Price synchronization and cost pass-through in multiproduct firms: Evidence from Danish producer prices *

Luca Dedola[†], Mark Strøm Kristoffersen[‡] and Gabriel Züllig[§]

September 19, 2019

Abstract

This paper studies price adjustment in a novel monthly dataset of prices of multiproduct firms, merged with firm-level balance sheet and cost data. Menu costs models of multiproduct firms have been shown to be able to generate real effects of monetary policy as large as those of standard time dependent models because of within-firm price synchronization. Microeconomic evidence on actual price decisions of multiproduct firms is thus crucial to understand the monetary transmission. First, we document key descriptive properties of the pricing dynamics across firms, finding that these statistics are broadly invariant to the number of goods firms produce. However, we show that the (unconditional) size distribution of price changes is similar to that generated by multiproduct firm models with menu costs common across goods, but also some degree of time dependence. Second, we consider the extent of price synchronization in multiproduct firms using a discrete choice framework. We find that upwards and downwards adjustments are differently affected by fundamentals, and that there is strong synchronization of adjustment decisions within firms relative to the industry, especially as the number of goods increases. Third, we exploit the richness of our dataset to estimate the pass-through of cost shocks along intensive and extensive margins, modelling them jointly to address endogenous selection bias. Concerning the extensive margin, we finding significant state-dependence in the decision to adjust prices. Concerning the intensive margin conditional on adjustment, we find that the medium run pass-through of energy costs is close to 1 and thus basically complete, and much larger than the medium-run pass-through of the costs of imported inputs. For the latter we also find significant heterogeneity across firms with fewer and more products.

JEL classification: D22, E31, F41

Keywords: multiproduct firms, producer prices, cost shocks

---

*Preliminary and incomplete draft. We are grateful to Kim Abildgren, Saroj Bhattarai, Antoine Bouet, Olivier Coibion, Wouter Den Haan, Sebastian Heise, Francesco Lippi and Joseph Vavra for helpful comments and discussions. Furthermore, we thank seminar participants at Danmarks Nationalbank and the Universities of Bordeaux and Copenhagen for helpful comments and suggestions. We would also like to thank Janni Stavad and Nicklas Elversøe from Statistics Denmark for their substantial help and effort regarding the Danish PPI data. The views expressed in this paper do not necessarily reflect the views of the European Central Bank or Danmarks Nationalbank. All remaining errors or omissions remain our own.

[†]European Central Bank and CEPR, e-mail: luca.dedola@ecb.int

[‡]mark.stroem.kristoffersen@gmail.com

[§]Danmarks Nationalbank and University of Copenhagen, e-mail: gaz@nationalbanken.dk

1 Introduction

We study price adjustments by multiproduct firms in a novel monthly micro price dataset, merged with firm balance sheet and cost data. Specifically, we use monthly producer price micro-data from the dataset that is used to compute the producer price index (PPI) by the Danish statistical office.[1] A crucial feature of the data that makes it relevant to an analysis of pricing by multiproduct firms is that there is substantial variation in the number of goods across more than 1,000 firms. This allows us to study how price setting features vary with the number of goods. The theoretical literature on price setting has pointed out that state-dependent price setting models of multiproduct firms can generate real effects of monetary policy as large as those in time-dependent price setting models (see e.g. Midrigan (2002), Alvarez and Lippi (2014), Alvarez et al. (2016)).

PPI micro data are especially useful to analyze in light of the above literature, as noted by Bhattarai and Schoenle (2014), since they are consistent with the basic assumption of virtualy all price-setting model in macroeconomics, where it is actually producing firms that set prices (rather than retailers whose prices are comprised in the CPI). A similar analysis of producer pricing decisions is not feasible with CPI data since the CPI sampling procedure maps to stores, so-called "outlets", which may sell goods from any number of firms, including imports. This makes pricing a complicated web of decisions that involves the whole distribution network. Moreover, it is generally also not possible to identify the producing firms for specific CPI items. In contrast, a further advantage of our dataset is that we can link prices to balance sheet and cost data at the firm level.

A second advantage of PPI micro data, relative to consumer prices, is that they contain very few "sales" prices (namely very short-lived price changes that are quickly reverted, see e.g. Bils and Klenow (2004) and Nakamura and Steinsson (2008)). For this reason, PPI microdata do not necessitate any necessarily ad-hoc "filtering" to make them amenable to interpretation through the lens of standard price setting models (we confirm this feature in our dataset below).

Our main findings are as follows. First, we document key descriptive properties of the pricing dynamics across firms, finding that these statistics are broadly invariant to the number of goods firms produce, in contrast with Bhattarai and Schoenle (2014). However, we show that the (unconditional) size- distribution of price changes is similar to that generated by multiproduct firm models with menu costs common across goods, but also some degree of time dependence, similarly to that studied by Alvarez et al. (2016).

Second, we consider the extent of price synchronization in multiproduct firms. We provide strong evidence of synchronization of price adjustment decisions using a multinomial logit discrete

[1] See Nakamura and Steinsson (2008) for a description of US PPI data; PPI microdata of other European countries were analyzed in Vermeulen et al. (2012).

choice framework. We find that upwards and downwards adjustments are differently affected by fundamentals, and that there is strong synchronization of adjustment decisions within firms relative to the industry, especially as the number of goods increases.

Third, we exploit the richness of our dataset to estimate price adjustment to cost shocks along extensive and intensive margins, modeling them jointly in the two-step econometric framework of Bourguignon et al (2007), specifically devised to address endogenous selection bias. Concerning the extensive margin, we build on our previous results from the multinomial logit, finding signifcant state-dependence in the decision to adjust prices. Concerning the intensive margin conditional on adjustment, we find that the medium run pass-through of energy costs is close to 1 and thus basically complete, and much larger than the medium-run pass-through of the costs of imported inputs. For the latter we also find significant heterogeneity across firms with fewer and more products.

The rest of the paper is organized as follows. Next section describes our datasets (though details are relegated to the appendix), while Section 3 presents some commonly used descriptive statistics on price changes, where we focus on the multiproduct dimension of firms. Section 4 reports the results of our empirical analysis of the extensive and intensive margin of price adjustment to energy and import costs.

2 Danish producer price data and multiproduct firms

The Danish PPI contains a large number of monthly price quotes for individual "items", that is, particular brands of products with certain time-persistent characteristics. These items, which we hence-forth dub "goods", are selected to represent the entire set of goods produced in Denmark and are sampled according to a multi-stage design — see Appendix A for details.

Two key differences relative to the US PPI data used in Bhattarai and Schoenle (2014) are that first, Danish PPI prices are collected at the firm/enterprise level rather than the establishment level ("price-forming units" usually defined to be "production entities in a single location", by the BLS); and second, that both domestic and export prices are reported. Both features of the data imply that relying on the US PPI micro data may actually lead to underestimating the number of products at the firm level.

2.1 Identifying and grouping multiproduct firms

The PPI data allow us to identify firms according to the number of goods produced by them. This distinction uses the firm identifiers and then counts the number of goods in the data for each firm and at any point in time. Following Bhattarai and Schoenle (2014), we then group the firms into the following five good bins according to the number of goods reported: i) firms with only one good at any point in time; ii) firms with 1 to 3 goods on average; iii) firms with more than 3 and up to 5 goods on average; iv) firms with more than 5 to 7 goods on average; and v) firms with more than 7 goods on average. Thus, firms in higher bins are likely to sell a greater number of goods than firms in lower bins. Since we have both domestic and export prices in our dataset, we group firms on the basis of the total number of goods they produce and on the number of domestic goods only. Note that we have many firms who have non-integer numbers of goods due to the averaging of the monthly number of goods for each firm. This is another reason to have bins.

In Table 1 we present descriptive statistics on firms according to the groups that we construct. The mean (median) number of goods per firm across these bins is 1 (1), 2.6 (3), 4.12 (4), 5.81 (5.79), and 19.0 (11.5) respectively. The product dispersion is comparable to that in the US PPI dataset, with the exception of the last bin for firms with more than 7 prices, which is more dispersed in our dataset. Observe that the Danish data contains 1,140 firms, compared to more than 28,000 in the US PPI.

The table also shows that while the majority of firms, around 80%, fall in bins 1 to 3, firms in bins 4 and 5 produce more goods, so that they account for a much larger share of prices than of firms. Firms in bins 4 and 5 set around 50% of all prices in our data. This is also comparable with the US PPI data. The distribution across bins is robust to only including goods sold in the domestic market. When grouping the firms according to the number of domestic goods they sell, goods of firms with up to 3 products represent a larger share of our sample, but prices set by firms with 5 or more products still make up 40% of the dataset.

Finally, regarding firm size, the table reports two statistics, mean and median employment at the firm level, where mean employment is defined both at the firm level and as employment per average number of goods per firm. Clearly, in line with the results in Bhattarai and Schoenle (2014), firms producing more goods do not have more employees per good, but they are overall larger than firms producing less goods.

3 Evidence on price adjustment and synchronization in multiproduct firms

We report the results from our empirical analysis in two parts below, following Bhattarai and Schoenle (2014). First, in contrast to the latter paper, we show there are insignificant differences in aggregate statistics on price changes, such as frequency, size, direction, and dispersion of price changes, across Danish producers with different numbers of goods. These findings are consistent with some specifications of the fixed costs of changing prices at the firm level in Alvarez and Lippi (2014), where those costs increase with the number of price changes, rather than being constant across them.

