Pricing Behaviour and Menu Costs in Multi-product Firms

Economica

Economica (2022)
doi:10.1111/ecca.12415

Pricing Behaviour and Menu Costs in Multi-product Firms

By WILKO LETTERIE* and ØIVIND A. NILSEN†

**Maastricht University* *†Norwegian School of Economics*

Final version received 8 February 2022.

This paper investigates the micro-foundations of pricing behaviour using monthly producer prices for Norwegian multi-product firms. We find both infrequent and many small price changes together with a high degree of within-firm synchronization. This points at fixed menu costs featuring scope economies, at additional linear and convex price adjustment costs, and at the presence of firm-specific shocks. The structural estimates and a simulation support the view that in order to understand pricing behaviour and the effectiveness of monetary policy, the analysis of multi-product firms and a richer price adjustment technology in the intermediate goods sector is valuable.

INTRODUCTION

In economics, the phenomenon of price rigidity has featured prominently on the research agenda for a long time. The classical menu costs are considered theoretically by Sheshinski and Weiss (1977, 1983). Typically, such physical costs are independent of the size of the price changes (Levy *et al.* 1997). In more recent years, models including menu costs featuring economies of scope have received substantial attention (Sheshinski and Weiss 1992; Midrigan 2011; Alvarez and Lippi 2014; Bhattarai and Schoenle 2014; Alvarez *et al.* 2016; Yang 2019; Stella 2020, Bonomo *et al.* 2020). In such a setting, there is a total fixed menu cost that is always incurred when a multi-product firm adjusts at least one price. A firm therefore has an incentive to synchronize its price changes if the menu cost function features scope economies. The attractiveness of a model featuring such scope economies stems from its ability to also explain two other important properties of price change data: infrequent price change and many small price changes in multi-product firms—a more realistic description of real-world firms than describing them as single-product firms.

Fixed price adjustment costs—menu costs—are meant to capture physical adjustment costs related to, for instance, producing new price lists, monthly supplemental price sheets, and informing and convincing interested parties (see Levy *et al.* 1997). Assuming the price adjustment costs to be fixed might be related to the fact that such physical costs are relatively easy to measure. However, as pointed out by Blinder *et al.* (1998, p. 522), Wolman (2007, p. 543) and more recently Tsoukis *et al.* (2011, p. 741), there are also implicit costs resulting from the unfavourable reaction of customers to large prices changes. In such a setting, processes related to changing the prices are costly, not the price change as such. Some of the costs related to changing prices depend on the size of the price adjustment.[1] The managerial costs related to decision-making and internal firm communications increase for larger price changes. The firm is also likely to incur higher costs of negotiation and communication with customers. Firms could also be reluctant to change prices due to competitive forces, especially when large price changes are involved. If customer demand is elastic, then a price increase implies a reduction of demand, and price reductions increase the risk of price wars. Zbaracki *et al.* (2004) observe that customer costs constitute close to 75% of total price adjustment costs, whereas managerial 'thinking' costs represent less than 25%, while the typical menu costs are rather unimportant.

On the other hand, using qualitative information, Blinder *et al.* (1998) claim to find very little support for the existence of proportional price adjustment costs. However, a majority of the empirical studies in this literature focus on supermarkets and retailers. Wolman (2007, p. 545) argues that in this industry, firms can observe prices of their competitors easily, and managerial costs potentially are therefore less important in the price adjustment process. Industrial corporations that produce specialized products often do not have easy access to information about competitors' prices. As we investigate firms in manufacturing industries, and in light of the previous discussion, it is still desirable to investigate the functional form of price adjustment costs.

We take advantage of the data used to construct the Producer Price Index (PPI) in Norway, which concern the most important prices for the manufacturing industries. The results should therefore be considered generalizable and relevant. The data consist of monthly observations over several years. Such high frequency data make it easier to uncover important price change moments related to inaction, the size of price changes, and within-firm price synchronization. This detail makes it possible to identify the structural parameters of the firm's optimization problem micro-econometrically.

