, the estimate of κ will be biased downward due to the negative correlation between \tilde{y}_t and u_t . Given the absence of good instruments for \tilde{y}_t , the estimation of κ using aggregate time-series data is problematic (Mavroeidis et al. 2014). Another identification issue involves trend inflation. The specification given by equation (1) presumes that trend inflation is constant. However, as emphasized by Hazell et al. (2022) and Jørgensen and Lansing (2023), shifts in trend inflation may confound the identification of the Phillips curve. For example, if trend inflation decreases as output declines and the regression model does not account for this correlation, the estimate of κ will be upwardly biased.

These identification challenges have led researchers to employ regional data to estimate κ . Recent examples include Hooper et al. (2020), McLeay and Tenreyro (2020), and Hazell et al. (2022).¹ In addition, Hooper et al. (2020) and Hazell et al. (2022) allow

¹Also relevant is Beraja et al. (2019), which uses regional data to identify wage Phillips curves.

for time fixed effects to control for shifting trend inflation. In the latter study, this identification approach yields an astonishingly small estimate of κ , which suggests that the Phillips curve is "flat". This view has become the conventional wisdom, at least for the pre-pandemic period.

The second set of considerations pertains to both the relevant measure of real activity that enters the Phillips curve and, consequently, the interpretation of the slope coefficient κ . In the underlying theory, firms optimize their pricing policies in response to current and anticipated movements in marginal cost. Thus, as emphasized by both Galí and Gertler (1999) and Sbordone (2002), the primitive form of the NKPC features real marginal cost (in percent deviations from trend) entering as the real activity variable. In fact, the conventional formulation of the NKPC in equation (1) only holds under specific conditions that establish a proportional relationship between marginal cost and the output gap. Among other things, wages must be perfectly flexible.² If these conditions are violated, then the output gap may not serve as an adequate proxy for real marginal cost, typically leading to a downward bias in the estimate of κ .³ Moreover, even if all conditions that establish a proportional relationship are approximately met, it is crucial to recognize that the output gap-based slope κ is ultimately the product of two parameters: the elasticity of inflation with respect to real marginal cost and the elasticity of marginal cost with respect to the output gap. The ability to separately identify the two coefficients is important for gaining a comprehensive understanding of inflation dynamics.

In this paper, we leverage a unique high-frequency micro dataset that provides information on prices, costs, and quantities of production to estimate the slope of the primitive form of the (marginal cost-based) NKPC. Similar to recent studies, we leverage the panel dimension of our dataset to enhance the identification approach: instruments are strong. However, our approach differs from prior studies as we utilize data at the individual firm-product level. By doing so, we are able to estimate firm-level pricing equations consistent with the underlying theory, which allow us to directly identify the primitive parameters that determine the slope of the marginal cost-based NKPC, including

²Indeed, it is for this reason that New Keynesian DSGE models with wage rigidity include the marginal cost-based Phillips curve in the system of equations as opposed to the conventional one (see Galí 2015 chapter 6 and the references therein).

³These considerations also extend to formulations of the conventional NKPC that utilizes the unemployment gap as a measure of economic activity instead of the output gap. They also apply to using an aggregate measure of real marginal cost such as the labor share.

the degree of price rigidity and strategic complementarities in price setting.⁴

Our estimates indicate that the slope of the marginal cost-based Phillips curve is high, suggesting a substantial pass-through of marginal cost into prices and inflation. This finding stands in stark contrast to the low estimates found in the existing literature using the conventional output gap-based formulation. To explore the origins of the difference, we use our firm-level data to estimate an output gap-based curve alongside the marginal cost-based formulation. Consistent with previous estimates, we find the former exhibits a substantially lower slope. We reconcile the wedge between the two slopes' estimates by showing that the implied elasticity of marginal cost with respect to the output gap is quite low. In other words, the slope of the conventional NKPC does not stem from a limited transmission of fluctuations in marginal cost to inflation, but rather from the weak connection between output gap and marginal cost.

Our analysis also stresses that an important reason why one would want to consider a marginal cost-based Phillips curve regards the transmission of supply shocks. These shocks can generate inflation via their influence on the natural level of output. Since the latter is not directly observable, it becomes problematic to use the output-based NKPC formulation to comprehend (and forecast) inflation dynamics during a period characterized by large supply shocks, as many would argue has been the case over the last few years. The primitive form of the NKPC does not suffer from this issue. Supply shocks have a direct impact on firms' marginal costs, and the slope of the marginal cost-based NKPC traces how their fluctuations impact inflation.

The paper proceeds as follows. In Section 2 we develop the theoretical framework that serves as the foundation of our estimation strategy. We start from first principles to derive a NKPC featuring nominal price rigidities and a general form of strategic complementarities in price setting, which encompasses both static oligopoly (Atkeson and Burstein 2008), dynamic oligopoly (Wang and Werning 2022), and monopolistic competition with variable elasticity of demand (Kimball 1995). With an eye toward the data work, we depart from the standard model by allowing for ex-ante heterogeneity across firms. In this environment, we derive the expression of the marginal-cost-based Phillips curve and its slope. As standard, the slope the curve depends on the parameter

⁴Notably, our identification procedure nests and enhances the one proposed by Hazell et al. (2022) for regional data, allowing us to address potential concerns related to shifts in trend inflation and inflation expectations by means of granular time fixed effects.

governing price rigidity. Moreover, due to imperfect competition, the slope also depends on the strength of strategic complementarities in firms' price-setting behavior. The presence of strategic complementarities implies that a relative price increase lowers a firm's desired markup, dampening the response of prices to movements in marginal cost.

Section 3 provides an overview of our data. Similarly to Amiti et al. (2019), we use data on the Belgian manufacturing sector. By merging data from different administrative sources, we collect information on firm-product-level output prices, quantities, and production costs. Notably, our the data is recorded at a quarterly frequency and the sample period spans the two decades (1999–2019) of market interactions between both domestic and foreign competitors. This dataset enables us to observe firms' own domestic prices, their competitors' prices, as well as a measure their marginal costs, all of which directly map to the theoretical objects.

In Section 4, we outline our identification strategy and clarify the benefits of using firm-level data to overcome the different identification challenges described above. While the conventional approach involves aggregating individual firm pricing decisions into the NKPC and then estimating the slope with aggregate data, we do the reverse. Leveraging our microdata, we estimate pricing relationships at the individual firm level to identify the structural parameters that determine the extent of price rigidity and the significance of strategic price complementarities. With estimates of these key parameters, we then aggregate to derive the implied slope of the primitive form of the Phillips curve.

We present the estimation results in Section 5. The estimates of the degrees of price rigidity and strategic complementarities are both sensible and robust. Our analysis suggests a substantial degree of price stickiness, implying that prices are fixed for about three to four quarters, on average. This figure closely aligns with the frequency of producer price adjustments reported by Nakamura and Steinsson (2008) in the US and almost exactly matches the frequency of price adjustments in the micro-data used to construct the official PPI series for Belgium. Our estimates also indicate a substantial degree of strategic complementarity in firms' price-setting behavior, which is broadly consistent with the empirical evidence in Amiti et al. (2019). The implied estimate of the slope of the marginal cost-based Phillips curve is in the range of 0.05 to 0.06. These numbers are an order of magnitude larger than the estimates of output gap-based or unemployment gap-based Phillips curve slopes available in the literature. Furthermore,

we capitalize on the granularity of our data and allow for sector-level heterogeneity in price-setting behavior in our empirical models to estimate sectoral Phillips curves. Our findings reveal a considerable variation in both the degree of nominal rigidity and strategic complementarity across sectors, which translate into heterogeneous pass-through of marginal costs to inflation.

Section 6 presents various exercises to assess the robustness of our findings. First, we demonstrate the resilience of our results to using alternative identification strategies that rely on variation in marginal costs driven by both aggregate demand and supply shocks. For the former, we use high-frequency shocks to monetary policy, while for the latter, we use high-frequency oil shocks. Second, we show that our estimates remain largely unchanged even when accounting for empirically plausible degrees of macroeconomic complementaries. Finally, we show that under reasonable conditions our identification approach and the resulting estimates of the NKPC slope remain valid regardless of whether we employ a Calvo framework or a conventional menu cost framework to model nominal rigidities.

In Section 7, we examine the implications of our model for aggregate inflation dynamics. As a direct test of our empirical approach, we demonstrate that the marginal cost-based model is capable of capturing the swings in year-over-year inflation for the aggregate manufacturing sector. We then show that the sectoral Phillips curves also perform well in capturing the year-over-year variations in their respective sectoral inflation rates.

In Section 8, we reconcile our estimates with those found in the literature. We first formalize the mapping between the marginal cost-based and output-based curves, which is mediated by the elasticity of marginal cost to output. Subsequently, we estimate the slope of the output-based NKPC using firm-level data. In line with earlier estimates (e.g., Rotemberg and Woodford 1997), we find that the output-based slope is low. We are able to identify the source as a low elasticity of marginal cost to output.⁵ Finally, we examine the transmission of oil shocks to illustrate the value of using the marginal cost-based Phillips curve for tracing the impact of supply shocks on inflation. Section 9 concludes.

⁵It should be noted that our estimates are based on pre-pandemic data. It could be that during the recent inflation surge, the output elasticity of marginal cost has risen, leading to a steepening of the output-based Phillips curve. If so, our decomposition would help decipher this phenomenon.

2 Theoretical framework

This section presents the theoretical framework that underlies our empirical analysis. We formulate the minimum structure required to produce firm pricing equations that allow us to identify the slope of the aggregate Phillips curve. The framework features heterogeneous firms competing under imperfect competition subject to nominal rigidity. Firms are granular. They internalize their impact on industry aggregates and are influenced by the pricing decisions of their competitors. This model generates a micro-founded New Keynesian Phillips curve, the slope of which is a function of the structural parameters that govern firms' pricing behavior.

2.1 Preferences and pricing behavior

Each period *t*, the economy is populated by heterogeneous producers (or firms), denoted by *f*, each operating in an industry $i \in I = [0, 1]$. We denote by \mathcal{F}_i the set of producers competing in industry *i*. While each firm is measure zero relative to the economy as a whole and hence takes aggregate expenditure as given, it might be large relative to its industry, and hence internalizes the effect of its pricing decisions on the consumption and price index of the industry.

Let P_{ft} be the price charged by each firm for a unit of its output, P_{it} the industry price index, $\varphi_{it} = {\varphi_{ft}}_{f \in \mathcal{F}_i}$ a vector of firm-specific demand shifters, and Y_{it} the real industry output. For any industry *i*, we consider an arbitrary invertible demand system that generates a residual demand function of the following form:⁶

$$\mathcal{D}_{ft} := d(P_{ft}, P_{it}, \varphi_{it}) Y_{it} \quad \forall f \in \mathcal{F}_i.$$
⁽²⁾

We assume firms face nominal rigidities as in Calvo (1983).⁷ Each period firms face a probability $1 - \theta$ of being able to change their price, independent across time and across firms, with $\theta \in [0, 1]$. Thus the price P_{ft} paid by consumers in any given period is either the (optimal) reset price set by a firm that is able to adjust, which we denote by P_{ft}^o , or the price charged in the previous period, P_{ft-1} .

⁶The focus on invertible demand systems is a mild technical requirement that rules out the case where firms produce goods that are perfect substitutes but encompasses any system featuring an arbitrary (but finite) elasticity of substitution.

⁷In Section 6 we show that our identification approach and the estimated NPKC remain valid if the data-generating process features Ss-style price adjustments as in conventional menu cost models.

The firms are to adjust their prices during the period do so with the objective of maximizing expected profits. Their pricing decisions consider both the pricing choices made by competitors and the impact of their own price adjustments on their residual demand and the industry-wide price index. Additionally, nominal rigidities generate forward-looking pricing behavior, as firms take into account that it might not be possible to adjust prices every period. As a result, the optimal reset price set by firms that are able to adjust is a weighted average of current and (expected) future nominal marginal costs and markups. Let $\Lambda_{t,\tau} := \beta^{\tau} \frac{U_{c,t+\tau+1}}{U_{c,t+\tau}}$ denote the stochastic discount factor, $TC_{ft} := TC(\mathcal{D}_{ft})$ the real total costs, and MC_{ft}^n the nominal marginal cost of firm f. Then the optimal reset price P_{ft}^o solves the following profit maximization problem:

$$\max_{P_{ft}^{o}, \{Y_{ft+\tau}\}_{\tau\geq 0}} \mathbb{E}_{t} \left\{ \sum_{\tau=0}^{\infty} \theta^{\tau} \left[\Lambda_{t,\tau} \left(\frac{P_{ft}^{o}}{P_{it+\tau}} \mathcal{D}_{ft+\tau} - TC(\mathcal{D}_{ft+\tau}) \right) \right] \right\},$$

subject to the sequence of expected demand functions $\{\mathcal{D}_{ft+\tau}\}_{\tau\geq 0}$ in (2). The FOC of the problem is:

$$\mathbb{E}_t \left\{ \sum_{\tau=0}^{\infty} \theta^{\tau} \Lambda_{t,\tau} D_{ft+\tau} \left[\frac{P_{ft}^o}{P_{it+\tau}} - (1+\mu_{ft+\tau}) \frac{MC_{ft+\tau}^n}{P_{it+\tau}} \right] \right\} = 0,$$
(3)

where μ_{ft} denotes the log markup. Let $\epsilon_{ft+\tau} := -\frac{\partial \ln \mathcal{D}ft+\tau}{\partial \ln P_{ft}^o}$ be the residual elasticity of demand faced by firm f. The desired net markup is then given by the Lerner index:

$$\mu_{ft+\tau} := \ln\left(\frac{\epsilon_{ft+\tau}}{\epsilon_{ft+\tau} - 1}\right). \tag{4}$$

According to equation (3), the optimal reset price depends on the expected path of marginal cost over the period the firm expects its price to be fixed, where θ^{τ} is the probability the firm expects its price to be fixed τ periods from now. Moreover, in finding the optimal reset price, the firm factors in how its pricing decision today affects the expected path of the desired markups.

2.2 Technology

Firms are heterogeneous in their production technologies. We assume that a unit of output of Y_{ft} is produced at a nominal marginal cost of:⁸

$$MC_{ft}^{n} = C_{it}\mathcal{A}_{ft}Y_{ft}^{\nu},\tag{5}$$

where C_{it} denotes the nominal unit cost of the composite input factor (e.g., wages and intermediate goods) that is independent of the scale of production; \mathcal{A}_{ft} is a firm-specific cost shifter that captures, among other idiosyncratic factors, heterogeneity in firm's production efficiency; ν is a parameter that pins down the short-term returns to scale of firms production technology, which are given by $(1/(1 + \nu))$.

In our benchmark model, we focus on the constant returns to scale case (v = 0). This assumption rules out macroeconomic complementarities due to the feedback of firms' pricing behavior into their respective marginal cost (see e.g. Galí 2015).⁹ In Appendix A.2 we present a general framework that allows for arbitrary aggregate returns to scale. In Section 6.2, we show that our estimates of the Phillips curve are robust as the empirical evidence is broadly consistent with the constant returns to scale assumption at both the sectoral and aggregate levels.

2.3 The optimal reset price

We log-linearize the FOC in equation (3) around the symmetric steady state with zero inflation.¹⁰ Denoting with lower-case letters the variables in logs, we obtain that the reset price satisfies:

$$p_{ft}^{o} = (1 - \beta\theta)\mathbb{E}_t \left\{ \sum_{\tau=0}^{\infty} (\beta\theta)^{\tau} \left(\mu_{ft+\tau} + mc_{ft+\tau}^n \right) \right\}.$$
(6)

As we show in Appendix A.1, the log-linearized desired markup is a function that depends inversely on the gap between the firms' own reset price and the price of its

⁸This functional form is rather general and consistent with standard production technologies used in the literature (see e.g. AIK and Hottman et al. 2016). As we show Appendix B.4, it nests Cobb-Douglas and CES as special cases.

⁹Macroeconomic complementarities can arise, for example, from roundabout production as in Basu (1995) or local input markets as in Woodford (2011).

¹⁰The choice of steady-state inflation is largely immaterial for our purposes but permits a lighter notation. We relax it in the empirical analysis, where we allow for both sector and industry-specific trends.

competitors, which we denote by p_{it}^{-f} . Formally:

$$\mu_{ft} - \mu = -\Gamma \left(p_{ft}^o - p_{it}^{-f} \right) + u_{ft}^{\mu}, \tag{7}$$

where $\Gamma > 0$ denotes the markup elasticity with respect to prices and u_{ft}^{μ} is a shock to the desired markup. Under weak assumptions, this relationship holds for standard imperfectly competitive frameworks, including static oligopoly (Atkeson and Burstein 2008), dynamic oligopoly (Wang and Werning 2022), and monopolistic competition with variable elasticity of demand (Kimball 1995). These frameworks share the property that, in equilibrium, a firm's elasticity of demand declines as its market share increases. Thus the presence of strategic complementarities in price-setting behavior implies that a relative price increase lowers a firm's desired markup, dampening the response of prices to movements in marginal cost.

Substituting the expression for $\mu_{ft+\tau}$ in the log-linearized FOC we obtain the following forward-looking pricing equation:

$$p_{ft}^{o} = (1 - \beta \theta) \mathbb{E}_t \left\{ \sum_{\tau=0}^{\infty} (\beta \theta)^{\tau} \left((1 - \Omega) (mc_{ft+\tau}^n + \mu) + \Omega p_{it+\tau}^{-f} \right) \right\} + u_{ft}, \tag{8}$$

where u_{ft} is a firm-specific *cost-push* shock.¹¹ The parameter $\Omega := \frac{\Gamma}{1+\Gamma}$ captures the strength of strategic complementarities and enters equation (8) in a way that reduces the responsiveness of prices to changes in marginal costs. If the elasticity of demand is constant—as it is in the textbook New Keynesian model with monopolistically competitive firms—so is the desired markup μ_{ft} . In this case, $\Omega = 0$ and the optimal pricing equation simplifies to the familiar formulation where the reset price is solely dependent on the current and future stream of marginal costs. Competitors' prices are then irrelevant.

2.4 The New Keynesian Phillips Curve

As we show in Appendix A.2, the log-linear aggregate price index is:

$$p_t = (1 - \theta)p_t^0 + \theta p_{t-1},\tag{9}$$

¹¹For some specifications of the demand system, the residual term u_{ft} depends on the vector of expected demand shifters ($\mathbb{E}_t \{\varphi_{it+\tau}\}$) and, in the case of non-stationary dynamic oligopoly, on the expected slope of the competitors' reaction function ($\mathbb{E}_t \{\partial p_{it+\tau}^{-f}/\partial p_{ft}^o\}$), which can vary over time depending on aggregate states. (See, e.g., the CES demand system in Appendix A.1.) This could potentially generate a correlation between p_{it}^{-f} and the residual term u_{ft} . Both possibilities are addressed in our baseline analysis, as well as in robustness tests, by using an instrumental variable approach.

with p_t and p_t^o denoting the aggregate price indexes implied by (2), which average across firms and industries. Let mc_t^n denote the aggregate log-nominal marginal cost. Define the aggregate real marginal cost and aggregate inflation as $mc_t^r = mc_t^n - p_t$ and $\pi_t = p_t - p_{t-1}$, respectively. Averaging the pricing equation in (8) across firms and industries and writing it in recursive form, we obtain an equation for the aggregate reset price:

$$p_t^o = (1 - \beta\theta) \left((1 - \Omega) (mc_t^n + \mu) + \Omega p_t \right) + \beta\theta \mathbb{E}_t p_{t+1}^o + \frac{\theta}{1 - \theta} u_t, \tag{10}$$

where u_t is an aggregate cost-push shock, defined in the Appendix. Combining equations (9) and (10) gives the primitive form of the NKPC curve:

$$\pi_t = \lambda \ \widehat{mc_t^r} + \beta \mathbb{E}_t \{ \pi_{t+1} \} + u_t, \tag{11}$$

which asserts that inflation depends on real marginal cost (in deviation from its steady state) and on expected future inflation. λ is the slope of the NKPC curve, defined by:

$$\lambda := \frac{(1-\theta)(1-\beta\theta)}{\theta}(1-\Omega).$$
(12)

Two observations are worth noting. First, the primitive form of the Phillips curve in equation (11) uses the log deviation of real marginal cost from its steady state as the relevant real activity variable. In contrast, the conventional formulation of the Phillips curve, displayed in equation (1), uses the output gap or unemployment to proxy for marginal cost. As we will discuss, the mapping between marginal cost and output gap is theoretically valid only under specific circumstances. Moreover, even when a proportionality between the two variables can be established, the elasticities of marginal cost to output gap and unemployment need not be equal to one. We return to these important points in Section 8.