Second, we provide evidence of synchronization of price adjustment decisions using a discrete choice framework. We find that upwards and downwards adjustments are differently affected by fundamentals, and that there is strong synchronization of adjustment decisions within firms relative to the industry, especially as the number of goods increases. Specifically, we find that, other things equal, the likelihood of an individual price cut (hike) rises with the number of positive (negative) changes in the other prices within a firm, consistent with common costs of changing prices.

3.1 Aggregate price change statistics

We first show that key aggregate price change statistics are broadly invariant to the number of goods produced by firms, and to their size. Recall that a crucial difference relative to the US PPI data used in Bhattarai and Schoenle (2014) is that Danish PPI prices are collected at the firm level rather than the establishment level, and both domestic and export prices are reported. On the one hand, both features of the data imply that the US PPI may actually understimate the number of products for a firm. On the other hand, there could a lower degree of complementarity in price changes within a firm than within an establishment. For the sake of comparability, we compute statistics for domestic prices only and for both domestic and export prices (which however we do not report to save space).

In relation to recent studies of price adjustment using CPI micro data from scanners at retail stores, it is worth noting the following aspects of the Danish PPI data at the outset, for what concerns sales and product replacements. First, while sales are important in the CPI data as documented by Nakamura and Steinsson (2008), Berardi et al. (2015) or to a lesser degree Wulfsberg (2016), they are not a major source of price adjustments in the PPI data. In order to check the relevance of sales in our data, we apply a sales filter similar to "filter B" in Nakamura and Steinsson (2008), where we define as a "sale" every price decrease that is fully reverted after 1, 2, or 3 months. This is the case for just 0.29% of all price observations or 3.79% of all price decreases. There is instead no evidence of "reversed sales", i.e. temporary price increases that are fully reverted according to "filter B". Figure 1 shows the average price index after price increases, decreases without sales and the identified sales prices separately. Interestingly, not only is the typical price decrease identified as a sale price much less persistent (by construction) than the typical non-sale price decrease, but it is also smaller. Therefore, we do not exclude sales prices from our analysis (but do control for them in our regressions in Section 3).

Second, Nakamura and Steinsson (2008) also show that for aggregate statistics on price changes, accounting for product substitutions can make a difference, especially in the CPI. In our PPI dataset, product replacements are less important since only 0.7 per cent of all price changes (including zero changes) and 0.8 per cent of all non-zero price changes are due to product replacements.

3.1.1 Frequency of price changes

We start by computing the frequency of price changes separately for each bin. This allows us to trace out how it varies with the number of goods per firm, a distinguishing feature of price setting models of multiproduct firms. We compute the frequency as the fraction of price changes for a representative good of a representative firm. We compute the frequency as the mean fraction of price changes during the life of a good. We do not count the first observation as a price change and assume that a price has not changed if a value is missing. For export prices, we consider price changes that occur both in Danish Kroner and in the currency in which the price is set according to the reporting firm, if the two differ. Also, we do not explicitly take into account issues of left-censoring of price-spells. For our purpose, it is most relevant that we apply our method consistently across all firms. After computing the frequency of price changes at the good level, we calculate the median frequency for all goods within the firm. Then, we report the mean, median, and standard error of frequencies across firms in a given good bin. We use this standard error to compute 95% confidence intervals throughout the paper.

A first key finding is that the frequency of price changes is broadly independent of the number of goods per firm. Figure 2a and 2b show this graphically for the sample including both domestic and export prices. The mean frequency in Figure 2a is slightly above 20% in both bin 1 and bin 5, and slightly below in bins 2, 3 and 4. The median frequency in Figure 2b is very similar across all bins, below 10%. Inverting these frequency values, this implies that the mean duration of a price spell is around 5 months across all bins, while the median duration is around 8 months across all bins. Interestingly, these values imply that prices change less frequently in the Danish PPI than in the US PPI.

Also, note that across all bins more than 65% of these changes (over all non-zero price changes) are positive price changes, as one would expect under trend inflation, and as summarized in Figure 3. Firms thus adjust prices upward with similar frequency independently of the number of goods they produce.

These results are very similar when we look only at domestic prices (if anything now prices of one-good firms in bin 1 seem to change more frequently than prices of multiproduct firms). Therefore, in general, Danish firms that produce a greater number of goods do not seem to change prices more frequently.

Second, we find a substantial seasonal component of PPI price changes, in striking similarity to Nakamura and Steinsson (2008). Figure 4 presents the median frequency (panel (a)) and the mean absolute size (panel (b)) of both price increases and decreases by month — whereas results for decreases are very similar whether we include or exclude sales. Four results stand out. First, the frequency of price changes declines monotonically over the first three quarters, and then is roughly constant. Second, in all four quarters, the frequency of price changes is largest in the first month of the quarter and declines monotonically within the quarter with the exception of September. This gives rise to the pattern of local peaks in the frequency of price changes in January, April, July, and October. Third, price increases play a disproportionate role in generating seasonality in price changes. Producer prices are twice as likely to change and increase in January than on average in other months of the year. Fourth, seasonality is much less apparent in the mean size of price increases and decreases, and if anything follows a different pattern than in the price change frequency. Mean price inceases are not larger in the months at the beginning of quarters, when the frequency is higher; price decreases are larger and more frequent in January.

Overall, these results suggest some time dependence of price changes, with possibly significant implications for monetary policy transmission. Olivei and Tenreyro (2007) show that the real effects of monetary policy shocks in the US differ depending on the quarter of the year in which the shock hits. They argue that seasonality in the flexibility of wages can explain their empirical findings. Our result that a disproportionate number of price increases are recorded in January could point to similar effects in Denmark and even in the euro area, as Álvarez et al. (2006) also find that prices are significantly more likely to change in January in the euro area. However, the size of price changes does not seem to be much larger in January, pointing to other mechanisms beyond large seasonal changes in firms' costs or demand.

3.1.2 Size and distribution of price changes

We next look at the distribution of price changes, finding no clear relationship with the number of goods produced in our data. Firms with different number of goods change their prices by similar amounts, and display a similar fraction of small price changes and dispersion of price changes.

We compute the size of price changes as the monthly log difference. Again, we compute this at the good level, take the median across goods in a firm, and then report the mean across firms in a good bin. Figure 5 shows that the mean absolute size of price changes is broadly constant at around 7% across bin 1 to 5. This feature holds even when we separate out the price changes into positive and negative price changes, conditional on adjustment. Figure 6 shows that the mean size of positive price changes and the mean size of negative price changes are unrelated to the number of goods. Thus, in general, firms that produce a greater number of goods adjust their prices by a similar amount, both upwards and downwards, to firms that produce a smaller number of goods.

In Figure 7 we also report the entire distribution of price changes for each different bin, following Alvarez et al. (2016). Specifically, to account for heterogeneity across goods we standardize by good category at the 2-digit level, by subtracting the category mean and dividing by the standard deviation. On each chart we superimpose the densities of the Standard Normal and the standardized Laplace distribution (both with unit variance). Recall that the Laplace distribution has a kurtosis of 6 and is thus more peaked than the Normal. It is apparent that the empirical distributions of standardized price changes are more peaked than the Normal across all bins. Their kurtosis is also close to that of the Laplace distribution, ranging from 4.7 for firms with more than 7 goods to 5.3 for firms with more than 1 good up to 3 goods — Bhattarai and Schoenle (2014) report much larger values for multiproduct firms. Interestingly, these distributions can be well approximated by the model with both random menu costs and firms with 4 or more goods studied in Alvarez et al. (2016).

Overall, we find that the distribution of price changes is very similar independently of the number of goods produced by a firm.

3.2 Synchronization in the likelihood of price changes

To further improve our understanding of the determinants of individual price adjustment, we go beyond providing aggregate statistics, and estimate a discrete choice model of pricing decisions. Specifically, as in Bhattarai and Schoenle (2014) we estimate a multinomial logit model for the decision to increase or decrease prices. This allows us to separately examine the relationship of upwards and downwards adjustment decisions with the explanators. The regression setup also provides us with a way to analyze synchronization of decisions within and across firms.

We find some important results: First, there is substantial synchronization of price changes within a firm which suggests a vital role played by complementarities in the cost of changing prices. Second, there is substantially more synchronization of individual adjustment decisions at the firm level relative to the industry. Third, we find evidence for elements of state-dependent pricing in response to variables related to costs.

3.2.1 Synchronization within firms and across industries

We estimate a multinomial logit model of the following form:

$$\Pr\left(Y_{ijt} = 1, 0, -1 \left| X_{ijt} = x\right.\right) = \Phi\left(\beta X_{ijt}\right),$$

where $Y_{i,j,t}$ is an indicator variable for positive, zero, or negative (log) price changes of good i produced by firm j at time t, with 0 as the base category, and $\Phi(\cdot)$ is the logistic function. We use estimates of β to report marginal effects on the change in the probability of adjustment, given one-standard-deviation changes around the mean of X. The logit model has the convenient property that the estimated coefficients take on the natural interpretation of the effect of the explanators on the probability of adjusting prices up or down over taking no action. We focus on estimation separately by bins which distills out trends of how multiproduct firms are different from single-product firms. However, since the results for the PPI as a whole are of independent interest, we estimate the model on pooled data across all good bins as well. This analysis also serves as a preliminary examination of selection into changing prices, which we carry out in the next section along with the estimation of the response to cost shocks.