Our study contributes to the literature in several ways. First, we specify a rather parsimonious model based on a very general price adjustment cost specification, including scope economies and proportional adjustment costs. This model is capable of predicting intermittent price changes within a firm, and that price changes can be both small and large. It also allows for price coordination, in line with Midrigan (2011), Alvarez and Lippi (2014), and Bhattarai and Schoenle (2014), due to scope advantages. Second, our empirical findings support the presence of scope economies in fixed menu costs but also point towards the relevance of proportional adjustment costs. Third, our analysis also contributes to the understanding of self-selection of firms into changing prices. The degree to which this phenomenon occurs affects non-neutrality of monetary policy (Golosov and Lucas 2007; Midrigan 2011; Carlsson 2017). Based on the results, we argue that economies of scope reduce firms self-selecting into a price change regime. The reason is that economies of scope in the pricing technology reduce the responsiveness of a single price to a change in its fundamentals. As a product price change also depends on whether it is beneficial to change the other prices set by the same firm, state dependency of price changes decreases. This pricing property implies that then monetary policy is likely to be non-neutral. Fourth, compared to most other empirical studies in which the number of products that a firm manufactures is predefined, we present an empirical approach that can exploit the richness of our data by incorporating each firm's number of products. Finally, our model is rather simple provided that the assumptions hold, and it can therefore be estimated with a rather straightforward and transparent estimation technique employing a maximum likelihood routine that incorporates a latent class approach.

This paper continues as follows. In Section I we present the data. The model is developed in Section II. The estimation method is depicted in Section III. We present the empirical results in Section IV. In Section V, we provide a simulation of our model, and we give some concluding remarks in Section VI.

I. The Data

The survey used to construct the Producer Price Index (PPI) provides monthly price observations. If Statistics Norway (SSB) regards a subset of the products to be important to obtain an accurate estimate of the price index, then data will be requested for these products only.[2] The selection of respondents is updated on a regular basis, in order to make sure that the indices are being kept relevant continuously compared to the development of the Norwegian economy (SSB 2013). Compulsory participation ensures a high response from the questioned firms. The gathered data collected through electronic reporting are subject to several controls aiming to identify extreme values and mistyping. Thus the data are of very high quality. The monthly prices are merged with annual firm-level information, using the firm identifiers in the price survey data. For all firms, there are a number of variables related to their economic activity, including employment numbers, wages and the like. The monthly frequency, long periods of consecutive price observations, multiple products per producer, securitization of the data by SSB, inclusion of non-listed firms and the possibility of linking our price data to cost data make our data quite unique.

A firm in our sample may operate on both domestic and export markets. We record only domestic prices to avoid that our results are driven by exchange rate changes and competitive forces on international markets. We focus on multi-product firms, and single-product firms are therefore disregarded from our final sample. The main reason for this choice is that we want to avoid a wrong categorization of firms. A firm for which one product price is observed might be a single-product firm, but it may also be a multi-product firm for which not all product price changes have been collected. In addition, we exclude multi-plant firms to ensure that pricing decisions are made at the level that we analyse.[3] Our final dataset covers the period 2004–9. After having excluded sectors producing capital goods—goods that may be rather different from other goods produced in the manufacturing industry—the number of observations in our dataset is 39,082. The numbers of firms, products and (two-digit NACE) sectors are 222, 855 and 16, respectively. On average, a firm provides information on about five products in the actual data. In an Online Appendix, we provide some more information on the data construction. Additionally, more detailed descriptives of our data can be found in Table A1 in the Appendix at the end of the paper.

Figure 1 shows the distribution of the monthly price changes. We see a mass-point of zero price changes (77% of the observations) consistent with findings in the established literature.[4] One potential explanation for such a mass-point could be scope economies in menu costs, but may also point at non-convex menu costs, as it is unlikely that shocks are absent.