Secondly, the slope of the NKPC, λ , is a function of the primitives governing firms' pricing behavior. As in standard New Keynesian models (e.g., Galí and Gertler 1999), high nominal rigidities and low discounting flatten the sensitivity of inflation to changes in real economic activity. Additionally, equation (12) shows how strategic complementarities also contribute to reducing the slope. Therefore, given a calibration of the discount factor β , estimates of the structural parameters θ and Ω pin down the slope of the Phillips curve.

Toward this end, we take the structural pricing equation (8) to the data. This exercise requires measures of prices and marginal costs, which we discuss in the next section. Notably, it is the use of firm-level data that permits the identification of the primitive

parameters.

3 Data and Measurement

We begin by introducing our dataset and highlighting its features that are relevant for measurement purposes. We then illustrate the procedure for constructing price and marginal cost measures using both product-level and firm-level data.

3.1 Data

We assemble a unique micro-level dataset that covers the manufacturing sector in Belgium between 1999 and 2019. A rare feature of our dataset is its ability to track, on a *quarterly basis*, manufacturing product-level prices and quantities sold in the domestic market by both domestic and foreign producers, as well as information on production costs for domestic producers. Our dataset is compiled from four administrative sources: PRODCOM, international trade data, VAT declarations, and Social Security declarations.

We obtain information on domestic firms from PRODCOM. This dataset allows us to observe firms' quarterly sales and physical quantities sold for each narrowly defined 8-digit manufacturing product. We use this highly disaggregated information to calculate domestic unit values (sales over quantities) at the firm-product level (PC 8 digit).¹² We obtain similar data on foreign competitors from the administrative records of Belgian Customs. Specifically, for each manufacturing product sold by a foreign producer to a Belgian buyer, we observe quarterly sales and quantity sold for different products (CN 8-digit), from which we compute unit values of foreign competitors in local markets.

We leverage detailed administrative data to measure firms' variable production costs at a quarterly frequency. Specifically, we obtain information on firms' purchases of intermediates (materials and services) from their VAT declarations. Additionally, we draw upon firms' Social Security declarations to obtain a measure of their labor costs (the wage bill).

¹²PRODCOM surveys all Belgian firms involved in manufacturing production with more than 10 employees, covering over 90% of production in each NACE 4-digit industry. The survey does not require firms to distinguish between production and sales to domestic and international customers. Therefore, we follow Amiti et al. (2019), henceforth AIK, and recover domestic values and quantities sold by combining information from PRODCOM with international trade data on firms' product-level exports (quantities and sales).

After applying standard data cleaning filters, our final sample includes 4, 499 firms observed over 84 quarters (1999:Q1–2019:Q4), totaling 129, 425 observations. Appendix B provides detailed information on the data sources and data cleaning procedures. Table 1 presents summary statistics of our dataset. Several features of the data are worth noting.

First, our dataset covers the lion's share of domestic manufacturing production in Belgium, spanning the entire size distribution. The average firm in our dataset employs 74 employees (measured in full-time equivalents) and has a domestic turnover (sales) of ϵ 6 million. The sales of the smallest firms in the sample are worth less than one-tenth of a thousandth of those generated by the largest producers.

Second, thought the paper we adopt a narrow industry definition based on 4-digit NACE rev. 2 codes, the standard sector classification system in the European Union. Based on this classification, we sort firms into 169 manufacturing industries, distributed across 9 manufacturing sectors.¹³ This classification optimally balances a coherent definition of the industry (which is mostly precise if narrow) with the ability to identify an appropriate set of competitors (both domestic and foreign) competing to gain market share in Belgium. Table 1 shows that the lion's share of the firms in our sample specializes in only one manufacturing industry. Even for those firms that operate in multiple industries, the contribution of the main industry to total firm revenues is, on average, 98% (median 100%). For the few multi-industry firms, in line with the theoretical framework, we treat each industry as a separate firm.

Third, the typical sector is characterized by a large number of firms with small market shares—the average within-industry share is approximately 1.5% on average, with a median of 0.5%—and a few relatively large producers. To the extent that these large firms internalize the effect of their pricing and production decisions on industry aggregates and strategically react to the pricing decisions of their competitors, the monopolistic competition benchmark would be a poor approximation. The theoretical framework introduced in the previous section explicitly accounts for this.

Fourth, although the largest firms have nontrivial market shares in their industries, they are small compared to the volume of economic activity of their macro sector

¹³The first four digits of the PRODCOM product classification coincide with the first four digits of the NACE rev. 2 classification and also to the first 4 digits of the CN product code classification used in the customs data. Following the official Eurostat classification system, we define manufacturing sectors by grouping 2-digit NACE rev. 2 codes. See Appendix B sectors' definitions.

	Mean	5 st pctle	25 th pctle	Median	75 th pctle	95 th pctle
Firm Employees	75.62	9.00	17.42	30.50	59.00	276.50
Firm Sales	6663.15	228.13	666.11	1486.44	3870.36	21946.60
Number of industries within firm	1.10	1.00	1.00	1.00	1.00	2.00
Within firm revenue share of main industry	98.23	86.80	100.00	100.00	100.00	100.00
Firm's market share within industry	1.54	0.06	0.22	0.52	1.30	6.02
Firm's market share within sector	0.21	0.01	0.02	0.05	0.14	0.70
Firm's market share within manufacturing	0.02	0.00	0.00	0.01	0.01	0.08
Number of consecutive quarters in sample	42.03	10.00	24.00	38.00	58.00	82.00

Table 1: Summary Statistics

Notes. The summary statistics reported in this table refer to the sample of domestic producers in PRODCOM. The firm's employees are measured in full-time equivalents. Firm sales are measured in thousands of Euros, rounded to the nearest integer. Within firm revenues shares and the firm's market shares are measured in percentages.

(e.g., textile manufacturing or electrical equipment manufacturing) and, even more so, compared to the volume of economic activity in the whole manufacturing sector in Belgium. It is therefore reasonable to assume that even the largest producers do not internalize the effect of their pricing and production decisions on the aggregate economy.

Finally, our data allow us to observe long time series of both prices and marginal costs. On average, we observe firms for approximately 10 consecutive years (42 quarters). This feature of the data is particularly important for identification purposes. As we discuss below, a long time series enables us to include unit fixed effects in our empirical models to control for time-invariant confounding factors without suffering from the classical Nickell bias that frequently complicates the estimation of dynamic panel models.

3.2 Measurement

We now describe the measures of prices and marginal cost that map to the theoretical counterparts in Section 2. Our measurement approach is guided by the data features

described above and builds on AIK, which uses a dataset similar to ours to study the pass-through of exchange rate shocks to prices. Appendix B provides a detailed description of the procedure used to construct all our variables.

Output prices. The key variable of interest is the domestic price of goods charged by firms in the local market (Belgium). Consistent with the notation used in the theoretical framework, we use the subscript *i* to denote an industry, *f* to denote a firm-industry pair, and *t* to denote time (quarters).¹⁴ We denote by s_{ft} the revenue share of the firm in the industry.

As in AIK, we compute the change in firm prices P_{ft}/P_{ft-1} , using the most disaggregated level allowed by the data. For domestic producers, the finest level of aggregation is a firm×PC 8-digit product code level. For foreign competitors, it is the importing-firm×source country×CN 8-digit product code level. Approximately half of the domestic firms in our sample are multi-product firms, meaning they produce multiple 8-digit products within the same industry. For these entities, we compute the price change by aggregating changes in product-level prices using a Törnqvist index:¹⁵

$$P_{ft}/P_{ft-1} = \prod_{p \in \mathcal{P}_{ft}} (P_{pt}/P_{pt-1})^{\bar{s}_{pt}}$$

In the formula above, \mathcal{P}_{ft} represents the set of 8-digit products manufactured by firm f, P_{pt} is the unit value of product p in \mathcal{P}_{ft} , and \bar{s}_{pt} is a Törnqvist weight computed as the average of the sale shares between t and t - 1: $\bar{s}_{pt} := \frac{s_{pt} + s_{pt-1}}{2}$. We then construct the time series of firms' prices (in levels) by concatenating quarterly changes.¹⁶

$$P_{ft} = P_{ft_f^0} \prod_{\tau = t_f^0 + 1}^t \left(\frac{P_{f\tau}}{P_{f\tau-1}} \right)$$

The normalization of the level of the firm's price index in the base year, $P_{f t_e^0}$, is one rationale for the inclusion

¹⁴Whenever a firm operates in multiple 4-digit industries, we treat each firm-industry pair as a separate unit in our sample. As we discussed, most firms operate in only one industry, and the main industry accounts for the lion's share of sales of multi-industry firms. Therefore, all our results are essentially unchanged if we restrict the sample to the main industry for each firm.

¹⁵The Törnqvist index coincides with the Cobb-Douglas index whenever sales shares are constant over time. Allowing for time-varying shares is empirically relevant given that market shares of individual firms vary over time due to changes in market conditions, entry and exit of firms, and other factors. The Törnqvist index also tends to be robust to measurement error and aggregation bias. As we discuss in Section 6, our results are nevertheless robust to the choice of different weights.

¹⁶ Let t_f^0 denote the first quarter when f appears in our data. Starting from a base period P_{ft^0} , which we can normalize to one, prices are concatenated using the formula:



Figure 1: Aggregate inflation: Official statistic vs. in-sample replication

Notes. This figure reports the official statistics of PPI year-over-year inflation for the domestic manufacturing sector in Belgium (gray dashed line) and the same measure constructed from micro data (solid black line). The latter is a Törnqvist index that aggregates across domestic products, firms, and industries in our final sample.

Using a similar approach, we construct the price index of competitors for each domestic firm by concatenating quarterly changes according to the following formula:

$$P_{it}^{-f}/P_{it-1}^{-f} = \prod_{k \in \mathcal{F}_i/f} (P_{kt}/P_{kt-1})^{\bar{s}_{kt}^{-f}}.$$

Here, $\bar{s}_{kt}^{-f} := \frac{1}{2} \left(\frac{s_{kt}}{1-s_{ft}} + \frac{s_{kt-1}}{1-s_{ft-1}} \right)$ represents a Törnqvist weight, which is constructed by averaging the residual revenue share of competitors in the industry at time t (net of firm f revenues) with that at time t - 1.¹⁷ It is important to note that the set of domestic competitors for each Belgian producer, denoted as \mathcal{F}_i , includes not only other Belgian manufacturers operating in the same industry but also foreign manufacturers that goods to Belgian customers.

Lastly, we construct an aggregate price index and a corresponding aggregate inflation series by aggregating product-level price changes across domestic products, firms, and industries. The correlation between our price index and the official

of firm fixed effects in our empirical specifications.

¹⁷As with the firm's price index, the level of the price index of competitor is constructed by normalizing the first period to one and concatenating quarterly changes. Also, in this case, the normalization is immaterial for estimation purposes as our empirical model always includes firm fixed effects.

manufacturing PPI is 0.97. Our inflation measure also aligns closely with the official statistic, as demonstrated in Figure 1. These aggregate statistics confirm the representativeness of our sample and validate our empirical procedure for constructing price indexes.

Marginal costs. The general cost structure outlined in equation (5) implies that firms' nominal marginal costs are proportional to its average variable costs, as follows:

$$MC_{ft}^{n} = (1+\nu)\frac{TVC_{ft}}{Y_{ft}}.$$
(13)

Taking logs of equation (13), we obtain a proxy of firms' marginal costs that has a measurable counterpart in our data. Various inputs are utilized simultaneously in the production of multiple goods, making the variation in total variable production costs (across the different product lines) the most suitable cost index for firms' pricing choices. Accordingly, we total variable costs (TVC_{ft}), as the sum of intermediate costs (materials and services purchased) and labor costs (wage bill), both measured at a quarterly frequency.

We obtain a firm-specific quantity index for domestic sales (Y_{ft}) by scaling a firm's domestic revenues by its domestic price index, such that $Y_{ft} = (PY)_{ft}/\bar{P}_{ft}$. For single-industry firms, \bar{P}_{ft} coincides with the firm-industry price index P_{ft} , which was discussed earlier. For multi-industry firms, we aggregate industry prices P_{ft} by using as weights the firm-specific revenue shares of each industry.¹⁸ Finally, returns to scale are not directly observable in the data. By applying a logarithmic transformation to equation (13), the inverse of the short-run returns to scale parameter, $\ln(1+\nu)$, enters as an additive term in our specifications. As we explain below, we control for this factor using different sets of fixed effects.

¹⁸Specifically, we apply the Törnqvist weight of each (4-digit) industry bundle *i* produced by firm *f* in quarter *t*, which is defined as $(s_{fit} + s_{fit-1})/2$, where s_{fit} is the share of revenues of the firm coming from sales in industry *i* in total sales across industries. The choice of \bar{P}_{ft} has essentially no impact on our estimation results because, as we have shown, the majority of the firms in our data operate in only one industry, and the sales of those firms that produce goods in multiple industries are typically concentrated mainly in their primary industry. All our results are robust to defining \bar{P}_{ft} as the price of the main industry or using other aggregation methods (such as an arithmetic average or a CES aggregator).

4 Identification strategy

In this section, we present the identification strategy that enables us to take the theoretical framework developed in Section 2 to the data. First, we show how to connect theoretical reset prices to observed prices to obtain forward-looking pricing equations that have measurable counterparts. Subsequently, we introduce different estimation methods to identify the structural parameters of the pricing equations. We conclude by discussing the identification challenges and how we tackle them.

4.1 Mapping the model to the data

Consider a firm entering period *t* before it learns whether or not it will be able to change its' price in *t*. Under the Calvo framework, the conditional expectation of the observed price given the reset price and the price in the previous period is given by $\mathbb{E}\{p_{ft}|p_{ft}^o, p_{ft-1}\} = (1-\theta)p_{ft}^o + \theta p_{ft-1}$. Define the error term $v_{ft} := p_{ft} - \mathbb{E}\{p_{ft}|p_{ft}^o, p_{ft-1}\}$ such that $\mathbb{E}\{v_{ft}\} = 0$, which captures ex-ante uncertainty regarding the outcome of the Calvo draw. Then:

$$p_{ft} = (1 - \theta) p_{ft}^{o} + \theta p_{ft-1} + v_{ft}.$$
 (14)

Leveraging the rational expectations assumption, we can use equation (8) to express the reset price p_{ft}^o in terms of (future) observable variables and replace it in the equation above to obtain to following population regression (up to an additive constant):

$$p_{ft} = (1-\theta)(1-\beta\theta) \sum_{\tau=0}^{\infty} (\beta\theta)^{\tau} \left((1-\Omega)mc_{ft+\tau}^n + \Omega p_{it+\tau}^{-f} \right) + \theta p_{ft-1} + \varepsilon_{ft},$$
(15)

where ε_{ft} is the composite error term, $\varepsilon_{ft} := v_{ft} + (1 - \theta)u_{ft} + (1 - \theta)(1 - \beta\theta)e_{ft}$, u_{ft} is the firm-specific cost-push shock, and e_{ft} is an expectational error such that $\mathbb{E}_t(e_{ft}) = 0$.

With microdata on prices and marginal costs, we take equation (15) to the data to identify the parameters that pin down the slope of the Phillips curve. Intuitively, averaging across firms that can and cannot adjust during a particular period, the elasticity of current prices to lagged prices identifies θ , the degree of price stickiness. Given θ and a calibration of the discount factor β , the complementarity parameter Ω can be identified either by variation in the expected present value of marginal costs or by variation in the expected present value of the competitors' price index, or both.

4.2 Empirical specifications

We present four alternative models to estimate the population equation (15). Individually, the different models differ in terms of the sources of variation used to identify the parameters, as well as in the assumptions made regarding the dynamics of marginal costs and prices. Together, they produce alternative estimates that allow us to assess the robustness of our identification framework.

Model A. Our first empirical model is the closest sample analog of the population regression in equation (15). Because the weights in front of the leads of marginal cost and competitors' prices decay at a rate $\beta\theta < 1$, we can approximate their present values with a truncated series as follows:

$$\sum_{\tau=0}^{\infty} (\beta\theta)^{\tau} m c_{ft+\tau}^{n} \approx \sum_{\tau=0}^{T-1} (\beta\theta)^{\tau} m c_{ft+\tau}^{n} + \frac{(\beta\theta)^{T}}{1-\beta\theta} m c_{ft+T}^{n} =: PV_{ft}^{T,mc},$$

$$\sum_{\tau=0}^{\infty} (\beta\theta)^{\tau} p_{it+\tau}^{-f} \approx \sum_{\tau=0}^{T-1} (\beta\theta)^{\tau} p_{it+\tau}^{-f} + \frac{(\beta\theta)^{T}}{1-\beta\theta} p_{it+T}^{-f} =: PV_{ft}^{T,p}.$$
(16)

In our baseline specification we truncate both series at T = 8 (2 years), but the results are robust to a choice of longer horizons.¹⁹ Replacing these expressions in (15) and calibrating β to the conventional value of 0.99, we obtain our first empirical model:

$$p_{ft} = (1 - \theta)(1 - \beta\theta) \left((1 - \Omega)PV_{ft}^{T,mc} + \Omega PV_{ft}^{T,p} \right) + \theta p_{ft-1} + \alpha_f + \alpha_{s \times t} + \varepsilon_{ft}.$$
(Model A)

Note that we augmented our empirical model with a battery of firm fixed effects (α_f) and sector-by-quarter fixed effects ($\alpha_{s\times t}$). In the data, we can identify nominal variables (both the price level and marginal cost) only up to a normalization.²⁰ The inclusion of firm fixed effects makes this normalization immaterial by estimating the parameters of using only within-firm variation.

The sector-by-quarter fixed effects serve several purposes. First, to keep the

¹⁹Assuming that mc^n and p^{-f} consist of permanent and transitory components and that the transitory components follow an AR(1) process with a persistence of $\rho < \frac{1}{\beta\theta}$, the truncation period *T* can be set such that $(\rho\beta\theta)^T \approx 0$. Thanks to the additional discounting by θ , the quantity $(\rho\beta\theta)^8$, evaluated at our benchmark estimates, is approximately one percent.