Our explanatory variables include three sets of covariates: First, we try to capture the extent of synchronization in price setting at the firm and the sectoral level. To this purpose, we use the fractions of positive and negative price changes within the same firm, excluding the price change of the good we are trying to explain. As suggested by the theoretical model of Alvarez and Lippi (2014), to try to distinguish between synchronization due to complementarity in costs of changing prices and in common marginal costs, we also use both the average of absolute price changes and the average price change of goods in the same firm. Second, we include the fraction of positive and negative price changes as well as the average price change in the same industry at the 2-digit NACE sector (excluding firm j). Furthermore, we include aggregate CPI inflation in the given month t, monthly dummies and dummies for export prices, sales, product replacements which we identify as changes in the base price at resampling as well as breaks (see Appendix for details). Finally, we control for the size of the firm by including the log number of employees.

Table 2 reports our results for each bin from 1-3 goods to more than 7 goods.[2] Estimates for price decreases are shown in the top panel, while estimates for price increases are shown in the bottom panel. We find that there is strong evidence of synchronization at the firm level. Specifically, the probability of reducing (raising) prices significantly increases with *both* the fraction of positive

[2] Within-firm synchronization in single-product firms is by definition impossible. Results for the overall sample are robust to including these firms by setting to zero the corresponding within-firm variables.

and negative price changes, and with the average of absolute price changes within the firm (but it decreases for price raises although there is substantial heterogeneity across firms with a different number of products). Conversely this probability declines (rises) in the average price change, also within the firm, with a similar marginal effect for price increases and decreases. The fractions of positive and negative price changes within the firm are especially large and significant across all bins. These results are strongly consistent with synchronization in price changes because of both firm-wide shocks to marginal costs, for the fraction of similarly signed price changes, and common costs of changing prices within the firm, for the fraction of opposite-signed price changes. While the former effects are stronger for price increases than price reductions, the latter effects are similar across negative and positive price changes. Moreover, the effect of price changes of the opposite sign within the firm increases in a statistically significant way with the number of goods produced by firms, although by much less than the effect of same-signed price changes, which almost doubles from the first bin (1-3 products) to the last bin (+7 products), for both price decreases and increases.

Conversely, we find significant but quantitatively smaller evidence of synchronization at the industry level. The probability of a negative (positive) price change increases with the fraction of negative (positive) price changes in the same industry, but it is in general much less affected by the fraction of price changes with the opposite sign. This evidence is consistent with industry-level synchronization because of common shocks to marginal costs across firms, as also shown by the significance of the effect of the average price change at the industry level, which is also larger than that of the average price change within the firm. Nevertheless, we find that firms with more than 7 products are more likely to decrease (increase) prices when the fraction of price increases (decreases) in their industry is larger.

Finally, a long-debated and important question in monetary economics concerns whether price setting is time-dependent or state-dependent. On the one hand, we find that there is substantial time-dependence in the probability of changing prices because of calendar effects, as already discussed above. Specifically, the probability of a price increase is significantly larger in January, April, July and October, than in other months, irrespective of the number of goods produced; conversely, the seasonal pattern for price decreases is not statistically significant. The marginal effects are about a factor three larger than the coefficient on the average industry price change.

On the other hand, we contribute some preliminary evidence of state-dependence by explicitly considering the role played by CPI inflation, an important aggregate variable, for pricing decisions. We find some evidence in support of what one would expect from a model where firms adjust prices in a state-dependent fashion. As Table 2 shows, the likelihood of a price decrease falls with higher CPI inflation while the likelihood of a price increase rises, by a similar amount. Interestingly, marginal effects of CPI are stronger than those of the average firm- and industry price changes. Specifically, a 1% increase in aggregate inflation is associated with a 0.7% higher chance that the firm increases the price in a given month; a similar fall in inflation increases the probability of a price cut by 0.5%. However, coefficients are not always significant across all bins, especially for firms with more than 7 products in the last bin. In the next section we expand our analysis of the determinants of both the probability and the size of price changes — the extensive and intesive margin — by bringing to bear information about firm-level cost shocks.

4 Evidence on the extensive and intensive margin of adjustment to cost shocks in multiproduct firms

In this section we analyze pass-through of cost shocks into prices of our sample of multiproduct firms, looking at both the extensive and the intensive margins of adjustment. We exploit the possibility of linking prices with balance sheet data via a (masked) firm identification number. We focus on two different kinds of costs, the first one with a common component across firms, namely the price of energy; the second with an idiosyncratic component, firm specific import prices. In estimating cost pass-through, we take into account the interaction between the extensive and the intensive margin, building on our multinomial logit results in Section 3.2. In the rest of this section we first review some helpful theoretical results on price adjustment to cost shocks; then we present evidence on the extensive and intensive adjustment margins. First, we find that shocks to energy prices and the prices of imported inputs significantly affect the probability of changing firm-level prices; however, selection bias, while statistically significant does not seem to be economically relevant. Second, conditional on price adjustment, estimated medium run pass-through after 12-18 months is quite different across these two cost shocks; it is more gradual but larger over the medium run for energy costs than the cost of imported inputs.

4.1 Cost pass-through under flexible prices

The general price setting equation under imperfect competition for the (static) optimal (log) price p^*_{ijt} postulates that it is a function of markup (μ_{ijt}) over marginal costs (mc_{ijt}):

$$p^*_{ijt} = \mu_{ijt} + mc_{ijt},$$

Under fairly general conditions, including separability of the firm-level demand for each product, (Amiti et al., forth., henceforth AIK) show that markups are a function of marginal costs and competitors' prices $p_{-j,t}$, so that in first differences we obtain the following pricing relation:[3]

$$\Delta p^*_{ijt} = \frac{1}{1+\Gamma}\Delta mc_{ijt} + \frac{\Gamma}{1+\Gamma}\Delta p_{-j,t}.$$

Marginal costs are generally unobservable, but under some fairly general assumptions, AIK show that they can be written as the sum of all variable input prices weighted by their respective shares in total variable costs at the firm level, plus a product-specific cost component. In our dataset, starting from 2008, we observe the following components of firms' variable costs: the annual shares of energy costs, ϕ^E_{jt}, and of imported inputs, ϕ^M_{jt}, together with monthly observations for the economy-wide price of energy Δp^E_t and the firm specific price of imported inputs Δp^M_{jt}; the total wage bill w_{jt} (and hours worked) and the total purchases of domestic inputs, v_{jt}, also at a monthly frequency, from VAT data, along with nominal sales r_{jt}.

Therefore, we can run the following regression for product price i of firm j:

$$\Delta p^*_{ijt} = \alpha + \beta^E \phi^E_{jt-1}\Delta p^E_t + \beta^M \phi^M_{jt-1}\Delta p^M_{jt} + \delta_1 \Delta \overline{p}_{i,-j,t} + \gamma X_{ijt} + \varepsilon_{ijt}$$

where X_{ijt} is a set of controls including $\Delta w_{jt}, \Delta v_{jt}$, and Δr_{jt} (see below for more details); Δp^E_t is the economy-wide price of energy and Δp^M_{jt} is the average change in the price of firm-specific imported inputs. Therefore the coefficients β^E, β^M can be interpreted as an estimate of the structural pass-through coefficient $\frac{1}{1+\Gamma}$. The other coefficients would not in general have a structural interpretation, given endogeneity and measurement errors; however, inclusion of those variables allows us to interpret the changes $\phi^E_{jt-1}\Delta p^E_t$ and $\phi^M_{jt-1}\Delta p^M_{jt}$ as orthogonal to all other controls (similarly to a Choleski ordering in a VAR in which the latter variables are ordered before Δp^E_t and Δp^M_{jt}). Given the above specification, we cannot use firm-time fixed effects because time-varying cost shares are also measured at the firm level.

This approach has nevertheless the limitation that in computing $\Delta \overline{p}_{i,-j,t}$ we cannot easily measure prices of foreign competitors (both for domestic prices and for export prices), so that the estimated β's may also reflect to some extent the elasticity of foreign competitors' prices to the energy and imported input shocks.

[3]When the demand for goods produced by multiproduct firms is not separable and has a different elasticity within the firm than across firms, then the good-specific markup cannot be easily expressed as simply a function of the prices of competitors of the same good. Conversely, the markup becomes a function of the sensitivity of the firm-specific demand for other goods, which in turn can be affected by competitors' prices in all these other markets.

4.2 Cost pass-through under price stickiness

When the above specifications are run in a sample in which prices change infrequently, the following two considerations are also important. First, including unchanged prices will bias the estimates downward (e.g. Berger and Vavra (forth.) formally show that the bias is proportional to the frequency of adjustment). This bias is present under both time-dependent and state-dependent pricing. Second, under state-dependent pricing, even conditioning on non-zero price changes, the above pass-through regression is biased by endogeneous selection into changing prices. Selection induces a positive correlation between $\phi^E_{jt-1}\Delta p^E_t$ and $\phi^M_{jt-1}\Delta p^M_{jt}$, and any idiosyncratic good-level shock; to wit: the price of a good receiving a large idiosyncratic productivity realisation Δz_{it} of the same sign as Δp^E_{jt} or Δp^M_{jt} is more likely to adjust. This selection bias is likely to be present at any horizon $t+k$ at which the price may or may not adjust. Therefore, it is important to take the extensive margin, the likelihood of price changes, into account when estimating cost pass-through at different time horizons. We take this into account by using a two step approach as in Bourguignon et al. (2007). Specifically, in the first step we model selection into price adjustment as a multinomial logit, similarly to section 3.2 above, while in the linear projections in the second step we include a "bias correction" based on the first step, among the controls X_{ijt}. We use exclusion restrictions to help in better identifying the model, by including some variables only in the multinomial logit estimation step, while excluding them from the second linear projection step, guided by theoretical considerations and the results in section 3.2 .