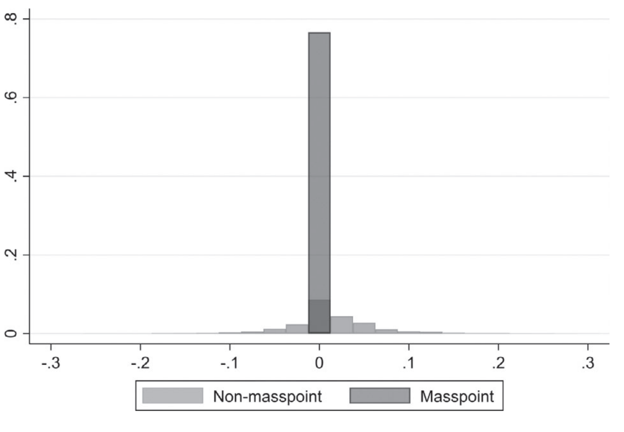

FIGURE 1. Monthly price changes ($\Delta P/P$). *Notes*: On the vertical axis, the frequency is depicted. The horizontal axis denotes the size of the price change rate. The observations are truncated at −0.30 and 0.30. This excludes 0.2% of the observations. The mass-point bar in dark grey denotes the share of observations with no price change from previous month (in total, 77% of the price changes are exactly equal to zero). The light grey denotes everything else, so excluding price changes equal to precisely zero. See also Table A1 in the Appendix for more descriptive statistics.

We provide three additional pieces of evidence shedding light on the within-firm price-change coordination.

First, Table 1 shows that most often, firms do not change a single price at all. In fact, at the firm level, the frequency of full price-change inaction is 69%. In about 18% of the observations, firms adjust all product prices. However, about 13% of the sample represents instances where within one firm, at least one price change and price inaction occur simultaneously. Hence synchronization does happen very often, but in a sizeable number of cases it is incomplete.

Second, Table 2 shows the mean values of the share of other products by the same firm with positive, no and negative price changes, respectively, conditional on whether the price change of an individual product is negative, zero or positive. A strong degree of price coordination within the firms is seen, with the largest shares along the diagonal in the table.

Third, Figure 2 reports to what extent the predicted probability of a price decrease, no price change or price increase is affected by the share of price changes of other goods within the firm. The predicted probabilities stem from an ordered probit model where the dependent variable is whether the price change of a single product j is negative, zero or positive.[5] The upper part of Figure 2 shows that as the fraction of downward price adjustments within a firm, excluding good j, increases, the probability of observing a downward price adjustment for good j increases as well. The probability of inaction is decreasing, while the probability of an upward price adjustment decreases. Correspondingly, the lower part of Figure 2 illustrates that when the fraction of upward price adjustments within a firm, excluding good j, increases, the probability of observing an upward price adjustment for good j increases while the probability of inaction decreases. The probability of a price drop decreases. Hence both panels of Figure 2 show a strong degree of within-firm synchronization.

TABLE 1
FULL AND PARTIAL PRICE SYNCHRONIZATION

	%
No price change at all	69.6
Partial synchronization	12.1
All prices change	18.3

Notes
The reported numbers are based on the individual price observations aggregated up to the firm level.

TABLE 2
MEAN SHARE OF OTHER PRICE CHANGES BY THE SAME FIRM

	Individual price		
	Negative price change	Unchanged price	Positive price change
Other negative price changes	0.472	0.015	0.185
Other inaction prices	0.211	0.963	0.188
Other positive price changes	0.317	0.022	0.627

Notes
This table shows the mean values of the share of other products by the same firm with negative price changes, no price changes and positive price changes, respectively, conditional on whether the price change of an individual product is negative, zero or positive.