²⁰See footnote 16.

notation light, we derived our equilibrium pricing equations by log-linearizing around a steady state with constant inflation. The inclusion of sector-by-quarter fixed effects makes our empirical models robust to the possibility of sector-specific trends in inflation and marginal costs, as the parameters of interest are identified using variation across firms that operate in the same sector at the same time. Second, shifts in aggregate or sectoral demand could also generate spurious co-movements of marginal costs and prices through general equilibrium forces. This would be the case, for example, if unobservable demand shocks affect wages and intermediate prices at either the aggregate or sectoral level. The fixed effects help to alleviate such potential concerns. Also, the inclusion of sector-by-quarter fixed effects allows us to account for changes in long-run expectations about inflation. As shown by Hazell et al. (2022), long-run expectations tend to positively co-move with the level of economic activity and therefore pose a threat to the identification of the structural parameters of interest.

Finally, in combination, the different layers of fixed effects account for the possibility of unobservable heterogeneity in production technologies at the sector level or firm level (e.g., heterogeneous returns to scale), which influence marginal costs but are not captured by our empirical proxy.

Model B. Model A uses information on both marginal costs and industry prices to identify the complementarity coefficient Ω . However, Ω can in principle be identified by either variation in either one of the two. Building on this insight, our second empirical model estimates the importance of strategic complementarities using only variation in marginal costs. Since most firms are small relative to their industry size, the variation in each firm's competitors' price index occurs primarily at the industry-year level (over 90%). As a result, we can use industry-by-quarter fixed effects ($\alpha_{i\times t}$) to absorb the present value of competitors' prices and estimate the following empirical model:

$$p_{ft} = (1 - \theta)(1 - \beta\theta)(1 - \Omega)PV_{ft}^{T,mc} + \theta p_{ft-1} + \alpha_f + \alpha_{i \times t} + \varepsilon_{ft}.$$
 (Model B)

Model B has a few advantages over Model A. First, it includes narrow industry-by-time fixed effects that better address concerns related to correlated demand shocks that could affect both marginal costs and prices of all producers in the industry. Second, it alleviates concerns related to the definition of the set of competitors for each firm (and therefore the definition and exact functional form of the competitors' price index), which could in principle be broader than the set of producers operating in the same industry.

Model C. A valuable feature of both Model A and Model B is that neither imposes stringent constraints on how firms form expectations about the dynamics of future marginal costs and industry prices. The flip side of this flexibility is that the estimating equations are highly nonlinear due to the presence of θ in both the present values and the coefficient in front of the lagged price, which might be demanding on the data. We thus explore an alternative approach that imposes additional structure on the dynamics of the right-hand side variables but allows us to derive more parsimonious and linear empirical specifications. Specifically, we assume that industry prices and marginal costs, in deviations from industry trends, follow an auto-regressive process of order one plus drift, with persistence parameters ρ^p and ρ^{mc} that are both less than $\frac{1}{\beta\theta}$. We then estimate the following system of linear equations, which represents our third empirical specification:

$$p_{ft} = (1 - \theta) \left((1 - \Omega) \frac{1 - \beta \theta}{1 - \beta \theta \rho^{mc}} mc_{ft}^n + \Omega \frac{1 - \beta \theta}{1 - \beta \theta \rho^p} p_{it}^{-f} \right) + \theta p_{ft-1} + \alpha_f + \alpha_{s \times t} + \varepsilon_{ft}$$

$$p_{it}^{-f} = \rho^p p_{it-1}^{-f} + \alpha_f + \alpha_{s \times t} + \varepsilon_{ft}^p$$
(Model C)
$$mc_{ft}^n = \rho^{mc} mc_{ft-1}^n + \alpha_f + \alpha_{s \times t} + \varepsilon_{ft}^m,$$

where the auto-correlation parameters ρ^{p} , ρ^{mc} are now jointly identified with the stickiness and complementary parameters.

Model D. Finally, as in our second empirical model, we can estimate an alternative version of Model C, by replacing the set of sector-by-quarter fixed effects with the narrower set of industry-by-quarter fixed effects and identify the complementarity parameter using only variation in marginal costs. This leads us to our fourth empirical model:

$$p_{ft} = (1 - \theta)(1 - \Omega) \frac{1 - \beta \theta}{1 - \beta \theta \rho^{mc}} mc_{ft}^n + \theta p_{ft-1} + \alpha_f + \alpha_{i \times t} + \varepsilon_{ft},$$

$$mc_{ft}^n = \rho^{mc} mc_{ft-1}^n + \alpha_f + \alpha_{i \times t} + \varepsilon_{ft}^{mc}.$$
(Model D)

4.3 Identification issues

The identification of the parameters of interest requires us to tackle econometric challenges arising from endogeneity of prices and imperfect measurement of marginal costs. First, because firms (that are able to adjust in the period) set their prices simultaneously, a firm's own price and its competitors' prices are both a function of unobservable demand shocks. Second, even if our data allows us to directly observe the different components of marginal costs, our measure is inherently noisy, possibly leading to attenuation bias in our estimates.

To address these issues, we estimate our empirical models via Generalized Method of Moments (GMM) by imposing a set of orthogonality conditions of the form:

$$\mathbb{E}\{\boldsymbol{z}_{ft}\cdot\boldsymbol{\varepsilon}_{ft}\}=0,\tag{17}$$

where z_{ft} is a vector of instruments. From the definition of ε_{ft} , the orthogonality condition is satisfied by any variable in the information set of the firm prior to resetting its price, at the beginning of time t. For the pricing equations in Model A, we include in the baseline instrument set a unit vector and lagged realizations of marginal costs and the price index of competitors, both of which are measured sufficiently in the past (8 quarters) to ensure orthogonality with respect to recent shocks.²¹ It is important to note that to identify the parameters of interest we can use variation in marginal costs driven by demand and supply simultaneously.²² This consideration makes lagged variables suitable instruments so long as they satisfy the orthogonality conditions. We show that this is indeed the case. We also show in Section 6.1 that our results are robust to a set of alternative instruments that reflect alternatively aggregate demand and supply shocks.

In addition, to estimate Model C we augment the set of instruments with one-quarter lags of the endogenous variables to estimate the autoregressive coefficients in the auxiliary equations. For models Model B and Model D, we only need instruments for marginal costs

²¹The presence of persistent firm-specific shocks can in principle raise concerns regarding the identification of the stickiness parameter θ , the coefficient attached to lagged prices. If there exists a positive serial correlation in these shocks, it has the potential to bias θ upward and, therefore, lead us to underestimate the slope of the Phillips curve. The evidence in Section 5 indicates that this is not the case as our structural estimate of θ matches an external estimate of the frequency of price adjustment obtained from complementary PPI micro data. A plausible explanation is that the persistence of demand shocks originates from a persistent industry component or a time-invariant firm component, which are both absorbed by the fixed effects included in our regression models.

²²By contrast, as we discuss in Section 8, we do need exclusively demand-side instruments to identify an output gap-based Phillips curve.

as the price of competitors is omitted due to the collinearity with the industry-by-quarter fixed effects.

Weak instruments represent a common issue encountered when estimating the slope of the NKPC using conventional approaches that rely on aggregate time-series regressions (Mavroeidis et al. 2014). This issue is due to the lack of sufficient variation in aggregate instruments, which leads to low statistical power and thus imprecise and possibly biased estimates of the slope coefficient. Our identification approach does not suffer from this: the micro-level instruments are powerful. To demonstrate this point, we regress the endogenous variables—marginal costs and competitors' prices—on our instruments, essentially producing what would be the first-stage regressions of a linear two-stage least squares model. Table 2 presents the estimation results. As we can see, all coefficients have the expected signs and are statistically significant. The high values of the Cragg-Donald F-statistic indicate that we can reject the hypothesis of weak identification at any standard confidence level.²³ It is also noteworthy that the instrument for marginal costs has practically no predictive power on the competitors' price index, and vice versa. This suggests that Model A and Model C identify the complementarity parameter using essentially orthogonal variation in the two endogenous variables.

The GMM estimation follows a two-step procedure. In the first step, the identity matrix is used as a weighting matrix to estimate the coefficients of the model. In the second step, we use the estimated coefficients to compute the efficient weighting matrix employed in the calculation of the standard errors. All our models weight observations by the Törnqvist weights, \bar{s}_{ft} . This implies that the weight assigned to each firm is the same as the weight assigned to it in the construction of aggregate price indexes, thereby ensuring that our estimates are representative from a macroeconomic standpoint. We cluster standard errors at the sector level to account for the potential correlation structure of error terms across firms in similar businesses. This choice is conservative but appropriate since it takes into account the firms' simultaneous exposure to demand and supply shocks.

²³We reach the same conclusion when using the test developed by Olea and Pflueger (2013), which is robust to heteroskedasticity, serial correlation, and clustering.

	Mod	lel A	Model B	Мос	lel C	Model D
	$PV_{ft}^{T,mc}$	$PV_{ft}^{T,p}$	$PV_{ft}^{T,mc}$	mc _{ft}	p_{it}^{-f}	mc _{ft}
mc _{ft-8}	0.133	-0.011	0.130	0.325	-0.017	0.285
p_{it-8}^{-f}	-0.034 (0.021)	0.439	(0.010)	-0.078 (0.037)	0.547 (0.044)	(0.002)
p_{ft-1}	0.242 (0.052)	0.096 (0.028)	0.276 (0.031)	0.327 (0.033)	0.088 (0.021)	0.335 (0.022)
Firm FE Sector×Time FE	y y	y y	У	y y	y y	у
Industry×Time FE			У			У
Weak IV F-stat	1,40	2.86	2,736.45	5,51	3.99	8,222.10

 Table 2: Relevance of the instruments

Notes. This table presents the estimate of different linear regression models that project the instruments on the endogenous variables. The set of instruments varies across models depending on the specific set of endogenous variables, as described in the paper. In all regressions, observations are weighted using Törnqvist weights. Standard errors, reported in parenthesis, are clustered at the sector level. We calculate present values using equation (16) calibrating θ to the estimates reported in Table 3. The discount factor is calibrated to $\beta = 0.99$. All models are estimated using the complete sample (N = 129, 425), although the number of observations in each column varies depending on the lag and lead structure employed in the estimation.

5 Estimation results

5.1 Structural parameters

We take our four empirical models to the data to estimate the structural parameters that govern firms' price-setting behavior. The results are reported in Table 3. Our estimates suggest a substantial degree of price stickiness, with a value of θ of approximately 0.7 that is precisely estimated and roughly similar across the different specifications. This estimate implies that, on average, prices remain fixed for approximately three to four quarters.²⁴ This finding closely aligns with the frequency of producer price adjustments reported by Nakamura and Steinsson (2008) in the US.

Notably, we have the opportunity to validate our model-based estimates of price

²⁴Given the Calvo structure, the average duration of prices is given by $\frac{1}{1-\theta}$.

rigidity by comparing them with an external estimate obtained from the microdata used by the Belgian statistical office to calculate the official manufacturing PPI. This alternative data source enables us to directly observe domestic producers' prices (as opposed to unit values, as in PRODCOM) for a representative basket of manufacturing products. In the PPI microdata, the price stickiness (measured as one minus the average probability of adjusting product prices from one quarter to the other) is 0.68 (standard error 0.001), which almost exactly matches the estimates of our structural models.

Our estimates also indicate a substantial degree of strategic complementarity in firms' price-setting behavior. The estimate of Ω is about 0.55, tightly estimated and robust across the different specifications. This estimate is consistent with those obtained by AIK using similar Belgian data but a different estimation approach. Combining the estimates of θ and Ω we can recover an estimate of the elasticity of a firm's own price to changes in competitors' prices, which is given by $\frac{\partial p_{ft}^o}{\partial p_{..}^{-f}} = \Omega(1-\theta) \left(\frac{1-\beta\theta}{1-\beta\theta\rho^p}\right)$ and approximately equal to 0.14. Different from AIK, our setting allows the elasticity to depend on both nominal rigidities and forward-looking pricing behavior. In flexible-price environments, the pass-through of shocks to competitors' prices into a firm's own price (Ω) would be about the same magnitude as the pass-through of shocks to a firm's own marginal cost $(1 - \Omega)$, as in AIK. Nominal rigidities influence this effect in two ways. First, only some firms can reset their prices, captured by $(1 - \theta)$. Second, firms that adjust their prices do so anticipating the path of future marginal cost over the period their price is expected to remain fixed. They thus adjust the marginal cost by the factor $\frac{1-\beta\theta}{1-\beta\theta\rho^p}$, which is increasing in the persistence of marginal cost, ρ^p . Finally, we find that variation in markups absorbs the almost entirety of firms' response to competitors' actions. Our estimates imply that the elasticity of markups with respect to competitors' prices, $\frac{\partial \mu_{ft}}{\partial p_{it}^{-f}} = -\Gamma\left(\frac{\partial p_{ft}^o}{\partial p_{it}^{-f}} - 1\right)$, is approximately 1, which is again consistent with the unitary pass-though found by AIK.

Finally, we find that the persistence of marginal cost is between 0.63 for Model C and 0.78 for Model D, and the persistence of the price index is 0.93. It is important to recognize that these parameters reflect persistent deviations from trends, that have been absorbed by inclusion of the time-by-industry fixed effects.

	Model A	Model B	Model C	Model D
	Р	Panel a: Structu	ural estimates	
θ	0.704	0.690	0.727	0.679
	(0.019)	(0.016)	(0.014)	(0.013)
Ω	0.539	0.651	0.529	0.625
	(0.058)	(0.071)	(0.030)	(0.086)
ρ^{mc}	~ /	()	0.633	0.778
1			(0.047)	(0.055)
$ ho^p$			0.936	
,			(0.004)	
Firm FE	у	у	у	у
Sector×Time FE	y	·	y	
Industry×Time FE	-	У	-	У
Number of equations	1	1	3	2
Number of moments	4	3	8	5
Hansen J-test χ^2	2.130	0.019	6.976	3.791
-	Pan	el b: Slope of tl	he Phillips cur	ve
λ	0.059	0.050	0.050	0.058
	(0.012)	(0.016)	(0.005)	(0.018)

Table 3: Structural estimates and slope of the Phillips curve

Notes. This table presents the estimates of Models A, B, C, and D, which are different approximations of the population regression in equation (15). For each model, panel a reports the estimates and standard errors of the structural parameters. Panel b reports the slope of the Phillips curve (λ) implied by the estimated parameters according to equation (11). In the first and second columns (models A and B), we calculate present values using equation (16). In the third and fourth columns (models C and D) we relay on the AR1 assumption. The discount factor is calibrated to $\beta = 0.99$. All models are estimated using the complete sample (N = 129, 425), although the number of observations in each column varies depending on the lag and lead structure employed in the estimation. In all regressions, observations are weighted using Törnqvist weights. Robust standard errors (reported in parenthesis) are clustered at the sector level.

5.2 The slope of the Phillips curve

By combining the estimates of the structural parameters we can now recover the slope of the marginal-cost-based Phillips curve. Table 3 shows that, across the different specifications, the estimated values of λ range from 0.05 to 0.06. As we noted earlier, these estimates are much larger than estimates of the conventional NKPC found in the literature, which typically display a magnitude three to ten times smaller in comparison. Note also that if it weren't for the substantial dampening effect of strategic complementarities, the slope of the marginal cost-based NKPC would be even higher: given our estimates of Ω , strategic complementarities cut the slope in half. In Section 8, we discuss how to reconcile our estimates of the marginal cost-based NKPC with a conventional output gap-based one.

5.3 Sectoral Phillips curves

Exploiting the granular nature of our data, we can estimate our pricing equations separately for individual sectors to recover estimates of sectoral Phillips curves. We use Model C to obtain sector-specific estimates of the structural parameters and of the implied slope of the NKPC.²⁵ (The estimates obtained from the other specifications are very similar.) Table 4 presents the results.

For all sectors, the estimation indicates plausible degrees of price stickiness, with values of θ ranging between 0.58 for the transport equipment sector to 0.78 for the chemicals sector. The estimates of the degree of complementarities are also sensible and quite heterogeneous. Some sectors, such as machinery and transport equipment, display almost complete pass-through of marginal costs, whereas others, such as food and textiles, exhibit significant adjustments of desired markups in response to economic conditions and competitive pressure.

Taken together, the sectoral estimates reveal a significant degree of heterogeneity in the slope of the sectoral NKPC, ranging from 0.04 percent for the food, beverage, and tobacco sector to 0.28 percent for the transport equipment sector. Nevertheless, all of

²⁵We consider nine sectors that jointly span the manufacturing sector in Belgium, excluding the "computer, electronic and optical products" sector due to an insufficient number of observations. Given the smaller sample size of the individual sectoral subsamples, we augment the set of instruments in each sectoral regression to include lags from t-1 to t-8 for both marginal cost and the price index of competitors to identify the parameters of the pricing equation and include lags from t-2 to t-8 to identify the autoregressive coefficients. The additional instruments enhance the identification power and the stability of the empirical estimates, particularly for sectors with smaller sample sizes.

Sector:	θ_s	Ω_s	λ_s	\$s
Transport equipment	0.582	0.077	0.281	0.033
	(0.009)	(0.015)	(0.019)	(0.012)
Electrical equipment	0.632	0.191	0.177	0.020
	(0.005)	(0.021)	(0.007)	(0.007)
Machinery equipment	0.668	0.087	0.154	0.019
	(0.021)	(0.041)	(0.029)	(0.012)
Wood, paper and printing	0.703	0.202	0.102	0.106
	(0.019)	(0.082)	(0.020)	(0.034)
Metals	0.728	0.127	0.091	0.099
	(0.016)	(0.056)	(0.013)	(0.016)
Rubber and plastic	0.714	0.367	0.075	0.106
	(0.018)	(0.047)	(0.012)	(0.011)
Textiles, apparel and leather	0.737	0.509	0.047	0.030
	(0.023)	(0.144)	(0.019)	(0.024)
Chemicals	0.779	0.330	0.043	0.275
	(0.011)	(0.031)	(0.005)	(0.066)
Food, beverages and tobacco	0.760	0.483	0.041	0.289
-	(0.018)	(0.048)	(0.010)	(0.031)

Table 4: Structural estimates and slope of the sectoral Phillips curves: Sectoral estimates.

Notes. This table presents the estimates of the structural parameters (θ and Ω), the implied slope of the NKPC (λ), and the sector-specific Törnqvist weight (\bar{s}) for different manufacturing sectors. The estimates are obtained using Model C. For each sector, observations are weighted in the regression using Törnqvist weights. Standard errors (in parenthesis) are robust to heteroskedasticity and autocorrelation at the firm level.

them are economically meaningful, confirming that movements in marginal costs are a key driver of inflation, even at the sectoral level.

6 Robustness analysis

Before moving to the aggregate implications of our estimates, we present a battery of exercises to test the robustness of our results and discuss some additional potential concerns with our identification strategy.

6.1 Aggregate instruments

The estimation of our models requires instruments for both marginal costs and the price index of competitors, owing to measurement and endogeneity concerns. We have demonstrated that the firm-level instruments used in our baseline specifications are powerful and have provided arguments supporting their validity. We have also shown that our estimates are robust to the inclusion of different sets of industry-by-quarter fixed effects, defined either at a higher or lower level of aggregation.