More in detail, in the first step we estimate a multinomial logit model of the following form:

$$\Pr\Big(Y_{ijt+k} = 1, 0, -1 \,\big|\, Z_{ijt} = x\Big) = \Phi\left(\eta Z_{ijt}\right),$$

where $Y_{i,j,t}$ is an indicator variable for positive, zero, or negative (log) price changes of good i produced by firm j between time t and $t+k$, with 0 as the base (no price change) category. Among the regressors Z_{ijt}, in addition to those we use in the second step, we include a similar set of covariates as in section 3.2; to enhance identification these variables are excluded from the linear projections in the second step. First, again as suggested by the theoretical model of Alvarez and Lippi (2014), we use the fractions of positive and negative price changes within the same firm, excluding the price change of the good we are trying to explain, together with the standard deviation of all price changes in the firm in the last 5 years, instead of the contemporaneous absolute average price change — this is due to the fact that we include also price changes of firms with only one reported product. Second, we include the fraction of positive and negative price changes in the industry, excluding the i-th good price. Furthermore, we include month fixed effects (dummies).

Figure 8 shows the dependent variable of this multinomial logit estimation, namely the fractions of cumulated zero, positive and negative price changes from one month (k=0) to 2 years (k=24). Interestingly, the fraction of a zero cumulated price change does not vanish even after 24 months; indeed as shown in the figure while decreasing over time, this fraction is above 0.5 after 12 months, and still close to 0.25 after 24 months. This implies that indeed selection bias may be potentially present even at these medium run horizons.

4.3 Dynamic selection and pass-through of costs of energy and imported input

We report results of estimating the above equation using local projections à la Jordà (2005), where the dependent variable is the cumulated price change of product i of firm j from period t to $t{+}k$, denoted $\Delta^k p_{ijt} = p_{i,j,t+k} - p_{ijt}$, conditional on it being non-zero over this time interval. On the right hand side, we include $\phi^E_{jt-1}\Delta p^E_t$ and $\phi^M_{jt-1}\Delta p^M_{jt}$ (in Danish kroner), as well as Δp_{-jt} (constructed using the first two digits of the product in the PPI database, for a total of 73 industries). Four lags of the first two regressors are also included. Furthermore, we include the above mentioned controls for firm-level costs, namely: total change in domestic purchases over last 3 months, change in hourly wage times the labor share, change in sales over last three months; we also include the monthly CPI inflation rate and the change in the Danish nominal effective (trade-weighted) exchange rate (NEER). Finally, we control for the following set of firm/product level time-invariant variables: the number of full-time equivalent employees and the number of products in the year, as well as dummies for replacement, sales, and export prices and industry fixed effects at the 2-digit level.[4] Following Bourguignon et al. (2007), we include "correction bias" terms from the first step estimation for each horizon; specifically we use variant 2 of the Dubin-McFadden approach, which does not restrict the correlation between the error terms of the selection step and linear projection step, but assumes that the conditional expectation of the latter is a linear function of know convolutions of the former.

Table 3 shows the results of the multinomial logit model for the horizon $k = 0$, which can be compared with those in Table 2, where the top panel reports results for price hikes and the bottom panel for price cuts. Observe that the sample is now shorter, starting in January 2008, because of the availability of wage data. In line with Table 2, we find strong evidence of synchronization within the firm, with very similar estimates of marginal effects of the fraction of same- and opposite-signed changes in the two specifications. While the effects of the industry-level fractions of positive and negative price changes are less strong and generally insignificant, those of the average industry price

[4]Given their computational complexity in the multinomial logit step we do not include firm-level fixed effects; we plan to explore their role in future revisions of the paper.

changes remain substantial. In further support of some degree of state-dependent pricing, we find that both the aggregate CPI and firm-level imported input prices significantly affect the probability of changing prices. While the former effects are now stronger for price decreases, the latter effects are quite symmetric between price cuts and hikes. Specifically, a decrease (increase) of 1% in the price of imported inputs raises the probability of a price cut (rise) by 0.3%. Conversely, we find that contemporaneous changes in energy prices do not significantly affect the probability of changing prices; point estimates may have even the wrong sign. Nevertheless, we find that the coefficient on $\phi^{E}_{jt-1}\Delta p^{E}_{t}$ has a ("correct") negative sign and is statistically significant for cumulated price declines over horizons k=1 to 24. The coefficient is positive for cumulated price increases from k=3, and statistically significantly from k=5. Also in line with Table 2, the effect of price changes of the opposite sign within the firm increases in a statistically significant way with the number of goods produced by firms, although by much less than the effect of same-signed price changes.

We then estimate a set of pass-through coefficients for each horizon k, $\beta^{E}_{k}, \beta^{M}_{k}$, conditional on non-zero price changes, using both our two-step procedure and standard OLS. Figure 9a reports the results for a change in energy cost, and figure 9b for import prices, respectively. The red solid lines represent the second-stage coefficient estimates at each horizon, the blue line the corresponding OLS estimates and the dark and light grey areas indicate 68% and 95% HAC robust confidence bands for the OLS estimate.[5]

The following results emerge. First, OLS and bias corrected point estimates are very similar, with the latter never falling out of the OLS confidence bands; moreover, point estimates are remarkably close for energy price shocks, and differ only a little more for import price shocks, where OLS are larger until around 6 months and then lower after 12 months. Therefore, even though we find that the bias correction terms in the second stage are significantly different from zero, at least after the first couple of months, state-dependence in the extensive margin does not translate into an economically significant OLS bias. Second, both cost shocks brings about a very persistent increase in prices, which are significantly affected even after 2 years, though their estimated response is basically stable after 6 months. Third, the two cost pass-through coefficients are however very different over time. The dynamic response to energy prices is small and insignificantly different from zero on impact and in the first couple of months after the shock, but quickly builds over time, reaching around 0.5 after 4 months and stabilizing around 1 after 6 months until 24 months. Conversely, the dynamic response to firm-specific import costs is positive and significantly different from zero

[5]Standard errors are clustered at the firm level. We don't report standard errors for the second stage in this version of the paper as it is not straightforward to take into account the uncertainty due to generated regressors in the second step. Nevertheless, those we have computed are still valid to test the null of zero regressor coefficients, which is rejected for all horizons $k > 3$, similarly to OLS estimates.

already on impact, at around 0.3, and pretty much stabilizes around this value at all horizons up to 24 months. Therefore, we find that estimates of medium run pass-through after 12-24 months are very different across these two cost shocks.

There could be several sources of the difference in the dynamic response of the two shocks, which we investigate next. First, it can be due to different persistence of the two shocks; however, both price changes seem to be well described as very transitory process, so that indeed exclusion of their lags does not materially affect our results. Second, the difference can be due to the fact that while we include the NEER change, we don't control for possible aggregate effect of energy prices. However, including the latter also does not materially affect results, as shown in figure 10. Finally, a source of the different estimates could be that there is more predictability in energy prices, beyond its own lags, for instance because they are persistently affected by other lagged macroeconomics variables. Therefore, we run a specification where we replace Δp_t^E with the series of oil supply shocks recently estimated by Baumeister and Hamilton (2019); arguably these shocks should be unforecastable by construction. Since we do not have measures of firm-specific oil intensity, we still interact the shocks with firm-specific energy shares ϕ_{jt-1}^E. However, energy intensity is likely to be much larger than oil intensity, so this could be a potential source of downward bias in our estimates. As shown in figure 11, now the response to oil shocks builds up more gradually and indeed stabilizes at a lower level, around 0.7-0.8 after 12 months — point estimates even decline somehow after 18 months. Despite lower point estimates, medium run pass-through is still much higher than that for import prices: OLS confidence bands contain 1 for horizons between 14 and 22 months but don't contain 0.4 already after 6 months up to 23 months. The quantitatively different medium run pass-through of the two cost shocks seems thus a fairly robust feature of the data.

We conclude this section by looking at any heterogeneity across firms with a low and a high number of products. For this purpose, we split the sample at the median value of 5 products. One question is whether firms with a lower number of products may display a more state-dependent behavior conditional on the cost shocks we examine, because of lower complementarities in changing prices. Another question is whether, conditional on changing prices, price adjustment is different across firms with a different number of products, because of for instance firms with more products are more likely to internalize the effects of changing a given price on the other markets.

Results for the two samples of firms are reported in figures 12 and 13. As shown by the first row of the figure, point estimates of the pass-through of firms with a low and a higher number of products, reported respectively in columns (a) and (b), are remarkably similar for energy cost shocks. Conversely, as shown by the second row of the figure, pass-through estimates substantially differ for import price shocks, with those of firms with less than 5 products in column (a) sensibly higher than those of the other firms in column (b), but still substantially below pass-through for energy shocks. This result is consistent with some kind of link between the lower pass-through of imported input prices and the firm-level number of products (or other features correlated with it, such as size). Nevertheless, it is not straightforward to relate the different price adjustment to energy and import costs even in the light of standard theories of price-setting by multiproduct firms.

5 Conclusions

[to be added]

References

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.

Álvarez, L. J., E. Dhyne, M. M. Hoeberichts, C. Kwapil, H. L. Bihan, P. Lünnemann, F. Martins, R. Sabbatini, H. Stahl, P. Vermeulen, and J. Vilmunen (2006): "Sticky Prices in the Euro Area: A Summary of New Micro Evidence," *Journal of the European Economic Association*, 4, 575–584.

Amiti, M., O. Itskhoki, and J. Konings (forth.): "International Shocks, Variable Markups and Domestic Prices," *Review of Economic Studies*.

Baumeister, C. and J. D. Hamilton (2019): "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks," *American Economic Review*, 109, 1873–1910.