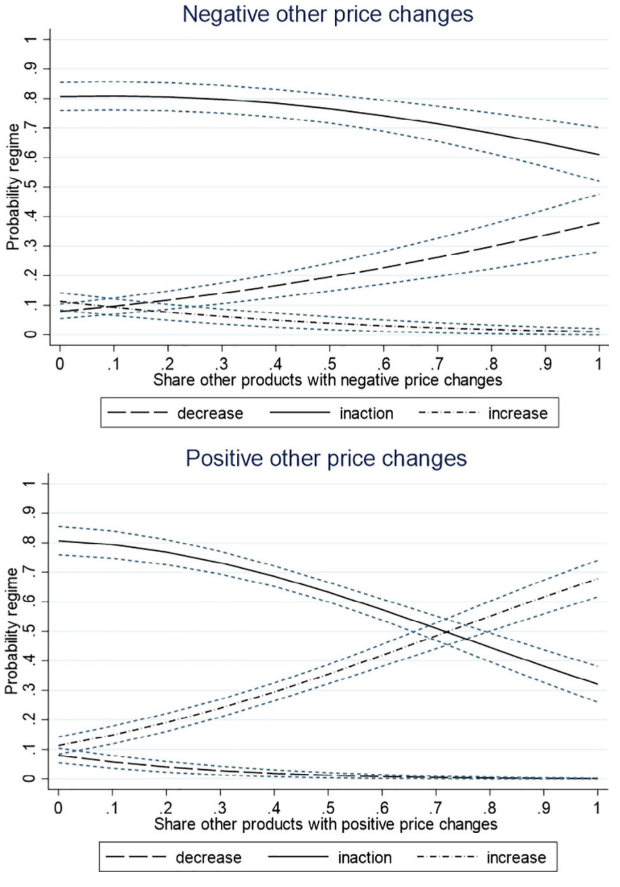

FIGURE 2. Probability of pricing regime conditional on share of other price changes in a firm.
Notes: The upper (lower) panel depicts how the probability of a price regime of a single product depends on the dynamics of the share of other products within the same firm experiencing a negative (positive) price change. The solid curves represent the predicted probability. The dashed curves define the confidence interval of the predicted probability. [Colour figure can be viewed at wileyonlinelibrary.com]

II. THE MODEL

Our model is meant to explain important aspects of micro-level pricing data that have been observed frequently: synchronization and infrequent price adjustment with both many small price changes and also large price adjustments.[6] The model nests two alternative explanations for these facts. First, the model is based on menu costs featuring economies of scope. Second, we consider a linear and convex cost component complemented with firm-specific shocks. More explicitly, the price adjustment costs function for prices is modelled as

$$C\left(\Delta P_{ijt}, P_{ijt-1}\right) = I\left(\Delta P_{ijt} \neq 0\right) \cdot \left(\frac{a}{m_{it}} + b \cdot \left|\Delta P_{ijt}\right| + \frac{c}{2} \cdot \left(\frac{\Delta P_{ijt}}{P_{ijt-1}}\right)^2 P_{ijt-1}\right), \tag{1}$$

where the index i refers to a firm, $\Delta P_{ijt} = P_{ijt} - P_{ijt-1}$ and P_{ijt} denotes the change and price, respectively, of product j in month t, m_{it} denotes the number of price changes by firm i in the same period, and $I(.)$ is an indicator function equal to 1 if the condition in brackets is satisfied, and 0 otherwise.

The first component of equation (1), a/m_{it}, captures traditional physical adjustment costs that are not related to the size of the price changes. These capture, for instance, producing new price lists, monthly supplemental price sheets, and informing and convincing interested parties (Levy *et al.* 1997). In our model, such a fixed cost of adjustment is given by a parameter a. At the same time, several studies suggest that firms obtain cost advantages when synchronizing price changes. Furthermore, simultaneous price changes and full or partial synchronization within the firms are observed in our data. To be able to replicate this pricing behaviour, we assume that menu costs allow a firm to obtain economies of scope and that the cost is deducted from the profit of the products subject to a price change. Hence, inspired by the existing models of menu costs featuring economies of scope, the fixed cost a is divided by the number of products with price changes, m_{it}. Thus the multi-product model goes beyond an approach based on simply the sum of N single-product firms. This also implies that the total fixed menu costs, a, do not depend on the number of price changes. One way of thinking about this is that each product manager who wants to change the product price for which he is responsible, may participate in gathering information. The efforts required for each product manager depend on how much these costs can be shared among all of the product managers who want to adjust a price. Hence the more product prices are involved, the less effort each product manager needs to put in, which is reflected by dividing the fixed cost a by m_{it}. In addition, customers need to be informed of price changes, and also the sales force needs to be knowledgeable. To some extent, such costs may be shared across various product accounts by a joint communication strategy.

We consider two additional adjustment costs types that depend on the price change size, also referred to as proportional adjustment costs. Linear costs are represented by $b \cdot |\Delta P_{ijt}|$, where $\Delta P_{ijt} = P_{ijt} - P_{ijt-1}$, and P_{ijt} denotes the price of product j in month t.[7] A convex cost component is given by the expression multiplied by the parameter c. The quadratic menu cost term

$$\left(\frac{\Delta P_{ijt}}{P_{ijt-1}}\right)^2 P_{ijt-1}$$

implies that larger price changes are very costly and provides incentives for the smaller price changes that we observe in the data descriptives.