We now estimate the structural parameters (and therefore the slope of the NKPC) using an alternative set of instruments that leverage variation in aggregate variables, specifically high-frequency monetary policy and oil shocks. The purpose of this exercise is threefold. First, it allows us to further assess the validity of our identification strategy, by confirming that our estimates are robust to using aggregate shocks as instruments. Second, it informs about the structural nature of our estimates, which, according to theory, should not be contingent upon the source of fluctuations-micro or macro, demand or supply-in marginal costs. Third, it provides us with a set of over-identifying restrictions to formally test for instrument validity.

As is standard, we measure monetary policy shocks by looking at surprises in asset prices linked to the movements of policy rates—in our case, the Euro area monetary policy rate. Specifically, we follow Gürkaynak et al. (2005) by leveraging variation in the overnight index swaps in a narrow window around ECB monetary policy announcements.²⁶ The construction of oil shocks is also standard: we follow the approach in Känzig (2021), which uses variation in crude oil futures prices in a narrow window around OPEC announcements.

The fluctuations in money and oil shocks capture shifts in aggregate demand and supply conditions, respectively, resulting in opposite movements in firms' marginal costs. To see this, Figure 2 shows the impulse–response functions of marginal costs to both shocks obtained via local linear projections. As we can see, on average, a positive money shock (tighter than expected monetary policy) significantly reduces the present value of future marginal costs. It does so by reducing demand, which in turn reduces marginal costs, given that firms face upward-sloping supply curves in input markets. In contrast, a positive oil shock (an unexpected increase in oil prices) shifts the present value of marginal

²⁶The time-series of aggregate money shocks is from Altavilla et al. (2019).



Figure 2: Local projections of the money and oil shocks on marginal cost

Notes. This figure displays the impulse response function of marginal cost to aggregate money and oil shocks estimated via local linear projections. The plot reports the coefficients b_h from the regressions $mc_{ft+h}^n - mc_{ft-1}^n = a_f + b_h S_{t-1} + \varepsilon_{ft+h}^S$ for $S \in \{MS, OS\}$ and $h = 1, \dots, 8$ quarters. The dark (light) gray shaded areas represent the 68 (95) percent confidence bands obtained from Newey-West standard errors with four quarters of correlation. All the regressions are weighted using Törnqvist weights.

costs upward. These results are consistent with the evidence on the aggregate effects of money and oil shocks produced by standard structural VAR models.²⁷

We construct our money and oil shock instruments to exploit both time-series and cross-sectional margins of variation. We first recover firm-specific loadings on each shock:

$$mc_{ft}^{n} = a_{f} + b_{f}^{m}MS_{t-1} + \varepsilon_{ft}^{m},$$

$$mc_{ft}^{n} = a_{f} + b_{f}^{o}OS_{t-1} + \varepsilon_{ft}^{o}.$$

The loadings of interest, b_f^m and b_f^o , are indexed by f as we estimate them running separate

²⁷See for example Gertler and Karadi (2015), Känzig (2021), and Gagliardone and Gertler (2023).

regressions for each firm. We then combine the firm-specific loadings with the aggregate shocks to obtain a set of firm-level money and oil shocks instruments: $\hat{MS}_{ft} := \hat{b}_f^m \cdot MS_{t-1}$ and $\hat{OS}_{ft} := \hat{b}_f^o \cdot OS_{t-1}$.²⁸ These instruments vary both in the time series and in the cross-section. Therefore, they are not absorbed by the time fixed effects in our model and we can use them to construct additional moment restrictions to identify the parameters.

	Money shocks		Oil shocks		
	(1)	(2)	(3)	(4)	
	Panel a: Structural Estimates				
θ	0.688	0.705	0.703	0.708	
	(0.014)	(0.005)	(0.007)	(0.006)	
Ω	0.719	0.534	0.748	0.528	
	(0.041)	(0.029)	(0.041)	(0.036)	
ρ^{mc}		0.782		0.801	
		(0.005)		(0.018)	
Firm FE	у	У	у	у	
Industry×Time FE	у	У	У	У	
Number of equations	1	1	3	2	
Number of moments	4	3	8	5	
Hansen J-test χ^2	6.138	6.864	9.454	6.659	
	Panel b: Slope of the Phillips curve				
λ	0.041	0.059	0.032	0.058	
	(0.010)	(0.003)	(0.007)	(0.006)	

 Table 5: Robustness: Alternative instruments

Notes. This table presents a set of regressions where we estimate our structural parameters, and the implied slope of the NKPC, using an alternative set of macro instruments. In columns (1)–(2) the additional instruments are a sequence of money shocks. In columns (3)–(4) a sequence of oil shocks. In columns (1) and (3) the estimates are based on Model B. Columns (2) and (4) are based on Model D. In columns (1) and (3) we calculate present values using equation (16). In columns (2) and (4) we reply on the the AR1 assumption. The discount factor is calibrated to $\beta = 0.99$. All models are estimated using the complete sample (N = 129, 425), although the number of observations in each column varies depending on the lag and lead structure employed in the estimation. In all regressions, observations are weighted using Törnqvist weights. Robust standard errors (reported in parenthesis) are clustered at the sector level.

²⁸In the literature on fiscal multipliers using regional panel data, Nekarda and Ramey (2011) and Nakamura and Steinsson (2014) use similar approaches to construct instruments for regional government purchases and regional military spending. See Murtazashvili and Wooldridge (2008) for conditions regarding the consistency of this class of instrumental variable estimators.

The intuition behind the exogeneity of these instruments is as follows. The aggregate components capture high-frequency surprises in financial asset prices that occur around the ECB and OPEC announcements. As financial assets are forward-looking, surprises observed around a sufficiently tight window after the announcements capture deviations from markets expectations, which are orthogonal to the information available prior to the announcement. The cross-sectional loadings capture the firm-specific average response of marginal costs to these orthogonal shocks.

We estimate Model B and Model D augmenting the set of instruments with current and four lagged values of \hat{MS}_{ft} and \hat{OS}_{ft} .²⁹ The results are reported in Table 5. As we can see, both the estimates of structural parameters and of the ones of the slope are in line with our baseline estimates reported in Table 3. The slope estimate obtained assuming a first-order autoregressive process for marginal costs (columns (2) and (4)) are identical $(\hat{\lambda} = 0.58)$ to the corresponding baseline estimate (Model D). The estimates are slightly lower—but well within the confidence bands—when we do not restrict the process for marginal cost (Model B). The similarity of the estimates across specifications is consistent with the underlying theory that the pass-through of marginal cost should be independent of its sources of variation.

A second important result emerging from Table 5 comes from the test statistics. The inclusion of oil and money shocks in the set of instruments generates a set of over-identifying restrictions, which we use for diagnostics. The bottom of panel a reports the Sargan–Hansen statistic over-identification test statistic: the low test statistics indicate that, across all specifications, we do not reject the null hypothesis that instruments satisfy the exclusion restrictions required by the moment conditions at conventional confidence levels.

²⁹Following the literature on local projections (Jordà 2005; Ramey 2016; Stock and Watson 2018), the set of instruments also includes the 8th-quarter lag of marginal cost. The rationale for doing so is that lag-augmenting the set of instruments makes inference more robust by accounting for the possibility of non-stationary marginal costs. As shown by Montiel Olea and Plagborg-Møller (2021), robust confidence intervals based on lag-augmented regressions have correct asymptotic coverage uniformly over the persistence in the data-generating process. In line with this, estimates obtained from using money or oil shocks as instruments only are noisier though consistent with the ones reported in Table 5. See Table A.3 in Appendix C.3.

6.2 Decreasing returns and macroeconomic complementarities

We derived our benchmark model under the assumption that the economy displays constant returns to scale in the aggregate. This assumption eliminates the possibility that macroeconomic complementarities might influence the slope of the NKPC. Specifically, if returns to scale are decreasing instead, price adjustments are dampened since marginal cost varies inversely with price, with the net aggregate effect of reducing the slope of the Phillips curve (see, e.g., Galí 2015). We now investigate the importance of this channel for our results.

In the general case with arbitrary aggregate returns to scale, the slope of the Phillips curve can be expressed as follows:³⁰

$$\lambda = \frac{(1-\theta)(1-\beta\theta)}{\theta}(1-\Omega)\Theta,$$

where the additional term $\Theta := \frac{1}{1+\gamma v(1-\Omega)}$ captures the role of macroeconomic complementarities that stem from decreasing returns, with *v* denoting short-run returns to scale (as in Section 3) and *y* denoting the elasticity of substitution across goods within industries. In our benchmark model with constant returns (*v* = 0), we have that $\Theta = 1$ and the slope reverts to that in equation (12). Decreasing returns (*v* > 0) reduce Θ , lowering the slope. Importantly, because Ω enters the expression for Θ , strategic complementarities reduce the importance of macroeconomic complementarities by reducing the importance of marginal costs in the determination of prices.

In Appendix B.3 we provide empirical evidence indicating that our results are robust to empirically reasonable departures from the assumption of aggregate constant short-run returns to scale. Following Lenzu et al. (2023), we perform a gross output production function estimation that allows us to obtain sector-specific estimates of returns to scale. Our estimates indicate that the returns to scale of the different sectors, and consequently in the aggregate, are close to unity. Specifically, the sectoral estimates range from 0.86 to 1.02, while the aggregate returns to scale are estimated to be approximately 0.96. This implies a value of \hat{v} of approximately 0.04.

Calibrating γ to 4, which implies aggregate steady-state markup between 1.3 and 1.4, and using our baseline estimate of $\hat{\Omega} = 0.55$, we obtain a value of $\hat{\Theta} = 0.94$. Thus

³⁰See Appendix A.2 for derivations.

macroeconomic complementarities imply a reduction of the slope of seven percent relative to our baseline estimate (from 0.06 to 0.056), which is well within the confidence bounds of our baseline estimates.

6.3 Menu costs

Our benchmark model assumes time-dependent pricing where the frequency of price adjustment is the one implied by Calvo model. In this section, we show that under certain reasonable conditions, our estimates of the slope of the Phillips curve remain valid if one instead models nominal rigidities using a menu cost setting where firms are free to adjust their price each period but are subject to a fixed adjustment cost. This exercise also informs us about how well the Calvo versus menu cost formulation captures price adjustment in our data.

In the conventional menu cost framework, the fixed cost of adjustment gives rise to an endogenous inaction region around the target price that is bounded by "Ss bands." As a result, price adjustments can be broken down into two components: shifts in the reset price given the adjustment frequency (the intensive margin) and shifts in the adjustment frequencies that correspond to shifts in the Ss bands (the extensive margin). As discussed by Caballero and Engel (2007), the extensive margin gives rise to a selection effect, wherein firms farthest away from their target price are more likely to adjust. The selection effect implies that, for a given price adjustment frequency, there will be greater price flexibility in the menu cost framework compared to the corresponding Calvo setup, which features only intensive margin adjustments.

Despite these differences, Auclert et al. (2022) argues that, when aggregate shocks are not too large, there exists an approximate observational equivalence in the aggregate price response to marginal cost shocks between models with Calvo rigidities and canonical models with menu costs (Golosov and Lucas 2007; Nakamura and Steinsson 2010). In particular, the authors show that, up to a first-order approximation of the price response, one can derive two "virtual hazard rates" (i.e., fictitious probabilities of keeping the price fixed between any two consecutive periods) that exactly replicate both the intensive and extensive margins of adjustments. It follows that, quantitatively, an average of the two virtual hazard rates can be used to calibrate a time-dependent model so that the aggregate response of the price level is similar to that of a menu cost model. Moreover, the average virtual hazard rate is approximately flat for a wide range of calibrations, which implies that the adjustment probability declines geometrically. It follows from these results that a Calvo model calibrated with the virtual frequency serves as a good approximation of canonical menu cost models.³¹

When the equivalence result holds, the frequency of adjustment is constant and we can write a general expression for the conditional expectation of a firm's price change:

$$\mathbb{E}\{p_{ft} - p_{ft-1} | p_{ft}^o, p_{ft-1}\} \approx \underbrace{(1 - \theta)(p_{ft}^o - p_{ft-1})}_{\text{Calvo term}} + \underbrace{(\theta - \tilde{\theta})(p_{ft}^o - p_{ft-1})}_{\text{Selection term}},$$

where $\tilde{\theta}$ is the average virtual frequency of adjustment under menu costs (and where Calvo term refers to the intensive margin and the selection term to the extensive margin). As noted above, the selection term captures the fact that adjusting firms are not a random sample of the population, but are exactly those whose reset price is farthest from their price in the previous period. By rearranging the expression above, we obtain the familiar expression for the expected price change:

$$\mathbb{E}\{p_{ft}|p_{ft}^o, p_{ft-1}\} \approx (1-\tilde{\theta})p_{ft}^o + \tilde{\theta}p_{ft-1}.$$
(18)

It follows from equation (18) that one can derive an approximate Phillips curve relationship even under a menu cost framework, where the virtual frequency $\tilde{\theta}$ replaces the actual frequency of price adjustments θ :

$$\pi_t \approx \frac{(1-\tilde{\theta})(1-\beta\tilde{\theta})}{\tilde{\theta}}(1-\Omega) \ \widehat{mc_t^r} + \beta \mathbb{E}_t \{\pi_{t+1}\} + u_t$$

Note that, as discussed in the work of Gertler and Leahy (2008), for a given frequency of price adjustments θ , the selection effects leads to larger price changes ($\tilde{\theta} \leq \theta$), and therefore the slope of the Phillips curve would be higher under a menu cost framework.

The arguments above clarify that, even if the data-generating process features Ss-style price adjustments, the population regression presented in equation (15) identifies correctly the virtual frequency $\tilde{\theta}$ (instead of the frequency of adjustment θ). Therefore, the estimated NKPC slope remains valid and unchanged under a menu cost framework.³²

³¹See also Gertler and Leahy (2008) for conditions under which an exact equivalence result holds.

³²We verify the claim using simulated data. Specifically, we simulate firm-level random walk processes for the reset price and generate time series of prices from a standard Calvo model and from a standard menu cost model, calibrated to have the same frequency of price adjustments θ . We then run regressions in equation (15) on the two simulated databases and confirm that the coefficient on the lagged price is exactly

This discussion also emphasizes why estimating the parameters using firm-level pricing regressions is preferable to calibrating θ based on the frequency of price adjustments: the former approach remains agnostic to, and robust against, the underlying price adjustment mechanism.³³

Our data allows us to take the analysis of price adjustments one step further and ask how important is the selection effect in rationalizing the data or, equivalently, how well does the Calvo model serve as an approximation. Two pieces of evidence suggest that selection does not appear to be a salient feature in our data. First, as discussed in Section 5, information on price changes in the PPI micro-data allows us to obtain a direct estimate of the empirical frequency of price adjustment. We found that it closely matches our estimate of $\tilde{\theta}$, and therefore implies a small role for selection in our sample.

A second piece of evidence comes from the observed kurtosis of price changes. Several papers—including Alvarez et al. (2016) and Alvarez et al. (2022)—discuss how Calvo and menu costs models produce observationally different distributions of price adjustments and, as a result, different values of kurtosis. In Golosov and Lucas (2007) and Nakamura and Steinsson (2010), the distribution of price changes is bimodal, with a large mass corresponding to the Ss bands. This implies a small value of kurtosis—between one and three–under standard parametrizations. In contrast, the size distribution of price adjustments under a Calvo mechanism is less dispersed, featuring a large mass of very small as well as some very large price changes, which implies a larger kurtosis in the order of six. We follow Midrigan (2011) and Alvarez et al. (2016) and measure the kurtosis of price adjustment in our PPI micro-data. We find a kurtosis of 5.4, which is again suggesting that the Calvo model offers a good approximation of price dynamics in our sample.

7 Inflation dynamics

Until now, we focused on estimating the forces that govern firms' pricing behavior and obtained estimates of the primitive form of the slope of the NKPC. In this section, we assess the capacity of our model to capture the inflation dynamics observed in the data. We begin

equal to θ for the Calvo model and is lower than θ for the menu cost model, with the coefficient on the reset price always summing to one. We thank Man Chon Iao for providing the simulated data.

³³We thank Ludwig Straub for the insightful discussions and suggestions on this point.
studying whether our estimates can explain the aggregate time series of producers' prices in the Belgian manufacturing sector. Moving to a finer level of aggregation, we then assess the model's ability to rationalize sectoral inflation dynamics.

7.1 Aggregate inflation

To derive an expression for aggregate inflation, we first close the model by assuming that nominal marginal cost follows a random walk with drift, as is consistent with the empirical evidence.³⁴ We then use the equation for the price index (equation (9)) and the equation for the reset price (equation (10)), along with the restriction on the process for the nominal marginal cost, to obtain the following reduced-form expression for quarterly inflation (see Appendix A.3 for derivations):

$$\pi_t = \tilde{\lambda} \left(m c_t^n - p_{t-1} \right) + \alpha + \theta u_t, \tag{19}$$

where $\tilde{\lambda} \equiv \tilde{\lambda}(\theta, \Omega)$ is an analytical function of the structural parameters, α captures trend inflation, and u_t is the aggregate cost-push shock. According to equation (19), quarterly inflation is increasing in current nominal marginal cost scaled by the lagged price level, consistent with the theory presented earlier. As before, the sensitivity of inflation depends upon the primitive pricing parameters θ and Ω .

Finally, to obtain an expression for the model contribution to year-over-year inflation, we iterate equation (19) backward three periods and suppress the cost-push shocks:

$$\pi_t^{y-y} = \sum_{\tau=0}^3 \tilde{\lambda} (1 - \tilde{\lambda})^{\tau} (mc_{t-\tau}^n - p_{t-4}) + \alpha^{y-y}.$$
 (20)

According to equation (20), year-over-year inflation depends on a distributed lag of nominal marginal cost scaled by the price level at t - 4.

The black line in Figure 3 plots year-over-year inflation in the data, that is the time series of manufacturing inflation constructed using the producer price index data. The red line in Figure 3 depicts the model-implied inflation series according to equation (20).³⁵ The

³⁵In practice, to account for cross-sectional heterogeneity in the primitive parameters, we construct

³⁴We regress mc_t^n on one lag and instrument with the second lag to reduce downward bias due to measurement error. We find that the estimated autoregressive coefficient is $\hat{\rho}^{mc} = 0.987$ (0.015), with Newey-West standard errors in brackets. Additionally, the Dickey-Fuller test does not reject the null hypothesis of unit root with Z = -1.639 and p-value = 0.463. Notice that this estimate is different although consistent with those in Table 3, as those estimates should be interpreted as the persistence of deviations from trend due to the inclusion of time fixed effects.



Figure 3: Aggregate inflation dynamics

Notes. This figure compares the inflation dynamics in the data to the model-implied one. The black line represents manufacturing producer price inflation in the data. The red line is the model-implied manufacturing producer price inflation (π_t^{y-y}) , constructed as a Törnqvist-weighted average of the sectoral model-based inflation rates (π_{st}^{y-y}) . The latter are constructed using the formula in equation (20) and the parameter estimates from Table 4.

difference between the black and red lines is the component of inflation due to cost-push shocks.

Overall, the model effectively tracks the broad swings in inflation over our sample period. It accounts for half of the variation in inflation ($R^2 = 0.5$) with a correlation of 0.71. Particularly noteworthy is its ability to capture the drop in inflation during the 2008 financial crisis and the sharp run-up in 2016. Additionally, the model successfully captures the consistent decline in inflation from 2011 to 2016, although not fully capturing its amplitude.