Berardi, N., E. Gautier, and H. Le Bihan (2015): "More Facts about Prices: France Before and During the Great Recession," *Journal of Money, Credit and Banking*, 47, 1465–1502.

Berger, D. and J. Vavra (forth.): "Shocks vs. Responsiveness: What Drives Time-Varying Dispersion?" *Journal of Political Economy*, 103, 60–82.

Bhattarai, S. and R. Schoenle (2014): "Multiproduct Firms and Price-setting: Theory and Evidence from U.S. Producer Prices," *Journal of Monetary Economics*, 66, 178–192.

Bils, M. and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985.

Bourguignon, F., M. Fournier, and M. Gurgang (2007): "Selection Bias Corrections Based on the Multinomial Logit Model: Monte Carlo Comparisons," *Jounral of Economic Surveys*, 21, 174–205.

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

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464.

Olivei, G. and S. Tenreyro (2007): "The Timing of Monetary Policy Shocks," *American Economic Review*, 97, 636–663.

Statistics Denmark (2017): "Documentation of statistics for Accounts Statistics for Non-Agricultural Private Sector," Tech. rep.

——— (2018): "Documentation of statistics for Purchases and Sales by Firms," Tech. rep.

——— (2019): "Documentation of statistics for Producer and Import Price Index for Commodities," Tech. rep.

Vermeulen, P., D. Dias, M. Dossche, E. Gautier, I. Hernando, R. Sabbatini, and H. Stahl (2012): "Price Setting in the Euro Area: Some Stylised Facts from Individual Producer Price Data," *Journal of Money, Credit and Banking*, 44, 1631–1650.

Wulfsberg, F. (2016): "Inflation and Price Adjustments: Micro Evidence from Norwegian Consumer Prices 1975-2004," *American Economic Journal: Macroeconomics*, 8, 175–194.

A Data description

A.1 Producer price micro data

We use the confidential microdata underlying the Danish producer and import price index for commodities compiled by Statistics Denmark. It covers the time period from January 1993 until June 2017. The producer and import price index for commodities is based on approximately 6,400 prices per month across 1,050 different commodities, reported by selected producers and importers in Denmark, see also Statistics Denmark (2019). Approximately 3,500 prices are used for calculating the producer price index, approximately 2,900 prices are used for calculating the import price index. The most important firms within selected areas are requested to report prices in order to ensure that the producer and import price index covers at least 70 percent of Danish production and imports.

The population covers all commodities that are imported or produced in Denmark for the domestic market or export, with the exception of some well-defined exemptions. Some commodities are not included because the turnover is too small and some commodities are not included because of the nature of the commodities.

Statistics Denmark undertakes great efforts to adjust for quality changes and product substitutions so that only true price changes are measured. When a product is substituted, Statistics Denmark re-computes the base price, and therefore we are able to identify these replacements. They constitute only 0.7 per cent of all prices changes (including zero price changes) and 0.8 per cent of all non-zero price changes. We include these in the baseline results we report, but control for identified product replacements in regressions. Goods are defined relatively narrowly in our dataset, as products are classified using the 8-digit combined nomenclature (CN). The first 6 digits of the CN codes correspond to the World Harmonized System (HS). We address breaks in product classifications by identifying changes in product codes within a firm which do not lead to a change in the price. The vast majority of identified breaks coincides with the months where Statistics Denmark re-defines product categories. The breaks constitutes only 0.04 per cent of all price changes (including zero changes), and *per construction* 0 per cent of all non-zero price changes. Similar to product replacements, we include these incidents in the baseline results we report, but control for identified breaks in regressions.

The prices used for the index are actual prices, which means that the prices must include all possible discounts. Therefore, list prices do not apply unless the prices never include discounts. A distinction is made between the prices of imported commodities and the prices of commodities for the domestic market or the export market:

- Imported commodities: Actual transaction prices (in some cases transfer prices) c.i.f. excluding all duties and taxes on the goods as far as possible on the 15th day of the month. For the firms reporting import prices, we calculate a firm-level import price index using the equally weighted average log differences in each month month.

- Danish commodities for the domestic market or export: Actual transaction price (in a few cases transfer prices) ex producer excluding VAT and excise duties as far as possible on the 15th of the month.

There are 6700 unique goods recorded and some point in the sample. The number of recorded goods in the average month is 2600, the highest (lowest) amount of unique goods is 3500 (1700). On average, the price of a good is reported for 115 subsequent months. For 89% of goods we observe the price for every month in at least two subsequent years, for almost half the goods we have at least 10 years of price reporting, and for more than 500 goods we observe the price for every month between 1993 and 2017. The amount of goods entering and leaving the sample tend to peak only in months resampling.

A.2 Firm registers

We combine the pricing data with annual firm-level data from Statistics Denmark's accounts statistics for the Danish business sector in the period from 1996 to 2016. A firm is identified at the enterprise level, i.e. the legal unit, see also Statistics Denmark (2017). The primary industries, the financial sector and the public sector are excluded.

There are 1,140 unique firms in the dataset, of which 640 are linked to a price observation in the average month. The number of firms which report prices fluctuates between 520 and 755 (in 2007). We match 80% of all price observations to a firms balance sheet over the whole sample. For the period after 2005, more than 90% of price observations can be linked to firm observations.

Income statement items we use include total sales and profits, from which we impute total cost. Firms report the total amount spent on purchasing energy throughout a year. The mean (median) spending on energy as a share of total cost is 1.7% (1.09%). Furthermore, we observe the number of employees in full-time equivalents, firm age (for a subsample of 81% of the firms), as well as expenditure on imported goods. We calculate the latter as a share of approximated total cost: The mean (median) import intensity is 27% (23.1%).

A.2.1 Monthly sales and purchases

For all firms covered by the Danish VAT system, we have information on purchases and sales, see also Statistics Denmark (2018). The data contains information on total sales and total purchases from 2001 to 2017, with the category of imported purchases reported separately starting in 2002.

The monthly frequency of this dataset allows us to leverage the high frequency of the pricing data. However, some firms to not report on a monthly basis, whereas the annual turnover of a firm determines its VAT declaration frequency. The frequency is monthly if the amount exceeds DKK 50 million, quarterly in the interval between DKK 5 million and DKK 50 million, and half-yearly if it is less than DKK 5 million (and above DKK 50,000). Quarterly and semi-annual data are recalculated and spread onto months by Statistics Denmark using information for firms with monthly VAT reporting in the same industry (at the DB-127 level).

We match about two thirds of good price observations to an observation in this FIKS registry, which is due to the shorter time span of the latter. Coverage for the period starting in 2002 is almost complete.

A.2.2 Firm payrolls

Danish firms register hours worked by and total compensation of employees in the tax authority's *e-Indkomst*. The BFL registry is available from 2008 onwards, after which we match payroll information to all observations in the base price dataset based on anonymised firm identifiers. We compute the average hourly wage at the firm level as well as the labor share in the firm's cost structure from the firm registers, which is 25% on average.

B Tables

Table 1: Summary statistics

	All	1	1-3	3-5	5-7	7+
No. of firms	1140	146	548	231	128	87
Mean employment (FTE)	560.6	78.8	153.9	259.0	254.9	1508.9
Median employment (FTE)	155.0	42.6	65.5	138.6	148.3	483.1
Mean employment per good	68.2	78.8	62.1	63.4	45.0	94.1
Median employment per good	33.2	42.6	26.2	34.8	25.5	48.5
Mean age	30.5	27.8	27.1	31.0	29.5	34.3
Median age	26.0	25.0	25.0	28.0	24.0	29.0
Share of total prices	100.0	2.1	20.2	22.8	16.7	38.1
Mean no. of products	8.4	1.0	2.6	4.1	5.8	19.0
Std. err. no. of products	12.4	0.0	0.5	0.6	0.5	18.6
25th percentile	3.0	1.0	2.2	3.6	5.4	8.8
Median	5.0	1.0	3.0	4.0	5.8	11.5
75th percentile	8.3	1.0	3.0	4.7	6.0	16.6

Note: Summary statistics on distribution of firms and prices across distinct bins of single- and multiproduct firms.