We assume that each firm produces N_{it} goods, presuming monopolistic competition. The decision problem concerns product price changes maximizing the present value of discounted cash flows:

(2)

$$V(P_{it}, A_{it}, B_{it}) = \max_{\Delta P_{ijt+1}, j \in \{1, N_{it}\}} E_t\left(\sum_{d=0}^{\infty} (\beta)^d \left(\sum_{j \in \{1, N_{it}\}} \left(\pi\left(A_{ijt+d}, B_{ijt+d}, P_{ijt+d}\right) - C\left(\Delta P_{ijt+d+1}, P_{ijt+d}\right)\right)\right)\right).$$

The index i refers to a firm, the index j refers to a product, and the index t refers to a month. The symbols P_{it}, A_{it} and B_{it} *without* the subscript j denote vectors of the corresponding variables above for all the N products produced by the firm. For example, $P_{it} = \left(P_{i1t}, \ldots, P_{iN_it}\right)$. The expression $\pi\left(A_{ijt+d}, B_{ijt+d}, P_{ijt+d}\right)$ denotes the firm's revenue function net of wage costs for a product j at time $t + d$. The monthly discount rate is given by β. The expectations operator $E_t(.)$ is included due to the stochastic variables A_{ijt}

and B_{ijt} representing shocks to supply and demand of a product, respectively. In the model, $\Delta P_{ijt+1} = P_{ijt+1} - P_{ijt}$ is the decision variable. The realization of the shocks $A_{ijt+1}$and B_{ijt+1} in period $t+1$ comes after ΔP_{ijt+1} is determined. The menu cost function for prices, $C(.)$, is already presented in equation (1).[8]

We define q_{ijt} as a measure of how the expected value of the firm changes when the price of product j is increased by one unit. In fact,

$$(3) \quad q_{ijt} \equiv \frac{\partial V}{\partial P_{ijt}} = E_t \left(\sum_{d=0}^{\infty} (\beta)^d \left(\frac{\partial \pi\left(A_{ijt+d}, B_{ijt+d}, P_{ijt+d}\right)}{\partial P_{ijt+d}} - \beta \frac{\partial C\left(\Delta P_{ijt+d+1}, P_{ijt+d}\right)}{\partial P_{ijt+d}} \right) \right).$$

It represents the expected discounted value of marginal change in future profits minus the saved future menu costs. More details around this expression for q will be discussed later.

The first-order condition for price change equals

$$(4) \quad q_{ijt} - b \cdot I\left(\Delta P_{ijt} > 0\right) + b \cdot I\left(\Delta P_{ijt} < 0\right) - c\left(\frac{\Delta P_{ijt}}{P_{ijt-1}}\right) = 0.$$

A price will be changed if the benefits are larger than the costs associated with the adjustment:

$$(5) \quad q_{ijt} \Delta P_{ijt} > C\left(\Delta P_{ijt}, P_{ijt-1}\right).$$

Substituting $\Delta P_{ijt}/P_{ijt}$ from equation (4) into equation (5) informs us that prices behave according to the following rules:[9]

$$(6) \quad \frac{\Delta P_{ijt}}{P_{ijt-1}} = \frac{1}{c}\left(q_{ijt} - b\right) \text{ if } q_{ijt} \geq \sqrt{\frac{2 \cdot a \cdot c}{m_{it} \cdot P_{ijt-1}}} + b.$$

This expression tells that a price increase occurs if q_{ijt} is larger than the associated price change costs. Similarly, for a price reduction, we have

$$(7) \quad \frac{\Delta P_{ijt}}{P_{ijt-1}} = \frac{1}{c}\left(q_{ijt} + b\right) \text{ if } q_{ijt} \leq -\sqrt{\frac{2 \cdot a \cdot c}{m_{it} \cdot P_{ijt-1}}} - b.$$

The thresholds are trigger points defining a region of inaction. From equations (6) and (7), we also observe that small price changes are more likely with scope economies, since if the number of prices to be adjusted—that is, m_{it}—is large, then the threshold will be low. In that case, small shocks to q_{ijt} may induce small price changes.