We conclude with two considerations. First, cost-push shocks play an important role in explaining inflation dynamics. This is consistent with the results of earlier quantitative NK models.³⁶ Secondly, when benchmarking our model to the data, we chose not to introduce any form of lag-dependence in inflation, as is commonly done in the literature

the aggregate inflation rates aggregating model-based inflation rates for different sectors with sectoral Törnqvist.

³⁶For example, in Primiceri et al. (2006), cost-push shocks arising from variation in the desired price and wage markups account for about 70% of inflation volatility.

to improve the model fit (see, for instance, Galí and Gertler 1999, Jørgensen and Lansing 2023, and references therein).



Figure 4: Sectoral inflation dynamics

Notes. This figure assesses the fit of the sectoral Phillips curves. For different manufacturing sectors, it compares the inflation dynamics in the data to the model-implied one. The black lines represent the sectoral producer price inflation series. The red lines are the model-implied inflation rates (π_{st}^{y-y}) , constructed using the formula in the text and the parameter estimates in Table 4.

7.2 Sectoral inflation

We now turn our attention to a finer level of aggregation and assess the ability of our model to rationalize sectoral price dynamics. We use the estimates in Table 4 and follow the steps described in the previous section to obtain a sectoral version of equation (20), now linking sectoral inflation to sectoral parameters and real marginal costs. For each manufacturing sector, Figure 4 plots year-over-year PPI inflation (black line) against

our model-based measure of inflation (red line). Sectors are ordered from top-left to bottom-right according to the estimated steepness of the NKPC slope.

This exercise demonstrates that the model fits the sectoral data reasonably well. In sectors in which we estimated a steeper slope (e.g., "Transportation, electrical, and machinery equipment"), the correlation between the model and the data substantially increases. In other sectors (e.g., "Food, beverage, and tobacco" and "Chemical products") the Phillips curve is flatter and a larger share of the variation of inflation is imputed to cost-push shocks.

8 **Reconciliation with the conventional NKPC**

In this section, we use our firm-level data to estimate the slope of a conventional output gap-based NKPC, which-as we show below-is the product of the slope of the marginal cost-based NKPC and the output elasticity of marginal cost. This exercise helps us reconcile our estimates of a high slope for the marginal cost-based curve with the low estimates of the conventional formulation available in the literature.

As discussed in the introduction, the theory underlying the slope of the conventional formulation is exact only under very restrictive circumstances that permit establishing a proportionality between marginal cost and output. Accordingly, in order to estimate a conventional Phillips curve, we follow the approach adopted by the literature, which postulates—at least implicitly—an approximate log-linear reduced-form relation between marginal cost and output. We estimate the elasticity between the two variables to be low, which reconciles the difference between the two slopes.

8.1 The output gap as the real activity measure

Let $\widetilde{y_{ft}^r}$ denote the (log) deviation of *firm-level* real output y_{ft}^r from its flexible price value y_{ft}^{r*} , the level assuming the firm is free to adjust prices each period. Assume that there is an approximately proportional relationship between $\widetilde{y_{ft}^r}$ and $\widehat{mc_{ft}^r}$. Let σ^y denote the coefficient of proportionality (elasticity) and ε_{it}^y denote an approximation error due, for example, to wage rigidity in the industry:

$$\widehat{mc_{ft}^r} = \sigma^y \widetilde{y_{ft}^r} + \varepsilon_{it}^y.$$

As in the aggregate data, the challenge in constructing the output gap is that we do not directly observe y_{ft}^{r*} , the natural level of output. To make progress, suppose that y_{ft}^{r*} is the sum of an industry-specific output trend y_{it}^{r} and a residual ξ_{ft} . As is standard, this residual reflects factors influencing firms' productivity, assumed to be independent of demand shocks (Galí 2015). This implies the following approximate relation between nominal marginal cost and nominal output:

$$mc_{ft}^n = \sigma^y y_{ft}^n + z_{ft},$$

with $z_{ft} = (1 - \sigma^y)p_{it} + mc_{it}^r - \sigma^y(y_{it}^r + \xi_{ft}) + \varepsilon_{it}^y$ and mc_{it}^r denoting the industry-specific trend in real marginal cost. Following the same steps as in Section 2.4, we can derive the slope of the Phillips curve that uses the output gap as the forcing variable, denoted by κ :

$$\kappa \approx \lambda \cdot \sigma^y$$
,

where λ is the slope of the marginal cost-based Phillips curve, given by equation (12). The equation above shows that the wedge between the slope of the marginal cost-based Phillips curve and the output gap-based one is approximately equal to the elasticity between the two forcing variables.

Within this framework, a straightforward modification of our empirical design enables us to estimate the output-based slope of the Phillips curve and the implied elasticity σ^y . Using the equation above to express mc_{ft}^n as a function of y_{ft}^n , we obtain the following regression model:

$$p_{ft} = \varsigma \ y_{ft}^n + \theta p_{ft-1} + \alpha_f + \alpha_{i \times t} + \underbrace{\psi \xi_{ft} + u_{ft}}_{\text{error term}}.$$
(21)

where the reduced-form coefficient $\varsigma := (1-\theta) \frac{1-\beta\theta}{1-\beta\theta\rho^{mc}} (1-\Omega)\sigma^y$ captures the pass-through from output to prices, which depends on the elasticity σ^y . In this model, the fixed effects soak up the variation in z_{ft} , except for the productivity residual ξ_{ft} , which contributes to the error term.

We estimate the regression model in equation (21) using firm-level value added (revenues minus cost of intermediates inputs) as a measure of nominal output, y_{ft}^n . To draw a tight connection with previous studies, we address the identification problem using variation in output driven by demand-side shocks. Specifically, we estimate equation (21) via two-stage least squares (2SLS) using as instruments the high-frequency monetary policy shocks employed in Section 6.1. As previously discussed, these instruments are

	Marginal cost (1)	Output (2)	 Marginal cost (3)	Output (4)
ς	0.110 (0.030)	0.030 (0.007)	0.072 (0.032)	0.021 (0.004)
λ, κ	0.073	0.017	0.058	0.015
σ^y	(0.024)	(0.004) 0.226 (0.052)	(0.030)	(0.003) 0.253 (0.058)
Firm FE	У	У	У	у
Industry×Quarter FE Industry×3 years FE	У	У	У	У

Table 6: Marginal cost-based vs output-based slopes

Notes. This table presents various reduced-form pricing regressions that use either marginal cost or the output gap as a forcing variable. Columns (1) and (2) show the estimates of the reduced-form elasticity ς from the regression model in equation (21) along with the implied Phillips curve's slopes (either λ or κ , depending on the forcing variable). Column (2) also reports the estimates of the elasticity (σ^{y}) implied by the relation between the marginal cost- and output-based slopes. Columns (2) and (3) report similar estimates but for a variant of the regression model in equation (21) where the (high-frequency) industry-by-quarter fixed effects are replaced by the (low-frequency) industry-by-3years fixed effects. The reduced-form coefficients are estimated via 2SLS. The instruments are the contemporaneous and four lags of the high-frequency firm-specific money shocks (defined in Section C), and the two-year lag of the forcing variable. All the regressions are weighted using the Törnqvist weights. Robust standard errors are clustered at the sector level.

powerful as they leverage both high-frequency aggregate variation in money surprises and a firm's heterogeneous loadings on such shocks.

The estimation results are in Table 6 in column (1). In column (2) we report the estimates of a variant of the same regression model that features marginal cost as forcing variable instead of output (columns (1) and (3)), which allows us to compare how the reduced-form estimates and implied slope coefficients (κ versus λ) change depending on the forcing variable.³⁷ As we can see, the pass-through of output is substantially lower than that of marginal cost, leading to a significant difference in the slopes. Consistent with the structural estimates in Table 3, the marginal cost-based Phillips curve has a slope of 0.073. This estimate is almost four times as large as the estimated slope of the output-based Phillips curve (0.017), which is in line with prior estimates that rely on aggregate output

 $^{^{37}}$ Note that column (1) is essentially the same regression specification as Column (2) of Table 5, but estimated via 2SLS instead of GMM.

gap measures (e.g., Rotemberg and Woodford 1997). The wedge between the two slopes is driven by the low elasticity of marginal cost to value-added output. We can infer from the ratio of the two implied slopes that $\hat{\sigma}^y = \hat{\kappa}/\hat{\lambda}$ is approximately 0.23. Such a low sensitivity is consistent, for example, with a high degree of wage rigidity observed in the data (Alvarez et al. 2006).

One potential concern with the model in equation (21) is that the industry-by-time fixed effects might absorb the general equilibrium effects of factor prices on marginal cost, which would lead to a downward bias in the estimates of the output elasticity of marginal costs. The results presented in columns (3) and (4) suggest that this is unlikely the case. There, we re-estimated the regression model in (21) replacing the industry-by-quarter fixed effects with a lower-frequency counterpart: industry-by-3-year fixed effects. The inclusion of this alternative set of fixed effects allows us to control for medium-term trends while allowing for variation in factor prices at the business-cycle frequency. Overall, the results are robust.

To sum up, the analysis presented in this section helps us reconcile our Phillips curve estimates with the estimates of the conventional output-based Phillips curve available in the literature. Our results indicate that the pass-through from marginal costs to prices is high, but the flatness of the NKPC is likely due to a low sensitivity of marginal cost to output. These results call for more theoretical and empirical work focusing on understanding the structural relationship between output and marginal cost, particularly so given that the elasticity connecting the two could be time-varying and possibly nonlinear.

8.2 Oil shocks, marginal cost and inflation

We argued that it is challenging to assess the impact of supply shocks on inflation through the lenses of the output-based Phillips curve, without relying on a fully specified macroeconomic model that explicitly addresses the endogeneity of the natural level of output. By contrast, the impact of supply shocks on marginal cost is measurable. Therefore one can use the marginal cost-based Phillips curve to trace the effects of supply shocks on inflation.

To illustrate this point, we trace out the transmission of identified oil price shocks on inflation in the data. We then contrast the empirical impulse response functions with the theoretical ones generated by our model calibrated to the estimated parameter.



Figure 5: Dynamic effects of oil shocks

Notes. This figure shows the impulse response function of real marginal cost and price level to aggregate oil shocks estimated via local linear projections. The plot reports the coefficients b_h from the regressions $x_{ft+h} - x_{ft-1} = a_f + b_h OS_{t-1} + \varepsilon_{ft+h}$ for $x \in \{mc^r, p\}$ and $h = 1, \dots, 15$ quarters. The impact is normalized to a 15.7% increase in Brent crude oil price (one standard deviation). The dark (light) gray shaded areas are 68 (95) percent confidence bands obtained from Newey-West standard errors with four quarters of correlation. The red line is the model-based response of prices calculated by feeding in the path of marginal cost (with perfect foresight) to a Phillips curve, calibrated with $\lambda = 0.06$ and $\beta = 0.99$. All the regressions are weighted using Törnqvist weights.

Using the high-frequency oil price shocks described in Section 6, Figure 5 illustrates how movements in oil prices trigger movements in marginal cost and consequent changes in inflation dynamics. Panel a shows the response of aggregate real marginal cost to a one-standard-deviation shock that increases the price of Brent crude oil by roughly fifteen percent. In response to this shock, marginal cost rises by two to two and a half percent during the first two quarters. It then gradually returns to its pre-shock level. The black line in panel b shows that the the transmission of oil shock to marginal costs has substantial effects on inflation. The price level displays a delayed but significant and rather persistent three percent increase within the initial year of the shock.

Next, we assess the ability of our New Keynesian model to reproduce the inflation response to the oil shock, once the Phillips curve is calibrated using parameter estimates presented in the paper. We perform the following exercise. We assume firms have perfect foresight and feed our marginal cost-based Phillips curve a path of the expected real marginal cost that matches the one generated by the impulse response of marginal cost in Panel a. We then compute the implied price dynamic and plot it in Panel b (red dotted line). As we can see, the model performs well in capturing the inflation dynamics induced by the oil shock, as evidenced by the model-based impulse-response function consistently lying within the confidence bands of the impulse response estimated in the data. This exercise provides additional validation of empirical estimates and also demonstrates the utility of the marginal cost-based Phillips curve in tracing the effects of supply shocks on inflation.

9 Concluding remarks

We use disaggregated data to identify the slope of the primitive form of the New Keynesian Phillips curve, which features marginal cost as a relevant measure of economic activity. We differ from previous literature by taking the identification approach all the way to the firm-product level. In particular, our approach leverages a unique dataset that provides high-frequency information on prices, output, and production costs and enables us to identify the primitive determinants of the slope of the Phillips curve.

Our estimates of the primitive parameters reveal a substantial degree of price stickiness—prices remaining fixed for an average duration of three to four quarters—and highlight an important role for strategic complementarities, which dampen the response of prices to marginal cost by half. Nonetheless, at the aggregate level, our estimates imply a slope of the marginal cost-based Phillips curve in the range of 0.05 to 0.06, which indicates a substantial pass-through of marginal cost into inflation. Importantly, these estimates are notably larger than the available estimates of the slope of the conventional formulation of NKPC that features the aggregate output gap or the aggregate unemployment gap as measures of real activity. We reconcile our findings with this literature by presenting evidence that, despite the high pass-through of marginal cost to inflation, the elasticity of marginal cost with respect to the output gap appears to be quite low. While there has been a considerable amount of theoretical work on the connection between marginal cost and inflation, the same is not true for the relation between marginal cost and the output gap and unemployment. For example, recent research has suggested that the elasticity of marginal cost might have increased significantly during the post-pandemic inflation surge (e.g., Benigno and Eggertsson 2023). Understanding why the elasticity may have been low in the pre-pandemic period but high in the post-pandemic period is a fruitful topic for future research as more recent data becomes available.

Using oil shocks as an example, we also illustrate how the identification approach proposed in this paper, along with the use of the primitive formulation of the NKPC, provides a transparent way of examining the transmission of supply-side shocks to inflation. This is another issue of significant relevance, particularly in light of recent events. Studying the role of supply shocks in the post-pandemic period is another interesting topic that we leave for future research.

References

- ALTAVILLA, C., L. BRUGNOLINI, R. S. GÜRKAYNAK, R. MOTTO, AND G. RAGUSA (2019): "Measuring euro area monetary policy," *Journal of Monetary Economics*, 108, 162–179.
- ALVAREZ, F., H. LE BIHAN, AND F. LIPPI (2016): "The real effects of monetary shocks in sticky price models: a sufficient statistic approach," *American Economic Review*, 106, 2817–2851.
- ALVAREZ, F. AND F. LIPPI (2014): "Price setting with menu cost for multiproduct firms," *Econometrica*, 82, 89–135.
- ALVAREZ, F. E., F. LIPPI, AND T. SOUGANIDIS (2022): "Price setting with strategic complementarities as a mean field game," Tech. rep., National Bureau of Economic Research.
- Alvarez, L. J., E. Dhyne, M. Hoeberichts, C. Kwapil, H. Le Bihan, P. Lünnemann, F. Martins, R. Sabbatini, H. Stahl, P. Vermeulen, et al. (2006): "Sticky prices in

the euro area: a summary of new micro-evidence," *Journal of the European Economic Association*, 4, 575–584.

- Аміті, М., О. Ітякнокі, AND J. Konings (2019): "International shocks, variable markups, and domestic prices," *The Review of Economic Studies*, 86, 2356–2402.
- ATKESON, A. AND A. BURSTEIN (2008): "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 98, 1998–2031.
- AUCLERT, A., R. D. RIGATO, M. ROGNLIE, AND L. STRAUB (2022): "New Pricing Models, Same Old Phillips Curves?" Tech. rep., National Bureau of Economic Research.
- BASU, S. (1995): "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review*, 85, 512.
- BENIGNO, P. AND G. B. EGGERTSSON (2023): "It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve," Tech. rep., National Bureau of Economic Research.
- BERAJA, M., E. HURST, AND J. OSPINA (2019): "The aggregate implications of regional business cycles," *Econometrica*, 87, 1789–1833.
- CABALLERO, R. J. AND E. M. ENGEL (2007): "Price stickiness in Ss models: New interpretations of old results," *Journal of Monetary Economics*, 54, 100–121.
- CALVO, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.
- DE LOECKER, J., P. K. GOLDBERG, A. K. KHANDELWAL, AND N. PAVCNIK (2016): "Prices, Markups and Trade Reform," *Econometrica*, 84, 445–510.
- EUROPEAN COMMISSION (2014): "Quality Report on PRODCOM 2014 Annual Data," PRODCOM Working Group.
- GAGLIARDONE, L. AND M. GERTLER (2023): "Oil Prices, Monetary Policy and Inflation Surges," Tech. rep., National Bureau of Economic Research.
- GALÍ, J. (2015): Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications, Princeton University Press.

- GALÍ, J. AND M. GERTLER (1999): "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 44, 195–222.
- GANDHI, A., S. NAVARRO, AND D. A. RIVERS (2020): "On the Identification of Gross Output Production Functions," *Journal of Political Economy*, 128, 2973–3016.
- GERTLER, M. AND P. KARADI (2015): "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 7, 44–76.
- GERTLER, M. AND J. LEAHY (2008): "A Phillips curve with an Ss foundation," Journal of Political Economy, 116, 533-572.
- GOLOSOV, M. AND R. E. LUCAS (2007): "Menu costs and Phillips curves," *Journal of Political Economy*, 115, 171–199.
- GÜRKAYNAK, R. S., B. P. SACK, AND E. T. SWANSON (2005): "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements," *International Journal of Central Banking*, 1, 55–93.
- HAZELL, J., J. HERRENO, E. NAKAMURA, AND J. STEINSSON (2022): "The slope of the Phillips Curve: evidence from US states," *The Quarterly Journal of Economics*, 137, 1299–1344.
- HOOPER, P., F. S. MISHKIN, AND A. SUFI (2020): "Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating?" *Research in Economics*, 74, 26–62.
- HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 131, 1291–1364.
- JORDÀ, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American Economic Review*, 95, 161–182.
- JØRGENSEN, P. L. AND K. J. LANSING (2023): "Anchored inflation expectations and the slope of the phillips curve," in *Federal Reserve Bank of San Francisco Working Paper 2019-27*.
- KÄNZIG, D. R. (2021): "The macroeconomic effects of oil supply news: Evidence from OPEC announcements," *American Economic Review*, 111, 1092–1125.
- KIMBALL, M. S. (1995): "The quantitative analytics of the basic neomonetarist model," .

- LENZU, S., D. RIVERS, AND J. TIELENS (2023): "Financial Shocks, Productivity, and Prices," *Working paper.*
- MARSCHAK, J. AND W. H. ANDREWS JR. (1944): "Random Simultaneous Equations and the Theory of Production," *Econometrica*, 12, 143–205.
- MAVROEIDIS, S., M. PLAGBORG-MØLLER, AND J. H. STOCK (2014): "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve," *American Economic Journal: Journal of Economic Literature*, 52, 124–188.
- McLEAY, M. AND S. TENREYRO (2020): "Optimal inflation and the identification of the Phillips curve," *NBER Macroeconomics Annual*, 34, 199–255.
- MIDRIGAN, V. (2011): "Menu costs, multiproduct firms, and aggregate fluctuations," *Econometrica*, 79, 1139–1180.
- MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): "Local projection inference is simpler and more robust than you think," *Econometrica*, 89, 1789–1823.
- MURTAZASHVILI, I. AND J. M. WOOLDRIDGE (2008): "Fixed effects instrumental variables estimation in correlated random coefficient panel data models," *Journal of Econometrics*, 142, 539–552.
- NAKAMURA, E. AND J. STEINSSON (2008): "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123, 1415–1464.
- ---- (2010): "Monetary non-neutrality in a multisector menu cost model," *The Quarterly Journal of Economics*, 125, 961–1013.
- ——— (2014): "Fiscal stimulus in a monetary union: Evidence from US regions," American Economic Review, 104, 753–792.
- NEKARDA, C. J. AND V. A. RAMEY (2011): "Industry evidence on the effects of government spending," *American Economic Journal: Macroeconomics*, 3, 36–59.
- NICKELL, S. (1981): "Biases in dynamic models with fixed effects," *Econometrica*, 49, 1417–1426.
- OLEA, J. L. M. AND C. PFLUEGER (2013): "A robust test for weak instruments," *Journal of Business & Economic Statistics*, 31, 358–369.