Table 2: Multinomial logit, price synchronization

	All	1-3	3-5	5-7	7+
Marg. effect on probability of decrease					
Fraction of pos. price changes in firm	2.44∗∗∗	1.85∗∗∗	2.26∗∗∗	2.12∗∗∗	2.76∗∗∗
	(0.00)	(0.03)	(0.00)	(0.05)	(0.06)
Fraction of neg. price changes in firm	3.95∗∗∗	2.57∗∗∗	3.87∗∗∗	4.01∗∗∗	5.26∗∗∗
	(0.00)	(0.03)	(0.00)	(0.05)	(0.05)
Fraction of pos. price changes in industry	0.14∗∗∗	0.038	-0.117	0.040	0.22∗∗∗
	(0.03)	(0.07)	(0.06)	(0.07)	(0.05)
Fraction of neg. price changes in industry	0.41∗∗∗	0.62∗∗∗	0.58∗∗∗	0.44∗∗∗	0.26∗∗∗
	(0.02)	(0.07)	(0.05)	(0.06)	(0.05)
Avg. price change in firm, excl. good	-0.09∗∗∗	-0.066	-0.04∗∗∗	-0.038	-0.261∗∗
	(0.00)	(0.04)	(0.00)	(0.04)	(0.09)
Avg. abs. price change in firm, excl. good	0.02∗∗∗	-0.004	0.04∗∗∗	0.041	-0.079
	(0.00)	(0.04)	(0.00)	(0.04)	(0.08)
Avg. price change in industry, excl. firm	-0.25∗∗∗	-0.22∗∗∗	-0.098	-0.137∗	-0.34∗∗∗
	(0.03)	(0.05)	(0.06)	(0.07)	(0.07)
CPI, log difference	-0.460∗∗	-0.536∗	-0.627∗	-0.612∗	0.170
	(0.14)	(0.26)	(0.28)	(0.29)	(0.30)
Marg. effect on probability of increase					
Fraction of pos. price changes in firm	6.18∗∗∗	4.37∗∗∗	5.96∗∗∗	6.31∗∗∗	8.30∗∗∗
	(0.00)	(0.03)	(0.00)	(0.05)	(0.05)
Fraction of neg. price changes in firm	2.79∗∗∗	2.07∗∗∗	2.59∗∗∗	2.65∗∗∗	2.81∗∗∗
	(0.00)	(0.03)	(0.00)	(0.06)	(0.06)
Fraction of pos. price changes in industry	0.35∗∗∗	0.46∗∗∗	0.51∗∗∗	0.214∗∗	0.26∗∗∗
	(0.03)	(0.07)	(0.07)	(0.08)	(0.06)
Fraction of neg. price changes in industry	0.044	0.053	-0.125∗	-0.134	0.153∗
	(0.02)	(0.08)	(0.06)	(0.07)	(0.06)
Avg. price change in firm, excl. good	0.10∗∗∗	0.095∗	0.04∗∗∗	-0.015	0.38∗∗∗
	(0.00)	(0.04)	(0.00)	(0.05)	(0.09)
Avg. abs. price change in firm, excl. good	-0.02∗∗∗	0.005	0.05∗∗∗	0.015	-0.251∗∗
	(0.00)	(0.05)	(0.00)	(0.05)	(0.09)
Avg. price change in industry, excl. firm	0.27∗∗∗	0.24∗∗∗	0.154∗∗	0.101	0.38∗∗∗
	(0.03)	(0.05)	(0.05)	(0.06)	(0.06)
CPI, log difference	0.69∗∗∗	0.695∗	0.960∗∗	0.261	0.427
	(0.16)	(0.28)	(0.31)	(0.32)	(0.33)
N	599310	157652	151956	112730	161751
R2	0.404	0.445	0.437	0.473	0.369

Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Note: Marginal effects (in percentage points) of a one standard deviation change in the regressor from the mean on the probability of increasing and decreasing the price relative to not changing the price. Exception: 1% in CPI inflation. Standard errors in parentheses. Further controls (not reported): Firm size, dummies for product replacement, sales, and exports, month fixed effects.

Table 3: Multinomial logit, first stage results

	All	1-5	5+
Marg. effect on probability of price increase			
Fraction of pos. price changes in firm	6.34***	5.27***	7.83***
	(0.04)	(0.04)	(0.06)
Fraction of neg. price changes in firm	2.74***	2.39***	2.87***
	(0.04)	(0.05)	(0.07)
Fraction of pos. price changes in industry	0.080	0.333**	0.037
	(0.06)	(0.11)	(0.08)
Fraction of neg. price changes in industry	-0.202**	-0.43***	-0.104
	(0.06)	(0.13)	(0.08)
Avg. price change in industry, excl. firm	0.14***	0.111**	0.114*
	(0.03)	(0.04)	(0.05)
Energy price change x lagged energy cost share	-0.371	-0.959*	0.607
	(0.38)	(0.48)	(0.61)
Import price change x lagged import cost share	0.28***	0.48***	0.108*
	(0.04)	(0.06)	(0.05)
CPI, log difference	0.557*	0.703	0.141
	(0.28)	(0.41)	(0.39)
Marg. effect on probability of price decrease			
Fraction of pos. price changes in firm	2.28***	1.93***	2.43***
	(0.04)	(0.04)	(0.06)
Fraction of neg. price changes in firm	4.09***	3.38***	5.22***
	(0.03)	(0.04)	(0.06)
Fraction of pos. price changes in industry	-0.25***	-0.61***	-0.027
	(0.06)	(0.11)	(0.07)
Fraction of neg. price changes in industry	-0.073	0.015	0.047
	(0.06)	(0.12)	(0.07)
Avg. price change in industry, excl. firm	-0.15***	-0.13***	-0.138**
	(0.03)	(0.04)	(0.04)
Energy price change x lagged energy cost share	-0.172	0.084	-0.492
	(0.34)	(0.43)	(0.55)
Import price change x lagged import cost share	-0.29***	-0.46***	-0.138**
	(0.04)	(0.06)	(0.05)
CPI, log difference	-1.00***	-1.33***	-0.516
	(0.27)	(0.39)	(0.37)
N	272372	128693	143679
R2	0.405	0.442	0.383

Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Note: Marginal effects (in percentage points) on increasing and decreasing the price relative to not changing the price. The variables of within firm and industry synchronization show the effect of a one standard deviation change of the regressor around its mean; the other variables do so for a 1% change in the input variable. Columns (2) and (3) split the sample along the median number of products. Standard errors in parentheses. The change in firm sales and domestic purchases over the past quarter as well as the change in the hourly wage rate interacted with the labor share account for other firm-level cost components (not reported). Further controls: Log firm size, dummies for product replacement, sales, exported and energy products, the change in the nominal effective exchange rate, month fixed effects.

C Figures

Figure 1: Average path of prices conditional on non-zero change in period 0

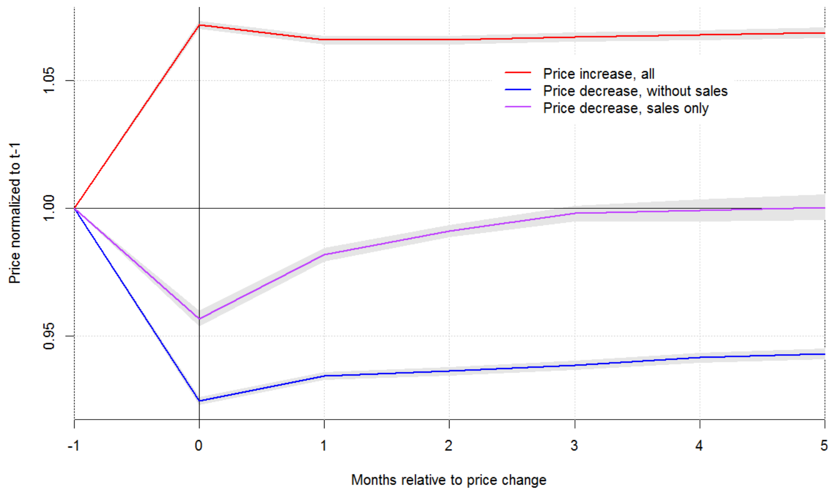

Note: Sales are identified using filter B from Nakamura and Steinsson (2008), i.e. we define as a sale all price decreases at time 0 that are fully reverted after 1, 2 or 3 months.

Figure 2: Frequency of price change

(a) Mean frequency of price change (b) Median frequency of price change

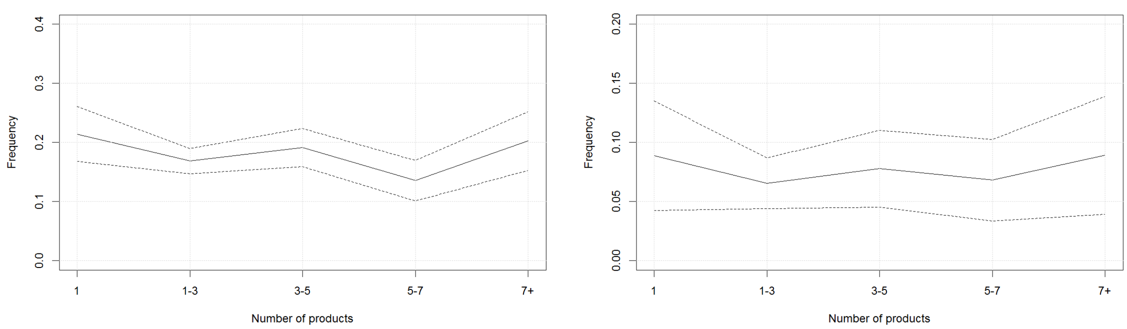

Note: Following Bhattarai and Schoenle (2014), we calculate the mean frequency of price changes with 95% bands. For each bin based on the number of goods, we first compute the frequency of price change at the firm-good level. Then, we compute the median frequency of price changes across goods at the firm level. Finally, we report the mean (a) and median (b) across firms in a given bin. Error bands are computed as ± 1.96 std.error across firms.

Figure 3: Mean fraction of positive price changes

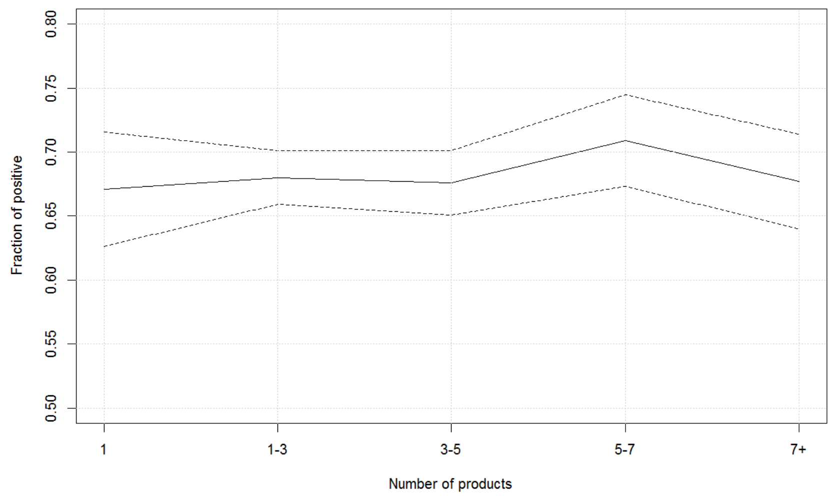

Note: Following Bhattarai and Schoenle (2014), we calculate the mean fraction of positive price changes with 95% bands. For each bin based on the number of goods, we first compute the number of strictly positive price changes over all non-zero price changes at the firm level. Then, we report the mean across firms in a given bin. Error bands are computed as ± 1.96 std.error across firms.

Figure 4: Seasonality of frequency and size of price changes

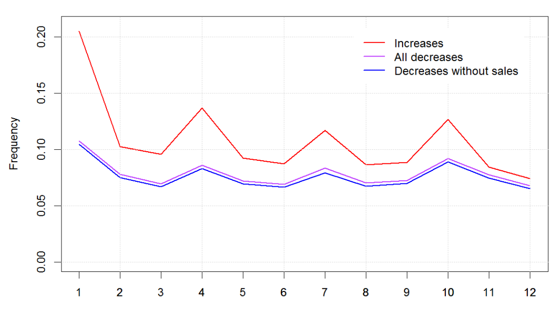
(a) Frequency of price changes by month

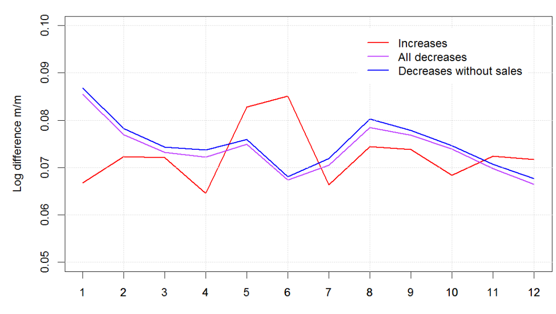
(b) Absolute size of price changes by month

Note: Mean frequency of price changes of firms per month of the year. Price changes (particularly increases) are most frequent in January, with local peaks at the first month of any quarter. Sales remain quantitatively minor and do not have a sesonal pattern different from regular price decreases.

Figure 5: Mean absolute size of price changes

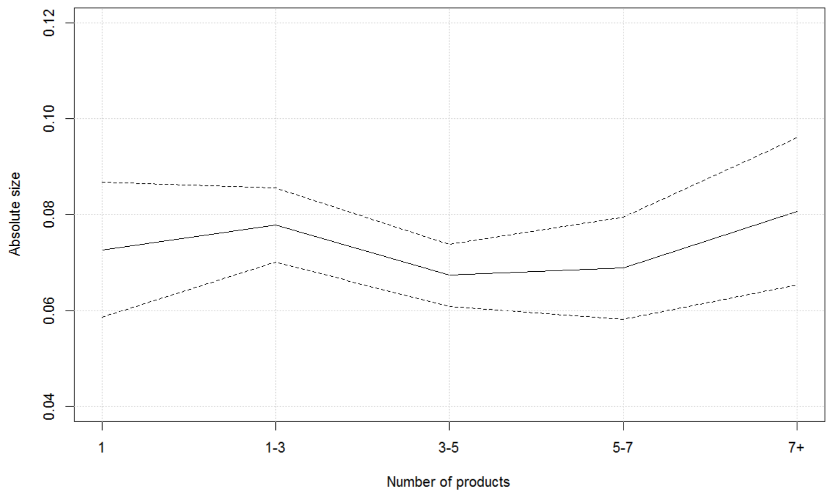

Note: Following Bhattarai and Schoenle (2014), we calculate the mean absolute size of price changes with 95% bands. For each bin based on the number of goods, we first compute the mean absolute size of price changes at the firm-good level. Then, we compute the median absolute size of of price changes across goods at the firm level. Finally, we report the mean across firms in a given bin. Error bands are computed as ± 1.96 std.error across firms.

Figure 6: Mean size of positive and negative price changes

(a) Mean size of positive price changes

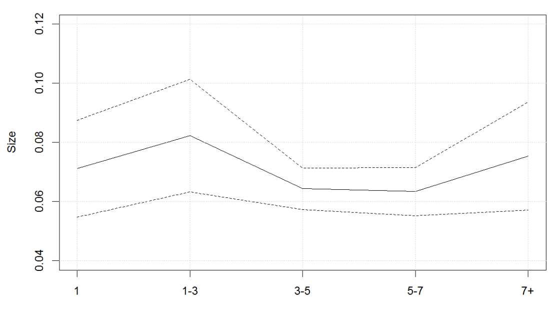

(b) Mean size of negative price changes

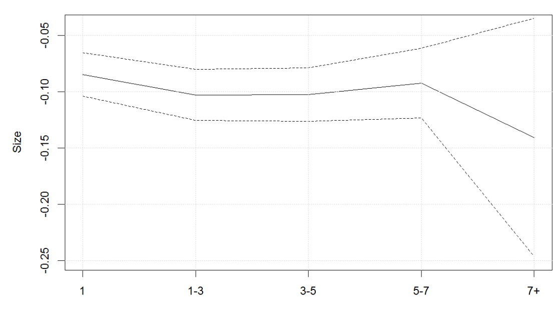

Note: Following Bhattarai and Schoenle (2014), we calculate the mean size of positive and negative price changes with 95% bands. For each bin based on the number of goods, we first compute the mean size of positive and negative price changes, separately, at the firm-good level. Then, we compute the median size of of positive and negative price changes across goods at the firm level. Finally, we report the mean across firms in a given bin. Error bands are computed as ± 1.96 std.error across firms.

Figure 7: Histograms of standardized price changes

(a) Firms with 1 product

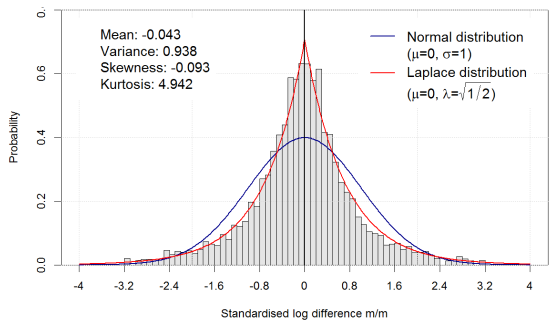

(b) Firms with 1-3 products

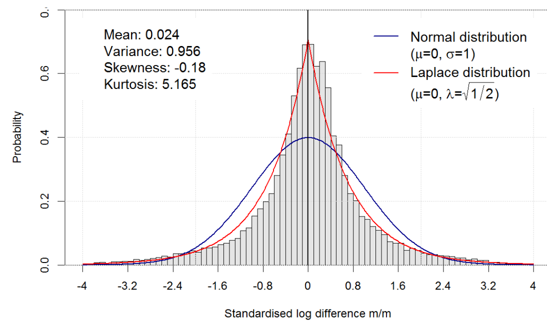

(c) Firms with 3-5 products

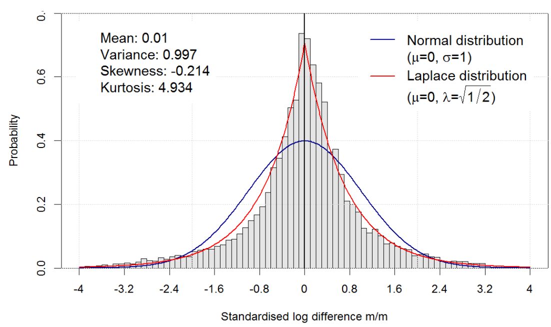

(d) Firms with 5-7 products

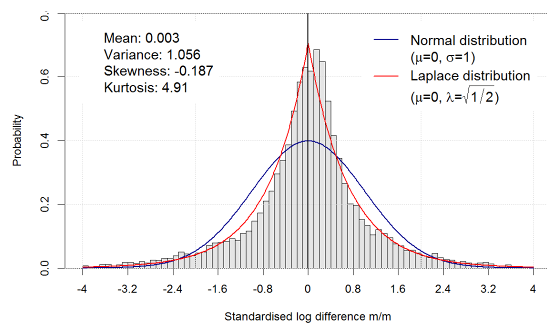

(e) Firms with more than 7 products

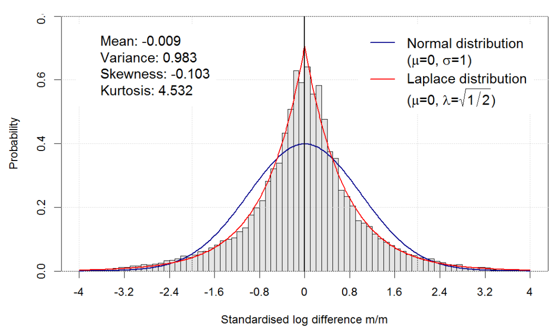

Note: Price changes are the log difference in price, standardized by good category. Price changes equal to zero are discarded. Following Alvarez et al. (2016), we drop price changes whose absolute value is smaller than 0.1 percent to avoid measurement error.