- PRIMICERI, G., E. SCHAUMBURG, AND A. TAMBALOTTI (2006): "Intertemporal disturbances," NBER Working Paper No. 12243.
- RAMEY, V. A. (2016): "Macroeconomic shocks and their propagation," *Handbook of Macroeconomics*, 2, 71–162.
- ROTEMBERG, J. J. AND M. WOODFORD (1997): "An optimization-based econometric framework for the evaluation of monetary policy," *NBER Macroeconomics Annual*, 12, 297–346.
- SBORDONE, A. M. (2002): "Prices and unit labor costs: a new test of price stickiness," *Journal of Monetary Economics*, 49, 265–292.
- STOCK, J. H. AND M. W. WATSON (2018): "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 128, 917–948.
- WANG, O. AND I. WERNING (2022): "Dynamic oligopoly and price stickiness," *American Economic Review*, 112, 2815–49.
- WOODFORD, M. (2011): Interest and Prices: Foundations of a Theory of Monetary Policy, Princeton University Press.

Anatomy of the Phillips Curve: Micro Evidence and Macro Implications

L. Gagliardone M. Gertler S. Lenzu J. Tielens

Appendix

A Derivations

This section provides additional information and derivations of the key equations presented in Section 2. We begin showing how the markup function in the paper maps to the markup functions under two prominent frameworks featuring imperfect competition. We then present the aggregation steps followed to derive the Phillips curve.

A.1 Derivation of Markup function

Dynamic oligopoly with nested CES preferences

Assume that there is a continuum of industries (indexed by *i*) and a finite number of firms N within each industry. Each firm is indexed by f (or j). Within each industry, firms compete à la Bertrand. In this environment, the price indexes for each industry P_{it} and the aggregate price index P_t are defined, respectively, as:

$$P_{it} := \left(\frac{1}{N} \sum_{f=1}^{N} (\varphi_{fit} P_{fit})^{1-\gamma}\right)^{\frac{1}{1-\gamma}}; \quad P_t := \left(\int_{i \in \mathcal{I}} (\varphi_{it} P_{it})^{1-\sigma} di\right)^{\frac{1}{1-\sigma}},$$

where φ_{fit} is a firm-specific demand shifter (firm appeal), and φ_{it} is an industry-specific demand shifter (taste shock). In what follows, we drop the subscript *i* is dropped when redundant and normalize the steady-state price level to simplify the notation. The demand function for firm $f \in i$ takes a nested CES form with the elasticity of substitution across industries $\sigma > 1$ and elasticity of substitution within industries $\gamma > \sigma$:

$$\mathcal{D}_{ft+\tau} = \left(\frac{\varphi_{ft+\tau}P_{ft}^{o}}{\varphi_{it+\tau}P_{it+\tau}}\right)^{-\gamma} \left(\frac{\varphi_{it+\tau}P_{it+\tau}}{P_{t+\tau}}\right)^{-\sigma} Y_{t+\tau}.$$
(A.1)

Firms internalize the dynamic effect of their choices on the industry price index and

on industry demand. Therefore, the residual elasticity of demand faced by firm f takes the following form:

$$\epsilon_{ft+\tau} := -\frac{\partial \ln \mathcal{D}_{ft+\tau}}{\partial \ln P_{ft}^o} = \gamma - (\gamma - \sigma) \frac{\partial p_{it+\tau}}{\partial p_{ft}^o}.$$
(A.2)

We can further characterize the derivative above. First, the price index of competitors of firm f is defined as:

$$P_{it}^{-f} := \left(\frac{1}{N-1} \sum_{j \neq f}^{N-1} (\varphi_{jit} P_{jit})^{1-\gamma}\right)^{\frac{1}{1-\gamma}}.$$

It follows that $P_{it}^{1-\gamma} = \frac{N-1}{N} \left(P_{it}^{-f} \right)^{1-\gamma} + \frac{1}{N} \left(\varphi_{ft} P_{ft}^{o} \right)^{1-\gamma}$. Next, using the definition of the industry price index P_{it} and denoting by $\zeta_{ft+\tau} := \frac{\partial p_{it+\tau}^{-f}}{\partial p_{ft}^{\circ}}$, its derivative with respect to the firms' reset price is given by:

$$\frac{\partial P_{it+\tau}}{\partial P_{ft}^o} = P_{it+\tau}^{\gamma} \left[\left(\frac{N-1}{N} \right) (P_{it+\tau}^{-f})^{-\gamma} \zeta_{ft+\tau} + \left(\frac{1}{N} \right) (\varphi_{ft})^{1-\gamma} (P_{ft}^o)^{-\gamma} \right].$$

Multiplying both sides by $\frac{P_{ft}^o}{P_{it+\tau}^o}$, we obtain:

$$\begin{aligned} \frac{\partial p_{it+\tau}}{\partial p_{ft}^{o}} &= \zeta_{ft+\tau} \left(\frac{N-1}{N} \right) \left(\frac{P_{it+\tau}^{-f}}{P_{it+\tau}} \right)^{1-\gamma} + \frac{1}{N} \left(\frac{\varphi_{ft+\tau} P_{ft}^{o}}{P_{it+\tau}} \right)^{1-\gamma} \\ &= \zeta_{ft+\tau} (1 - s_{ft+\tau}) + s_{ft+\tau}, \end{aligned}$$

where $s_{ft+\tau} := \frac{1}{N} \frac{P_{ft}^o \mathcal{D}_{ft+\tau}}{P_{it}Y_{it+\tau}} = \frac{1}{N} \left(\frac{\varphi_{ft+\tau}P_{ft}^o}{P_{it+\tau}} \right)^{1-\gamma}$ denotes the within industry revenue share of firm f, and $Y_{it+\tau} := \varphi_{it+\tau}^{1-\sigma} \left(\frac{P_{it+\tau}}{P_{t+\tau}} \right)^{-\sigma} Y_{t+\tau}$ is the industry demand. Replacing the expression for $\frac{\partial p_{it+\tau}}{\partial p_{ft}^o}$ into equation (A.2), we find that the within-industry elasticity of demand faced by firm f is given by:

$$\epsilon_{ft+\tau} = \gamma - (\gamma - \sigma) \left[\zeta_{ft+\tau} (1 - s_{ft+\tau}) + s_{ft+\tau} \right]. \tag{A.3}$$

The intuition behind this expression is straightforward. The stronger the reaction of competitors to a firm's price change—captured by $\zeta_{ft+\tau}$ —, the lower the residual elasticity of demand is. A low residual elasticity of demand, in turn, implies that the firm can sustain a higher markup in equilibrium. This result mirrors the one in the dynamic oligopoly environment in Wang and Werning (2022) and it nests a number of static environments featuring imperfectly competitive firms. In a static oligopoly, $\epsilon_{ft+\tau} = 0$ for $\tau > 0$. In

Atkeson and Burstein (2008)'s static Nash oligopoly, $\epsilon_{ft+\tau} = 0$ for $\tau > 0$ and $\zeta_{ft+\tau} = 0$. Under monopolistic competition, $N \to \infty$, which implies $\zeta_{ft+\tau} \to 0$ and $s_{ft+\tau} \to 0$.

We now use this result to derive the expression for the log-linearized desired markup in equation (7) in the paper. As is standard, we log-linearize around a symmetric Nash steady state (Atkeson and Burstein, 2008).³⁸ Log-linearizing the elasticity in (A.3) around the steady state we obtain the steady state residual demand elasticity

$$\epsilon = \gamma - (\gamma - \sigma) \frac{1}{N},$$

which corresponds to the expression in Atkeson and Burstein (2008). In this model, the desired markup is given by the Lerner index $\mu_{ft+\tau} := \ln(\epsilon_{ft+\tau}/(\epsilon_{ft+\tau}-1))$. Log-linearizing this expression and substituting the expression for steady state residual demand elasticity we obtain the expression for the log-linearized desired markup (in deviation from steady state) in equation (7):

$$\mu_{ft+\tau} - \mu = -\Gamma \left(p_{ft}^o - p_{it+\tau} \right) + u_{ft+\tau}^{\mu},$$

where $\Gamma := \frac{(\gamma - \sigma)(\gamma - 1)}{\epsilon(\epsilon - 1)} \frac{N - 1}{N} > 0$ denotes the markup elasticity with respect to prices and $u_{ft}^{\mu} := -\frac{(\gamma - \sigma)(\gamma - 1)}{\epsilon(\epsilon - 1)} \ln \varphi_{ft} + \frac{\gamma - \sigma}{\epsilon(\epsilon - 1)} \frac{N - 1}{N} \zeta_{ft},$

captures residual variation in the markup that depends on the demand shifters and changes in the slope of competitors' reaction function.

Finally, using these expressions, we can show how to obtain the pricing equation in (8). Log-linearizing the industry price index and ignoring constants, we get:

$$p_{it} = \frac{N-1}{N} p_{it}^{-f} + \frac{1}{N} (\ln \varphi_{ft} + p_{ft}^{o}).$$

³⁸The symmetry assumption is standard in the literature (e.g., Midrigan (2011) and Alvarez and Lippi (2014)), which eases the notation but is largely immaterial for our estimation purposes. Relaxing this assumption would imply that firm-specific steady-state demand elasticities, ϵ_f . In this case, the estimates of the parameters of our pricing equations should be interpreted as average across firms. The assumption of Nash steady state, also standard the literature, implies that $\zeta_{j,\tau} = 0$ at the steady state for all *js* and τs . This comes with some loss of generality, but two points can be made. First, as shown by Wang and Werning (2022), one can write a "behavioral" model with the weaker assumption that $\mathbb{E}{\{\zeta_{j,\tau}\}} = 0$ for all *js* and τs that delivers, under specific values for the elasticities σ and γ , a pass-through of shocks to marginal cost into prices that is qualitatively the same as the one produced by the Nash model. Secondly, these considerations also apply to our empirical analysis, as we directly estimate the parameters (Γ , in particular) rather than the underlying elasticities.

Substituting in equation (6) for the markup and rearranging, we obtain:

$$p_{ft}^{o} = (1 - \beta\theta)\mathbb{E}_t \left\{ \sum_{\tau=0}^{\infty} (\beta\theta)^{\tau} \left((1 - \Omega)(mc_{ft+\tau}^n + \mu) + \Omega p_{it+\tau}^{-f} + (1 - \Omega)u_{ft+\tau}^{\mu} \right) \right\}, \quad (A.4)$$

where, as in the paper, $\Omega := \frac{\Gamma}{1+\Gamma}$. This parameter denotes the relative weight on the price index of competitors (p_{it}^{-f}) and captures the importance of strategic complementarities. When Ω is close to one, firms are not strategic and only look at their marginal cost when resetting prices. In particular, $\Omega \to 0$ as $N \to \infty$, which is the monopolistic competition case. The residual in equation (8) $u_{ft} := (1 - \beta \theta)(1 - \Omega)\mathbb{E}_t \left\{ \sum_{\tau=0}^{\infty} (\beta \theta)^{\tau} u_{ft+\tau}^{\mu} \right\}$ is therefore a firm-specific cost-push shock, function of the expectation of future demand shifters and slopes of competitors' reaction function.

Monopolistic competition with Kimball preferences

Assume that the industry output Y_{it} is produced by a unitary measure of perfectly competitive firms using a bundle of differentiated intermediate inputs Y_{ft} , $f \in i$. The bundle of inputs is assembled into final goods using the Kimball aggregator:³⁹

$$\int_0^1 \Upsilon\left(\frac{Y_{ft}}{Y_{it}}\right) df = 1,$$

where the function $\Upsilon(\cdot)$ is strictly increasing, strictly concave, and satisfies $\Upsilon(1) = 1$.

Taking as given the industry demand Y_{it} , each firm minimizes costs subject to the aggregate constraint:

$$\min_{Y_{ft}} \int_0^1 P_{ft} Y_{ft} df \qquad \text{s.t. } \int_0^1 \Upsilon\left(\frac{Y_{ft}}{Y_{it}}\right) df = 1.$$

Denoting by ψ the Lagrange multiplier of the constraint, the first-order condition of the problem is:

$$P_{ft} = \psi \Upsilon' \left(\frac{Y_{ft}}{Y_{it}}\right) \frac{1}{Y_{it}}$$
(A.5)

Define implicitly the industry price index P_{it} as:

$$\int_0^1 \phi\left(\Upsilon'(1)\frac{P_{ft}}{P_{it}}\right) df = 1$$

where $\phi := \Upsilon \circ (\Upsilon')^{-1}$. Evaluating the first-order condition (A.5) at symmetric prices,

³⁹For simplicity we now abstract from taste shocks. See footnote 11 and the derivations in the previous section.

 $P_{ft} = P_{it}$, we get $\psi = \frac{P_{it}Y_{it}}{\Upsilon'(1)}$. Replacing for ψ , we get the demand function:

$$\frac{P_{ft}}{P_{it}} = \frac{1}{\Upsilon'(1)} \Upsilon'\left(\frac{Y_{ft}}{Y_{it}}\right).$$
(A.6)

Therefore, the demand function faced by firms when resetting prices is:

$$\mathcal{D}_{ft+\tau} = \left[(\Upsilon')^{-1} \left(\Upsilon'(1) \frac{P_{ot}^o}{P_{it+\tau}} \right) \right] \left(\frac{P_{it+\tau}}{P_{t+\tau}} \right)^{-\sigma} Y_{t+\tau}$$

Taking logs of equation (A.1) and differentiating, we get the residual elasticity of demand:

$$\epsilon_{ft+\tau} := -\frac{\partial \ln \mathcal{D}_{ft+\tau}}{\partial \ln P_{ft}^o} = -\frac{\Upsilon'\left(\frac{Y_{ft+\tau}}{Y_{it+\tau}}\right)}{\Upsilon''\left(\frac{Y_{ft+\tau}}{Y_{it+\tau}}\right) \cdot \left(\frac{Y_{ft+\tau}}{Y_{it+\tau}}\right)}$$
(A.7)

We now use this result to derive the expression for the log-linearized desired markup in equation (7) in the paper, under monopolistic competition with Kimball preferences. As above, for ease of exposition, we focus on the symmetric state state. Denote the steady state residual demand elasticity by $\epsilon = -\frac{\Upsilon'(1)}{\Upsilon''(1)}$ and by by ϵ' the derivative of the residual demand elasticity $\epsilon_{ft+\tau}$ in (A.7) with respect to $\frac{Y_{ft+\tau}}{Y_{it+\tau}}$, evaluated at the steady state:

$$\epsilon' = \frac{\Upsilon'(1) \left(\Upsilon''(1) + \Upsilon''(1)\right) - \left(\Upsilon''(1)\right)^2}{\left(\Upsilon''(1)\right)^2} \le 0.$$
(A.8)

The equation above holds with equality if the elasticity is constant (e.g., under CES preferences). Also in this model, the desired markup is given by the Lerner index. Log-linearizing the Learner index around the steady state and using equation (A.8), we have that, up to a first-order approximation, the log-markup (in deviation from steady state) is equal to:

$$\mu_{ft+\tau} - \mu = \frac{\epsilon'}{\epsilon(\epsilon - 1)} \left(y_{ft+\tau} - y_{it+\tau} \right)$$

Finally, log-linearizing the demand function (A.1) and using it to replace for the difference log difference in output, we obtain:

$$\mu_{ft+\tau} - \mu = -\Gamma\left(p_{ft}^o - p_{it+\tau}\right)$$

where, in the case of Kimball preferences, the sensitivity of the markup to the relative price is given by $\Gamma := \frac{\epsilon'}{\epsilon(\epsilon-1)} \frac{1}{\Gamma''(1)}$.

Notice that, without loss of generality, $p_{it+\tau} = p_{it+\tau}^{-f}$ because of the continuum of firms within an industry. Substituting into the pricing equation (6) and rearranging leads to the expression equation (7).

Finally, following the same steps as the previous section, we obtain $\Omega := \frac{\Gamma}{1+\Gamma}$ and the corresponding mapping to the pricing equation in (8).

A.2 Aggregation and the Phillips Curve

Suppose $N < \infty$ and order firms in each industry from 1 to N.⁴⁰ ⁴¹ The aggregate price index (in log-linear terms) is:

$$p_t = \int_{i \in I} \left(\frac{1}{N} \sum_{f=1}^N p_{fit} \right) di,$$

(In the paper, we dropped the industry subscript for ease of notation.) Denote by A_{ft}^{\star} for $f \in \{1, ..., N\}$ the set of industries in which the *f*-th firm can adjust. The price index can then be rewritten as:

$$p_t = \frac{1}{N} \sum_{f=1}^N \left(\int_{i \in \mathcal{I}/A_{ft}^\star} p_{fit-1} di + \int_{i \in A_{ft}^\star} p_{fit}^o di \right),$$

where we are using the fact that firms that cannot adjust set their price to their t - 1 level, whereas firms that can adjust set their price to their optimal reset price.

Since A_{ft}^{\star} has measure $1 - \theta$, and the identity of firms that adjust is an i.i.d. draw from the total population of firms, using the law of large numbers for each $f = \{1, ..., N\}$ across industries we have that:⁴²

$$\frac{1}{N}\sum_{f=1}^{N}\int_{i\in\mathcal{I}/A_{ft}^{\star}}p_{fit-1}di=\theta\int_{i\in\mathcal{I}}\left(\frac{1}{N}\sum_{f=1}^{N}p_{fit-1}\right)di=\theta p_{t-1}$$

and

$$\frac{1}{N}\sum_{f=1}^N\int_{i\in A_{ft}^\star}p_{fit}^odi=(1-\theta)\int_{i\in I}\left(\frac{1}{N}\sum_{f=1}^Np_{fit}^o\right)di.$$

⁴⁰Notice that the same argument goes through with minor modifications but heavier notation for $N_i \neq N$ for a non-zero measure of industries. In general, heterogeneity of the parameters can be accommodated by repeating the same argument for each group of homogeneous industries with non-zero measure and then taking weighted averages of different industries. See for example Wang and Werning (2022), appendix C2.

⁴²The i.i.d. assumption implies that: $\int_{i \in B \subseteq [0,1]} p_{fit} di = Pr(B) \int_{i \in I} p_{fit} di.$ Notice also that $\int_{i \in [0,1]} \left(\frac{1}{N} \sum_{f=1}^{N} p_{it}^{-f}\right) di = \int_{i \in [0,1]} \left(\frac{1}{N} \sum_{f=1}^{N} \left[\frac{N}{N-1} p_{it} - \frac{1}{N-1} p_{fit}\right]\right) di = p_t.$

⁴¹Letting $N \rightarrow \infty$, all results hold under Kimball preferences.