Figure 8: Fraction of cumulated price changes from 1 to 24 months

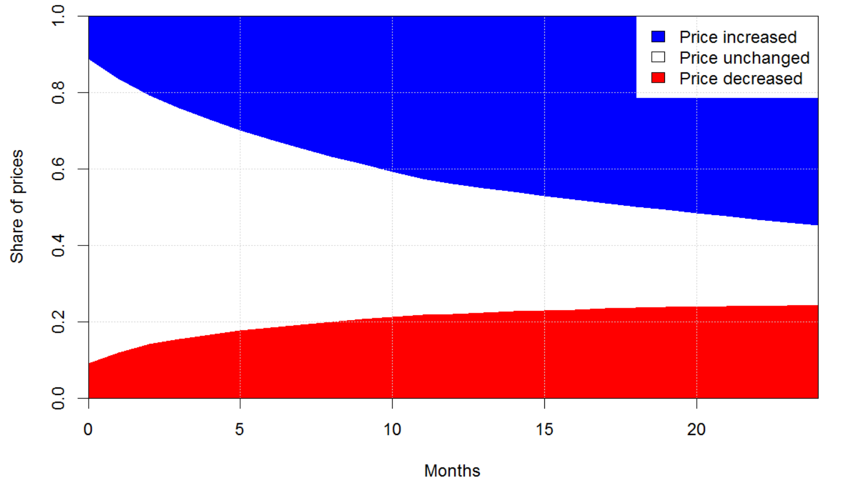

Note: For every horizon k, this figure depicts the probability of having changed (increased or decreased) the price between month 0 and k.

Figure 9: Pass-through coefficients from local projections (Baseline)

(a) Energy cost pass-through

(b) Import price pass-through

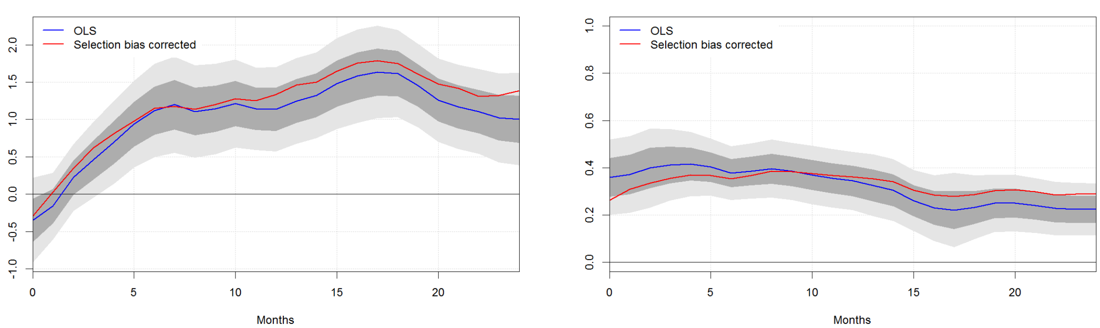

Note: Estimated coefficients of a change in input cost interacted with the share of the input of total cost, conditional on that the price has changed. Blue solid lines show point estimates of local projections of the cumulative price change until k months ahead using OLS, red solid lines correct for a possible selection bias. 68% and 95% confidence bands in grey. Further controls (not reported): Lagged values in the input price change, (absolute) average price change at the firm level excluding the good, the average price change of competitors excluding the firm, firm size, dummies for product replacement, sales, and exports, time fixed effects.

Figure 10: Pass-through (including aggregate change in energy price)

(a) Energy cost pass-through

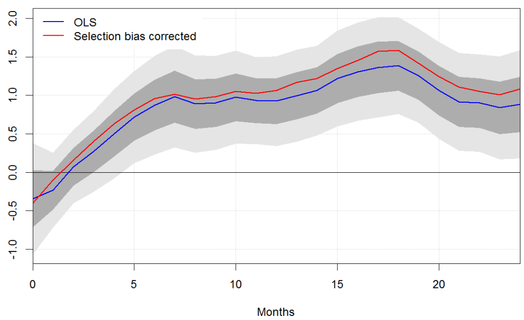

(b) Import price pass-through

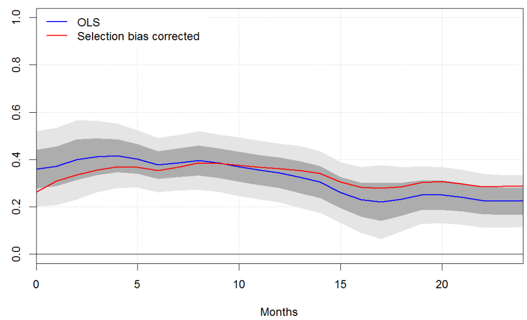

Note: Estimated coefficients of a change in input cost interacted with the share of the input of total cost, conditional on that the price has changed. Blue solid lines show point estimates of local projections of the cumulative price change until k months ahead using OLS, red solid lines correct for a possible selection bias. 68% and 95% confidence bands in grey. Further controls (not reported): Change in the economy-wide price of energy, lagged values in the input price change, (absolute) average price change at the firm level excluding the good, the average price change of competitors excluding the firm, firm size, dummies for product replacement, sales, and exports, time fixed effects.

Figure 11: Pass-through (Oil supply shock)

(a) Oil price surprise pass-through

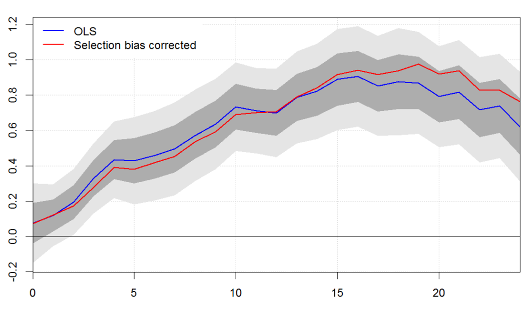

(b) Import price pass-through

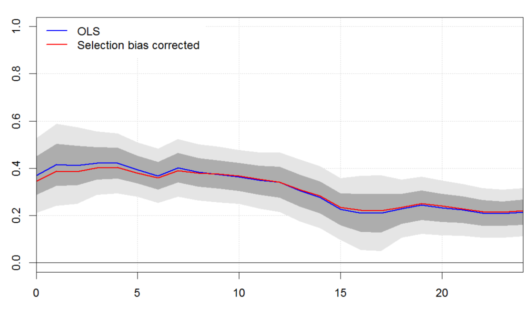

Note: Estimated coefficients of a change in input cost interacted with the share of the input of total cost, conditional on that the price has changed. The input cost in panel (a) is the oil supply shock series provided by Baumeister and Hamilton (2019). Blue solid lines show point estimates of local projections of the cumulative price change until k months ahead using OLS, red solid lines correct for a possible selection bias. 68% and 95% confidence bands in grey. Further controls (not reported): Lagged values in the input price change, (absolute) average price change at the firm level excluding the good, the average price change of competitors excluding the firm, firm size, dummies for product replacement, sales, and exports, time fixed effects.

Figure 12: Pass-through of energy cost shocks by number of products

(a) Less than 5 products (b) More than 5 products

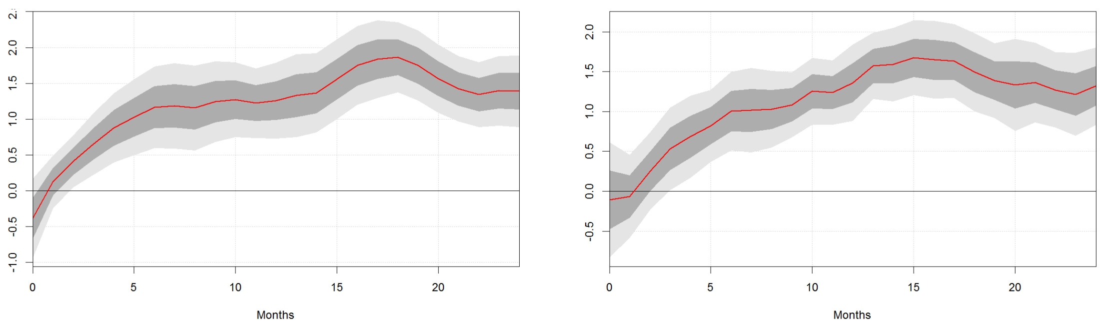

Figure 13: Pass-through of import cost shocks by number of products

(a) Less than 5 products (b) More than 5 products

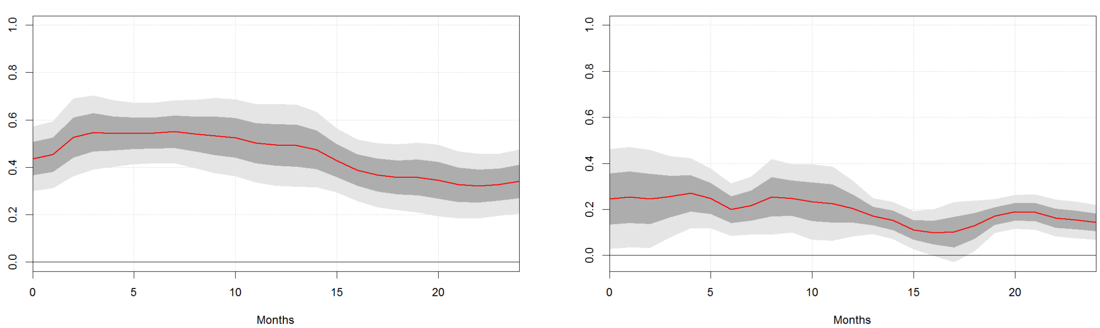

Note: Estimated coefficients of a change in input cost interacted with the share of the input of total cost, conditional on that the price has changed. Red solid lines show point estimates of local projections of the cumulative price change until k months ahead. 68% and 95% confidence bands in grey. Further controls (not reported): Lagged values in the input price change, (absolute) average price change at the firm level excluding the good, the average price change of competitors excluding the firm, firm size, dummies for product replacement, sales, and exports, time fixed effects.