Defining $p_t^o \coloneqq \int_{i \in I} \left(\frac{1}{N} \sum_{f=1}^N p_{fit}^o \right) di$ as the average reset price in the economy, we obtain $p_t = \theta p_{t-1} + (1-\theta) p_t^o,$

which is equation (9) in the paper.⁴³

Next, we replace the aggregate reset price, p_t^o , with an expression that depends on aggregate marginal costs and prices. Using the definition of firm-level marginal cost in equation (5), in log-terms we have that:

$$mc_{fit}^n = c_{it} + a_{fit} + vy_{fit}.$$

The average marginal cost in the industry is $mc_{it}^n := \frac{1}{N} \sum_{f=1}^N mc_{fit}^n$, which implies:

$$mc_{it}^n = c_{it} + a_{it} + vy_{it}.$$

Combining the two equations above and subtracting the (log) industry price index on both sides, we obtain an expression related real marginal costs to cost shifters and output:

$$mc_{fit}^{r} = mc_{it}^{r} + (a_{fit} - a_{it}) + v(y_{fit} - y_{it}).$$

We use the demand function to express the log output deviation, $y_{fit} - y_{it}$, in terms of log price prices. In the case of CES preferences (see equation (A.1)), we obtain:

$$mc_{fit}^{r} = mc_{it}^{r} + (a_{fit} - a_{it}) - \gamma v (p_{fit}^{o} - p_{it}) - \gamma v \ln \varphi_{fit},$$

where *y* denotes the within-industry elasticity of substituting parameter.⁴⁴

We then proceed with the following steps in order: we manipulate equation (A.4) to express the reset price in recursive form, decompose firm-level nominal marginal cost into firm-level real marginal cost and the industry price index prices, and finally use equation (A.2) to replace for firm-level real marginal cost:

$$\begin{split} p_{fit}^{o} &= (1 - \beta\theta) \left((1 - \Omega) (mc_{fit}^{n} + \mu) + \Omega p_{it}^{-f} + (1 - \Omega) u_{fit}^{\mu} \right) + \beta\theta \mathbb{E}_{t} p_{fit+1}^{o} \\ &= (1 - \beta\theta)\Theta \left((1 - \Omega) \widehat{mc_{it}^{r}} + \Omega p_{it}^{-f} + (1 - \Omega) (1 + \gamma v) p_{it} + (1 - \Omega) u_{fit}^{\mu} \right) \\ &+ \beta\theta \mathbb{E}_{t} p_{fit+1}^{o} + (1 - \beta\theta)\Theta (1 - \Omega) \left(a_{fit} - a_{it} - \gamma v \ln \varphi_{fit} \right), \end{split}$$

where $\Theta := \frac{1}{1+\gamma\nu(1-\Omega)}$ captures macroeconomic complementarities due to aggregate returns to scale in production. The last term $a_{fit} - a_{it} + \gamma\nu \ln \varphi_{fit}$ is such that

⁴³Notice that $p_t = \theta p_{t-1} + (1 - \theta) p_t^o$ holds with Kimball preferences as well up to a first-order approximation around the symmetric steady state.

⁴⁴A similar expression holds under monopolistic competition with Kimball preferences. In this case, γ is replaced with the corresponding elasticity of relative output to relative prices, $1/\Upsilon''(1)$.

 $\int_{i \in I} \left(\frac{1}{N} \sum_{f=1}^{N} (a_{fit} - a_{it} + \gamma v \ln \varphi_{fit}) \right) di = 0.$ This follows from the i.i.d. assumption on price adjustments, which implies that the average firm-level shifter of resetting firms coincides with the unconditional average.

Finally, averaging across firms and industries, we have that the aggregate reset price is then given by:

$$p_t^o = (1 - \beta\theta) \left((1 - \Omega) \widehat{\Theta mc_t^r} + p_t \right) + \beta\theta \mathbb{E}_t p_{t+1}^o + \frac{\theta}{1 - \theta} u_t,$$

where $u_t := \frac{(1-\theta)(1-\beta\theta)}{\theta}(1-\Omega) \int_{i\in I} \left(\frac{1}{N} \sum_{f=1}^N u_{fit}^{\mu}\right) di$ is an aggregate cost-push shock. Subtracting p_t from both sides and using the log-linearized price index:

$$p_t^o - p_t = (1 - \beta\theta)(1 - \Omega)\widehat{\Theta mc_t^r} + \beta\theta(\mathbb{E}_t p_{t+1}^o - p_t) + \frac{\theta}{1 - \theta}u_t$$
$$\Rightarrow \frac{\theta}{1 - \theta}\pi_t = (1 - \beta\theta)(1 - \Omega)\widehat{\Theta mc_t^r} + \beta\theta\mathbb{E}_t\left(\frac{\theta}{1 - \theta}\pi_{t+1} + \pi_{t+1}\right) + \frac{\theta}{1 - \theta}u_t$$

Rearranging one obtains the marginal cost-based Phillips curve:

$$\pi_t = \lambda \Theta \ \widehat{mc_t^r} + \beta \mathbb{E}_t \pi_{t+1} + u_t$$

where $\lambda := \frac{(1-\theta)(1-\beta\theta)}{\theta}(1-\Omega)$ is the slope. The equation above highlights that macroeconomic complementarities also mediate the pass-through of marginal cost to prices via Θ . Under the assumption of constant aggregate returns to scale, we have that $\Theta = 1$, and the Phillips curve simplifies to equation (11). This condition is exactly verified when $\nu = 0$, but also when $\Omega = 1$.

A.3 Derivations of inflation dynamics

r

Ignoring the intercept, the system of equations is given by:

$$p_t^o = (1 - \beta\theta)((1 - \Omega)mc_t^n + \Omega p_t) + \beta\theta\mathbb{E}_t p_{t+1}^o + \frac{\theta}{1 - \theta}u_t,$$

$$p_t = (1 - \theta)p_t^o + \theta p_{t-1},$$

$$nc_t^n = mc_{t-1}^n + \varepsilon_t^{mc}.$$
(A.9)

We guess and verify using the method of undetermined coefficients that the solution is of the form:

$$p_t^o = \Xi(mc_{t-1}^n + \varepsilon_t^{mc}) + (1 - \Xi)p_{t-1} + \frac{\theta}{1 - \theta}u_t,$$

$$p_t = \tilde{\lambda}(mc_{t-1}^n + \varepsilon_t^{mc}) + (1 - \tilde{\lambda})p_{t-1} + \theta u_t,$$

where Ξ and $\tilde{\lambda}$ are the coefficients to be determined. Plugging the guessed solution into the system gives the following restrictions on the parameters:

$$\begin{split} \Xi &= (1 - \beta \theta)(1 - \Omega + \Omega \tilde{\lambda}) + \beta \theta (\Xi + (1 - \Xi) \tilde{\lambda}), \\ \tilde{\lambda} &= (1 - \theta) \Xi. \end{split}$$

We select the solution that implies that system (A.9) has exactly one eigenvalue larger than one in modulus. This gives the following values for the parameters in terms of primitives:

$$\begin{split} \Xi &= \frac{\beta\theta(2-\Omega(1-\theta)-\theta)+\Omega(1-\theta)-1}{2\beta(1-\theta)\theta} \\ &+ \frac{\sqrt{(-\Omega(1-\theta)(1-\beta\theta)-\beta(2-\theta)\theta+1)^2+4\beta(1-\Omega)(1-\theta)\theta(1-\beta\theta)}}{2\beta(1-\theta)\theta}, \\ \tilde{\lambda} &= (1-\theta)\Xi. \end{split}$$

Rearranging the guessed solution for $p_t = \tilde{\lambda}mc_t^n + (1 - \tilde{\lambda})p_{t-1} + \theta u_t$ and adding back the intercept, we obtain equation (19). Repeated substitution for p_{t-1} gives equation (20).

B Data and Measurement

B.1 Data sources and data cleaning

In this section we describe the different administrative sources used to assemble our micro-level data.

We use the information in PRODCOM to compute the quarterly change in productand firm-level prices and to define the boundary of markets (industries) in which firms compete. PRODCOM is a large-scale survey commissioned by Eurostat and administered in Belgium by the national statistical office. The survey is designed to cover at least 90% of domestic production value within each manufacturing industry (4-digit NACE codes) by surveying all firms operating in the country with (a) a minimum of 20 employees or (b) total revenue above 4.5 million euros (European Commission 2014). Firms are required to disclose, on a monthly basis, product-specific physical quantities (e.g., volume, kg., m^2 , etc.) of production sold and the value of production sold (in euros) for all their manufacturing products.

Products are defined in PRODCOM by an 8-digit PC code (e.g., 10.83.11.30 is "Decaffeinated coffee, not roasted", 10.83.11.50 is "Roasted coffee, not decaffeinated", and

10.83.11.70 is "Roasted decaffeinated coffee"). Industries are defined by the first four digits of the product codes (e.g. Processing of tea and coffee is "Processing of tea and coffee"). Sectors are defined by the first two digits of the product codes (AC is "Manufacture of food products, beverages, and tobacco products"). The industry and sector definitions follow the NACE classification as the first four digits of PRODCOM codes are identical to the first four digits of the NACE classification.

In the raw data, there are approximately 4,000 product headings distributed across 13 manufacturing sectors. The PC product codes have been revised several times between 1999 and 2019, with a substantial overhaul in 2008. We use the conversion tables provided by Eurostat and firm-specific information on firms' product baskets to harmonize the 8-digit product codes across consecutive quarters and harmonize 4-digit industry codes over time.⁴⁵ In most cases, the conversion tables provide a unique mapping of the 8-digit product codes across consecutive years. In a limited number of cases, the mapping is many-to-one, one-to-many, or many-to-many. The many-to-one mapping is straightforward. The one-to-many and many-to-many could be problematic. We are able to deal with most of these cases using information on the basket of products produced by each firm.⁴⁶ In a limited number of cases (less than 0.1% of the sample) we do not have sufficient information to resolve the uncertainty regarding the mapping. We drop these observations from the sample. Table A.1 reports the list of manufacturing sectors and their 2-digit PC codes.

We aggregate monthly information at the quarterly level and construct product-level prices (unit values) by dividing the product-level sales by the product-level quantity sold. As explained in the paper, we are interested in domestic prices, that is prices charged by producers in Belgium. PRODCOM does not require firms to separately report distinguishing between production and sales to domestic and international customers. Therefore, we follow AIK and recover domestic values and quantities sold by combining information from PRODCOM with data on firms' product-level exports (quantities and

⁴⁵The official conversion tables are available at https://ec.europa.eu/eurostat/ramon. The harmonization of the industry code essentially consists in harmonizing the to NACE rev. 1 industry, used before 2008, to the NACE rev. 2 industry codes, used from 2008.

⁴⁶For example, consider a case when the official mapping indicates that product 11.11.11.11 in year t could map to either 22.22.22.21 or 22.22.22.22 in year t + 1. Suppose two firms, f_1 and f_2 , report in period t sales of product 11.11.11.11 in year t. If f_1 reports only sales of 22.22.22.21 and f_2 only reports sales of 22.22.22.22 in year t + 1 we infer that we should map 11.11.11.11 to 22.22.22.21 for the former and map 11.11.11.11 to 22.22.22.22 for the latter.

sales) available through Belgian Customs (for extra-EU trade) and through Intrastat Inquiry (for intra-EU trade).⁴⁷ We use the official conversion tables provided by Eurostat to map the CN product codes classification used in the international trade data to the PRODCOM product code classification.⁴⁸ In the majority of the cases, the CN-to-PC conversion involves either a one-to-one or many-to-one mapping, which poses no issues. We drop the observations that involve one-to-many and many-to-many mappings. These account for less than 5% of the observations and production value.

We apply the following filters and data manipulations to the PRODCOM data. First, following AIK, we only keep firms' observations in a given quarter if there was a positive production reported for at least one product in the quarter. This avoids large jumps in the quarterly values due to non-reporting for some months by some firms. In the rare cases when a firm reports positive values but quantities are missing, we impute quantity sold from the average value to quantity ratio in the months where both values and quantities are reported. Second, we require firms to file VAT declarations and Social Security declarations (as explained below): these two data sources are needed to measure firms' marginal costs.

The second important use of international trade data is to obtain information on international competitors selling their manufacturing products in Belgium. For each domestic firm, the merged Customs–Intrastat data reports the quantity purchased (in Kg) and sales (converted to Euros) of different manufacturing products (about 10,000 distinct CN product headings) purchased by Belgian firms from each foreign country. As is standard, we define a foreign competitor as a foreign country–domestic buyer pair. For each foreign competitor, we aggregate the product-level sales and quantity sold at the quarterly level (the reporting is monthly in the raw data) and compute quarterly prices (unit values) by taking the ratio of the two.⁴⁹

⁴⁷Importantly, in constructing our measure of domestic sales we address issues related to carry-along-trade, which might overstate the amount of production of firms that import products that are destined for immediate sales. See Appendix B.2 in AIK for details.

⁴⁸The first six digits of the CN product classification codes correspond to the World HS classification system.

⁴⁹Some CN codes change over time (although to a smaller extent relative to the PC codes). We use the official conversion tables, also available on the Eurostat website, to map CN product codes across consecutive years. We only make an adjustment if the code is a one-to-one change between two years. We do not take into account changes in PC codes that involve splitting into multiple codes or multiple PC codes combining into one code. Effectively, these changes in the PC codes are treated as if new products are generated.

We leverage data from two administrative sources to measure firms' total production (turnover) and variable production costs at a quarterly frequency. Belgian firms file VAT declarations to the Belgian tax authority that contain information on the total sales of the enterprise as well as information on purchases of raw materials and other goods and services that entail VAT-liable transactions, including domestic and international transactions. The coverage of the VAT declarations is almost universal, with a limited number of exceptions that affect the reporting of sole proprietorship and self-employed and therefore mostly do not apply to the firms surveyed by PRODCOM.⁵⁰ We obtain information on employment and labor costs (wage bill) from the Social Security declarations filed on a quarterly basis by each Belgian firm to the Department of Social Security of Belgium.

We sum firm-quarter level expenses on intermediates and labor to obtain a measure of total variable costs, which we use in the construction of firms' marginal costs. We multiply these costs by the ratio of total manufacturing sales (from PRODCOM) over total sales (from the VAT) to adjust for the fact that some firms also have some production outside manufacturing.⁵¹

Finally, we apply the following data-cleaning steps to deal with missing values and outliers. (i) We focus on manufacturing industries, defined by the NACE rev.2 2-digit codes 15–36, dropping from our sample all product headings that correspond to mining and quarrying and all product codes corresponding to industrial services. (ii) We drop observations referring to firms whose sales from manufacturing products (as measured in PRODCOM) is lower than seventy percent of total firm-level sales (as reported in the VAT declarations). This ensures that our sample includes firms' whose real activity is primarily, if not entirely, in manufacturing. (iii) We exclude firms that operate in the following sectors: "Coke and refined petroleum products" sector, "Pharmaceuticals, medicinal chemical, and botanical products" sector, and "Other manufacturing and repair and installation of machinery and equipment" sector from our samples. In the first sector,

 $^{^{50}}$ Enterprises file their VAT declaration online, either on a monthly or a quarterly frequency, depending on some size-based thresholds. Smaller enterprises (turnover < 2.5M euros excl. VAT) can choose to file at the monthly or quarterly frequency. Larger enterprises file monthly. In the case of multiple plants or establishments under one VAT identifier, the declaration is filed as a single file for that VAT identifier. We aggregate all monthly declarations at the quarterly level.

⁵¹As mentioned below, we conservatively drop observations referring to firms whose manufacturing sales are lower than seventy percent of total sales. In the remaining sample, the ratio has a mean of 0.94 and a median of 0.97, confirming the extensive coverage of PRODCOM.

we do not have a sufficient sample size to ensure a robust estimation of the parameters of interest. The second sector is highly subsidized and a significant fraction of output is produced by subsidiaries of international corporations. The last sector is a residual grouping that consists of firms producing diverse and varied products for which it is difficult to define an appropriate set of competitors. (iv) We keep only observations for which we are able to compute product-level price indexes, the corresponding quantity indexes, competitors' price indexes, and marginal costs. (v) We drop observations for which the quarter-to-quarter change of either the firm-level price index or marginal costs is greater than 100% in absolute value. (vi) Finally, for each firm-industry pair that enters our dataset discontinuously we keep only the longest continuous time-spell. This ensures that each time series used in the estimation has no gaps, which would force us to interpolate making assumptions about prices and marginal costs when the data is not recorded.

Sector	Sector definition	NACE Rev.2	
5000	Sector definition	2-digits codes	
CA	Food products, beverages and tobacco products	10-12	
CB	Textiles, apparel, leather and related products	13-15	
CC	Wood and paper products, and printing	16-18	
CE	Chemicals and chemical products	20	
CG	Rubber and plastics products,	00.00	
	and other non-metallic mineral products	22-23	
CH	Basic metals and fabricated metal products,	24 25	
	except machinery and equipment	24-20	
CI	Computer, electronic and optical products	26	
CJ	Electrical equipment	27	
СК	Machinery and equipment n.e.c.	28	
CL	Transport equipment	29-30	

Table A.1: List of manufacturing sectors

Notes. This table reports the list of manufacturing sectors in our sample and the corresponding 2-digit NACE rev. 2 codes.

B.2 Construction of price indexes

We construct a set of indexes that capture price changes in manufacturing goods at different levels of aggregation (firm-industry, firm, industry, individual manufacturing sector, and whole manufacturing sector).

Firm-industry price index. The main variable of interest is the price of domestically sold manufacturing products at the firm-industry level, P_{ft} , for both domestic and foreign producers. We construct this variable using information on prices changes at the most disaggregated level allowed by the data.

Due to repeated product code revisions, it does not exist a consistent 8-digit product code taxonomy across the entire sample period.⁵² Therefore, we compute the sequence of price changes across consecutive time periods (t and t + 1) by mapping the product codes at t + 1 to their corresponding codes at t, aggregate them at the firm-industry level, and recover the time series of the firm-industry price index (in levels) by concatenating quarterly price changes.

Specifically, denote by \mathcal{P}_{ft} the set of products manufactured by firm f and by P_{pt} the price (unit value) of a given product $p \in \mathcal{P}_{ft}$. We first compute the price change for each product, P_{pt}/P_{pt-1} , appropriately care of any change in product codes. Following the data cleaning procedure in AIK, in the construction of the product-level price changes, we drop product-level observations with abnormally large price jumps in a given quarter $(P_{pt}/P_{pt-1} > 3 \text{ or } P_{pt}/P_{pt-1} < 1/3)$. Then, we construct firm-industry price change as a Törnqvist index:

$$P_{ft}/P_{ft-1} = \prod_{p \in \mathcal{P}_{ft}} (P_{pt}/P_{pt-1})^{\bar{s}_{pt}},$$
(A.10)

where \bar{s}_{pt} is a Törnqvist weight computed as the average of the sale shares between *t* and t-1: $\bar{s}_{pt} := \frac{s_{pt}+s_{pt-1}}{2}$.⁵³ Finally, we use the sequence of quarterly price changes to construct the time series of firm-industry's prices (in levels):

$$P_{ft} = P_{ft_f^0} \prod_{\tau = t_f^0 + 1}^{t} \left(P_{f\tau} / P_{f\tau-1} \right), \tag{A.11}$$

where t_f^0 denotes the first quarter when f appears in our data and P_{ft^0} is the price level in that quarter. We normalize P_{ft^0} to one for all firm-industry pairs f in our dataset. As discussed in the paper, this normalization is immaterial for our empirical analysis as any level-effects are absorbed by the firm-industry fixed effects included in all our empirical specifications.

⁵²See Appendix B.1 for additional information on the data.

⁵³This index accounts for the presence of multi-product firms, by averaging across products produced by the same firm in a given industry. The Törnqvist weights, \bar{s}_{pt} , give larger weights to those produces that account for a larger share of firms turnover.

Firm price index. As discussed in the paper, the vast majority of firms in our data operate in only one (4-digit) industry, which implies that the firm-industry price index, P_{ft} , and the firm price index, \bar{P}_{ft} , coincide. Yet, in a in a limited number of cases, it becomes necessary to construct a firm's price index that aggregates across different firm-industry price indexes. In doing this, we construct the firm-level price index \bar{P}_{ft} following method similar to the one outlined above. Specifically, we construct a Törnqvist index that aggregates across price changes of individual (4-digit) industry bundles $i \in I_f$ produced by firm f in quarter $t: \bar{P}_{ft}/\bar{P}_{ft-1} = \prod_{i \in I_f} (P_{fit}/P_{fit-1})^{\bar{s}_{fit}}$, with Törnqvist weights defined as $\bar{s}_{fit} := (s_{fit} + s_{fit-1})/2$, where s_{fit} is the share of sales of industry i in the firms' total sales (across manufacturing industries). We then concatenate the quarterly price changes above to obtain the price index \bar{P}_{ft} , normalizing the level of the price index to one in the first quarter when the firm first appears in our dataset. Note that for single-industry firms the price index \bar{P}_{ft} coincides with with the firm-industry price index P_{fit} in (A.11).

Competitors price index. Using a similar approach, we construct the price index of competitors for each domestic firm. We start computing quarterly price changes: $P_{it}^{-f}/P_{it-1}^{-f} = \prod_{k \in \mathcal{F}_i/f} (P_{kt}/P_{kt-1})^{\bar{s}_{kt}^{-f}}$, with $\bar{s}_{kt}^{-f} := \frac{1}{2} \left(\frac{s_{kt}}{1-s_{ft}} + \frac{s_{kt-1}}{1-s_{ft-1}} \right)$ denoting a Törnqvist weight constructed by averaging the residual revenue share of competitors in the industry at time *t* (net of firm *f* revenues) with that at time *t* – 1. We then concatenate the changes normalizing the level of the price index in the first period to one. Also, in this case, the normalization is immaterial for estimation purposes as our empirical model always includes firm fixed effects. Note that the set of domestic competitors for each Belgian producer, denoted in the paper by \mathcal{F}_i , includes not only other Belgian manufacturers operating in the same industry but also foreign manufacturers that belong to the same industry and sell to Belgian customers.

Industry, sector, and aggregate price index. We construct the industry-level, sector-level, and aggregate (manufacturing) price indexes by aggregating quarterly firm-level price changes. The formula to construct the percentage change in these price indexes is analogous to the one in (A.10), where now the Törnqvist weights assigned to each firm-industry price change, P_{ft}/P_{ft-1} , captures the (weighted) average market shares of the *f* in its own industry, sector, or manufacturing, respectively. Once again, the level of the indexes is constructed by concatenating changes and normalizing the level of the

price index to one for the first observation in the time series.

B.3 Cost functions

Equation (5) in the paper describes the functional form adopted to characterize firms' marginal costs as a function of output, costs, and technology. This function choice is consistent with variable cost functions of this form:

$$\Psi(Y_{ft}, W, P^m) = \frac{1}{(1+\nu)} C_{ft} \mathcal{A}_{ft} Y^{1+\nu}$$

where C_{ft} denotes the unit cost of a composite of variable input factors that is independent of the scale of production; \mathcal{A}_{ft} is a firm-specific cost shifter; (1+v) denotes the (short-run) elasticity of marginal costs to output. This cost function is rather general and consistent with some of the most empirically prominent functional forms of production technologies, including CES and Cobb-Douglas.

Consider Cobb-Douglas production technologies of this form $Y_{ft} = A_{ft} \left(L_{ft}^{\gamma^l} M_{ft}^{\gamma^m} K_{ft}^{\gamma^k} \right)$, where A_{ft} denotes the firms' total factor productivity (technical productivity or TFPQ). The parameters γ^l , γ^m , and γ^k represent the output elasticities of different inputs and their sum captures the (long-run) returns to scales in production. Solving the cost minimization problem, and omitting the subscripts to economize on notation, yields the cost function:

$$\Psi(Y, W, P^m, R) = (\gamma^l + \gamma^m + \gamma^k) \cdot C \cdot \mathcal{A} \cdot h(Y)$$

where $C := [(W)^{\gamma^l} (P^m)^{\gamma^m} (R)^{\gamma^k}]^{\frac{1}{(\gamma^l + \gamma^m + \gamma^k)}}$ is the cost to buy one unit of the composite inputs bundle; $\mathcal{A} := \left(A (\gamma^l)^{\gamma^l} (\gamma^m)^{\gamma^m} (\gamma^k)^{\gamma^k}\right)^{\frac{-1}{(\gamma^l + \gamma^m + \gamma^k)}}$ is a cost shifter inversely related to firms' productivity; the function $h(Y) := Y^{\frac{1}{(\alpha + \gamma)}}$ captures the curvature of the cost function with respect to output, which depends on returns to scale $(\gamma^l + \gamma^m + \gamma^k)$. Defining $\nu := 1/(\gamma^l + \gamma^m + \gamma^k)$ and differentiating with respect to output we obtain the marginal cost function in the paper.

Consider now generalized CES production technologies of this form $Y = A \left(\gamma^l (L)^{\varrho} + \gamma^m (M)^{\varrho} + \gamma^k (K)^{\varrho}\right)^{\frac{\eta}{\varrho}}$, where the parameter ϱ denotes the elasticity of substitution between different inputs. The distribution parameters, γ^l , γ^m , and γ^k (summing to unity) determine the relative shares of the factors in the total cost. The parameter η measures long-run economies of scale. Solving the cost minimization

problem we obtain the short-run cost function:

 $\Psi(Y, W, P^m) = \eta \cdot C \cdot \mathcal{A} \cdot h(Y)$ with unit cost $C := \left[\left(\gamma^l \right)^{\frac{1}{(1+\rho)}} \cdot W^{\frac{\rho}{(1+\rho)}} + \left(\gamma^m \right) \frac{1}{(1+\rho)} \cdot \left(P^m \right)^{\frac{\rho}{(1+\rho)}} + \left(\gamma^k \right) \frac{1}{(1+\rho)} \cdot \left(R \right)^{\frac{\rho}{(1+\rho)}} \right]$, cost shifter is $\mathcal{A} = \frac{1}{\eta \cdot A^{\frac{1}{\eta}}}$, and returns to scale function $h(Y) := Y^{\frac{1}{\eta}}$. Defining $v := 1/\eta$ and differentiating with respect to output we obtain the marginal cost function in the paper.

As is standard, our definition of marginal costs views labor and intermediates as flexible inputs, whereas capital is viewed as fixed in the short-run and only adjustable in the long-run. Given the above derivations, it is straightforward to show that the cost functions in the paper are consistent with general *short-run* cost functions, which nest Cobb-Douglas and CES as particular cases, that depend on a firm's *short-run returns to scale*.

Continuing with the Cobb-Douglas example, the cost-minimization problem of a firm choosing variable production factors (intermediates and labor) taking fixed inputs (capital) as given:

$$\min_{L,M} TC = TVC + FC = WL + P^mM + FC$$

s.t. $Y = \bar{Y}; \quad K = \bar{K}$
 $Y = A\left(L^{\gamma^l}M^{\gamma^m}K^{\gamma^k}\right)$

where $FC := R\bar{K}$, which is fixed in the short run. Solving the cost minimization problem we obtain the short-run cost function:

$$\Psi(Y, W, P^m) = (\gamma^l + \gamma^m) \cdot C \cdot \mathcal{A} \cdot h(Y)$$

where $C := [(W)^{\gamma^l} (P^m)^{\gamma^m}]^{\frac{1}{(\gamma^l+\gamma^m)}}$ is the cost to buy one unit of the composite flexible inputs bundle; $\mathcal{A} := (\tilde{A} \alpha^{\alpha} \gamma^{\gamma})^{\frac{-1}{(\alpha+\gamma)}}$ is the cost shifter, which is inversely related to firms' productivity ($\tilde{A} := A \cdot \bar{K}^{\beta}$); the function $h(Y) = Y^{\frac{1}{(\alpha+\gamma)}}$ captures the curvature of the cost function with respect to output, which now depends on the short-run returns to scale in production ($\gamma^l + \gamma^m$). Defining $v := 1/(\gamma^l + \gamma^m)$ and differentiating with respect to output we obtain the marginal cost function in the paper.

B.4 Estimates of returns to scale

We estimate the elasticities that determine the returns to scale of production (both shortand long-run) by performing production function estimations. We consider the following log-production function: $y_{ft} = \ln A_{ft} + f(l_{ft}, m_{ft}, k_{ft}; \gamma)$. Here, y_{ft} denotes firm-level output (physical quantity) produced by firm f in period t, and A_{ft} captures a firm's technical efficiency (TFPQ). $f(\cdot)$ is the log-gross output production function, which we model as a Cobb-Douglas aggregate of labor (l_{ft}) , intermediates (m_{ft}) , and capital (k_{ft}) . The vector of structural parameters to be estimated is denoted by $\gamma := \gamma^l, \gamma^m, \gamma^k$, which collects the output elasticities of the different inputs.

Following Lenzu et al. (2023), we construct a firm-level quantity index by deflating firm-level sales by the firm-level price index: $Y_{jt} = \frac{(PY)_{ft}}{\bar{P}_{ft}}$. Labor services are measured using the wage bill, and intermediates costs are measured as the total value of materials and services used in production. A measure of the capital stock is constructed from investments in fixed assets using the perpetual inventory method. We deflate labor, capital, and intermediate inputs using the corresponding industry-level producer price deflators.

We estimate the production function separately for each sector, following the approach developed in Lenzu et al. (2023), which combines the structural approach developed in Gandhi et al. (2020) with the control function approach developed by De Loecker et al. (2016) to control for differences in input quality across firms. This approach identifies the production function parameters by addressing the simultaneity bias that derives from the correlation between input choices and unobserved (to the econometrician) productivity (Marschak and Andrews Jr. (1944)), and it solves the identification problem that affects the estimates of the output elasticities of flexible inputs.⁵⁴ In line with the rest of our analysis, we perform the production estimation for each industry by weighting observations using within-industry sales-based Törnqvist weights.

Table A.2 presents the estimates of the output elasticities and returns to scale for individual manufacturing sectors and for the aggregate economy.⁵⁵ The latter is obtained as a sales-weighted average of the sectoral estimates. As discussed in Appendix B.4,

⁵⁴The details of the estimation routine are provided in the Appendix of Lenzu et al. (2023).

⁵⁵We are unable to perform the production function estimation for a handful of the sector "Computer, electronic and optical products" (CI) due to its small sample size.

the key estimates for our purposes are the ones regarding the elasticities of variable inputs, whose sum pins down the *short-run returns to scale* and determines the strength of macroeconomic complementarities. Consistent with the previous studies (see, e.g., Lenzu et al. (2023) and the references therein), our estimates indicate returns to scale in the ballpark of unity for most sectors and, therefore, in the aggregate.

	Output elasticities		Returns to scale		
Sector	Labor	Intermediates	Capital	Long-run	Short-run
	(<i>γ</i> ^{<i>i</i>})	(γ''')	(γ^{κ})	$(\gamma^{r} + \gamma^{m} + \gamma^{n})$	$(\gamma^{i} + \gamma^{m})$
CA	0.248	0.770	0.094	1.112	1.018
CB	0.201	0.748	0.061	1.010	0.949
CC	0.253	0.729	0.040	1.022	0.982
CE	0.080	0.794	0.124	0.999	0.874
CG	0.242	0.717	0.129	1.088	0.958
CH	0.250	0.721	0.147	1.119	0.972
CJ	0.322	0.646	0.120	1.088	0.967
CK	0.197	0.707	0.180	1.084	0.904
CL	0.149	0.796	0.075	1.020	0.945
Aggregate	0.209	0.749	0.104	1.062	0.958

Table A.2: Estimates of output elasticities and returns to scale

Notes. This table reports the within-sector average production function elasticities estimated following the approach in Lenzu et al. (2023), as described above. The first column indicates the manufacturing sector. The subsequent three columns report the estimates obtained from a quantity production function estimation. The following two columns report the long-run and short-run returns to scale. Each row corresponds to a different manufacturing sector. The last row is a sales-weighted average of the sectoral estimates.

C Additional empirical results and robustness

C.1 Nickell bias

As first discussed in Nickell (1981), the presence of firm fixed effects in our dynamic panel specifications can in principle introduce a bias in the estimation of the stickiness parameter θ . Nevertheless, since the bias goes to zero at rate 1/T where T is the time length of the panel dataset, and the average number of quarters in which we observe a firm is T = 47, the baseline results should be valid. We confirm this by estimating the most parsimonious of the specifications, Model D, on the sample of firms that we are able to

track for at least 70 consecutive quarters. The average firm in this subsample is observed for 78.5 consecutive quarters.

Over this sample, we estimate $\hat{\theta} = 0.701 \ (0.010)$, not statistically different from the baseline estimates in Table 3, and consistent with the frequency of price adjustments in the PPI data. The other parameters estimates, $\hat{\Omega} = 0.437 \ (0.048)$ and $\hat{\rho}^{mc} = 0.806 \ (0.028)$, are also broadly consistent with our previous estimates. These estimates imply a slightly steeper slope $\hat{\lambda} = 0.074 \ (0.011)$.

C.2 Reduced-form models

Our theoretical framework builds on a demand system in which the elasticity of demand perceived by a firm is a function of its own price relative to the industry expenditure function. As we noted earlier, this rather general restriction is consistent with some of the most prominent and empirically tractable demand systems. It implies that the elasticity of markups to a firm's own reset price and to competitors' prices is the same (equation (7)), or equivalently, that marginal costs and competitor prices enter the reset price equation with coefficients of $(1 - \Omega)$ and Ω , respectively. We now relax this assumption and show that we cannot reject the hypothesis that it holds *a posteriori*.

In particular, we estimate the following linear model using two-stage least squares:

$$p_{ft} = \omega^{mc} mc_{ft}^n + \omega^p p_{it}^{-f} + \theta p_{ft-1} + \alpha_f + \alpha_{s\times t} + u_{ft}.$$
(A.12)

Model (A.12) is essentially a reduced-form version of Model C. As in our previous analysis, the parameter θ captures the degree of nominal rigidity. The elasticities of interest, ω^{mc} and ω^{p} , capture the extent to which shocks to marginal costs and competitors' prices are passes through to a firm's own prices while accounting for nominal rigidities and forward-looking behavior. Importantly, through the lens of our model, these two elasticities also provide us with two independent estimates of the complementarity parameter Ω —one based on variation in marginal costs and the other based on variation in the price of competitors:

$$\hat{\Omega}^{mc} := 1 - \frac{\hat{\omega}^{mc}}{1 - \hat{\theta}} \frac{1 - \beta \hat{\theta} \rho^{mc}}{1 - \beta \hat{\theta}}; \quad \hat{\Omega}^{p} := \frac{\hat{\omega}^{p}}{1 - \hat{\theta}} \frac{1 - \beta \hat{\theta} \rho^{p}}{1 - \beta \hat{\theta}}.$$

We take model (A.12) to the data, instrumenting the endogenous variables with the

set of instruments used in our baseline model.⁵⁶ We find that the estimated degree of price stickiness is $\hat{\theta} = 0.732$ (standard error 0.019), which is essentially the same as the baseline estimates presented in Section 5. The estimates of the reduced-form elasticities are $\hat{\omega}^{mc} = 0.082 \ (0.037)$ and $\hat{\omega}^{p} = 0.097 \ (0.028)$, both precisely estimated.

Calibrating ρ^{mc} and ρ^{p} to the estimates of Model C reported in Table 3 and β to 0.99, we use the formula above to recover the two independent estimates of the complementarity parameter: $\hat{\Omega}^{mc} = 0.400 \ (0.242)$ and $\hat{\Omega}^{p} = 0.424 \ (0.120)$.⁵⁷ Based on these estimates, we fail to reject the null hypothesis that the elasticity of markups to a firm's own reset price and to competitors' prices are equal ($\Omega^{mc} = \Omega^{p}$), with a Z-statistic of 0.09 and a p-value of 0.928. This finding is consistent with AIK's results and implies that, even after adjusting for nominal rigidities and forward-looking pricing behavior, firms' pricing behavior in data is largely consistent with the class of models featured in our theoretical framework.

C.3 Aggregate instruments

In section 6.1 we presented a battery of robustness exercises that use aggregate instruments (oil and money shocks) to identify the parameters of interest. Following the literature on local projections, the set of instruments also includes the 8th-quarter lag of the marginal cost. The rationale for doing so is that lag-augmenting the set of instruments makes inference more robust by accounting for the possibility of non-stationary marginal costs. As shown by Montiel Olea and Plagborg-Møller 2021, robust confidence intervals based on lag-augmented regressions have correct asymptotic coverage uniformly over the persistence in the data-generating process.

Table A.3 presents a set of estimation results where the parameters are estimated using only variation in the aggregate instruments, but exclude the the 8th-quarter lag of the marginal cost. Consistent with the arguments in Montiel Olea and Plagborg-Møller (2021), the estimates obtained from using money or oil shocks as instruments only are somewhat noisier—but they are still largely consistent—with the one reported in 5.

⁵⁶The first-stage regressions coincide with the fifth and sixth column in Table 2 (Model C).

⁵⁷We found that estimates of the persistence parameters ρ^{mc} and ρ^{p} remain largely unaffected when estimated by two-stage least squares. If anything, taking into account the uncertainty associated with these parameters would decrease the Z-statistics of the test presented above.

	Money Shocks (1)	Oil Shocks (2)	Lag of MC (3)	
Dep. Var: p_{ft}	Panel a: Reduced-form estimates			
$mc_{f_t}^n$	0.141	0.188	0.122	
Ji	(0.036)	(0.077)	(0.032)	
p_{ft-1}	0.697	0.677	0.702	
- 5	(0.020)	(0.678)	(0.033)	
Firm FE	у	у	у	
Industry×Time FE	У	У	У	
Cragg-Donald Wald F	756	457	8171	
Hansen J-test χ^2	3.842	4.723		
	Panel b: Slope of the Phillips curve			
λ	0.093	0.129	0.080	
	(0.028)	(0.065)	(0.026)	

 Table A.3: Robustness: Alternative instruments

Notes. Panel a shows estimates of Model D using two-stage least squares. The first two columns use money, \hat{MS}_{ft} , and oil, \hat{OS}_{ft} , instruments for the marginal cost, respectively. Both include the contemporaneous instrument and four lags of it. The last column uses the 2-year lag of the marginal cost, mc_{ft-8}^n , as an instrument. Panel b reports the slope of the Phillips curve implied by the estimates in Panel a, calibrating $\beta = 0.99$ and ρ^{mc} to the corresponding baseline estimates from Model D in Tables (3) and (5). All the regressions are weighted using the Törnqvist weights. Robust standard errors are clustered at the sector level.