For prices that are not adjusted, we have the condition

$$(8) \quad \frac{\Delta P_{ijt}}{P_{ijt-1}} = 0 \text{ if } -\sqrt{\frac{2 \cdot a \cdot c}{(m_{it}+1) \cdot P_{ijt-1}}} - b \leq q_{ijt} \leq \sqrt{\frac{2 \cdot a \cdot c}{(m_{it}+1) \cdot P_{ijt-1}}} + b.$$

Here, it is worth noting that a division by $(m_{it}+1)$ is present in the expression for the thresholds, compared to a division by m_{it} in equations (6) and (7).

Equations (6) and (7) also show that if $a = 0$ and $-b \leq q_{ijt} \leq b$, then the firm will not adjust its price. Hence linear costs induce infrequent price change. Strikingly, if $a = 0$, then we will see small price changes in the data. Minor deviations from the thresholds $q_{ijt} \geq b$ and $q_{ijt} \leq -b$ will induce small price changes. Finally, if the model also includes a firm-specific shock captured by q_{ijt}, then we see immediately that the model generates synchronization of price changes even if $a = 0$, and economies of scope are absent. This implies that a model where price adjustment costs are linear and convex, and where shocks are firm-specific, is capable of explaining three key features of micro-level data: infrequent price changes, small adjustments and synchronization. If fixed costs are present, that is, $a > 0$, then small price changes are infrequent, and the tails of the price change distribution will become thicker. Higher fixed costs cause lumpy price changes because the thresholds in equations (6) and (7) increase in absolute value. Then firms will not adjust prices for quite some time. Once adjustment takes place, the price change will be large.

To find a solution of our model requires that the price adjustment cost function is convex (i.e. $c > 0$) and twice differentiable in the price change everywhere except possibly at a price change equal to zero. The convexity assumption is clearly relevant for taking the first-order condition in equation (4). This condition holds in the case when the price change is non-zero, and is used to determine the size of the price change, that is, the intensive margin of the pricing decision. Given this solution for the size of the price change, the extensive margin can be found, indicating whether or not to change the price (see equation (5)). The extensive margin is based on a comparison of the value of adjustment versus the value of no adjustment, hence it is not important to have differentiability in the neighbourhood where the price change is equal to zero.[10] The limit case where $c = 0$ is a violation of the convexity and twice differentiable requirements. Thus other solution techniques (e.g. Dixit 1991; Alvarez and Lippi 2014) are necessary to understand the role of non-convex price adjustment costs in the absence of convex components (see also our discussion at the end of Section III). Note, however, that Nilsen and Vange (2019) observe, using a cut of the same data as in the present paper: ‘the probability of observing a price change, conditional that a price change also took place during the previous month, is 0.525’, while the unconditional probability of a price change is 0.200. The already presented substantial proportion of small price-change observations, that very large price changes are rare, together with the correlation over time, indicates the presence of convex costs. Thus the assumption $c > 0$ seems innocuous.

Coordination provides individual product managers the possibility to share the fixed menu costs. The fixed menu cost for a single price is a/m_{it}. This fixed menu cost is smallest if all prices of the firm are adjusted, that is, when $m_{it} = N_{it}$. Whether or not all prices are changed is determined by applying equations (6) and (7), where $m_{it} = N_{it}$. If these equations are satisfied, then all prices will be adjusted. If some prices are not meeting the requirement in equation (6) or (7) with $m_{it} = N_{it}$, then these prices not satisfying the condition will not be changed. They will remain unadjusted in this specific period, as the fixed menu cost per product price will only increase from now on, as it is divided by a smaller number of prices being changed, that is, $m_{it} < N_{it}$.

The next step in the optimization is to set m_{it} equal to the number of prices satisfying equations (6) and (7) in the previous optimization round. Now consider whether it is optimal to change the remaining product prices by checking whether the conditions in equations (6) and (7) are satisfied when applying the new number m_{it} in the thresholds. If some prices do not meet the requirements, then they are now also skipped from the set of price change candidates, and the optimization process will be repeated with a smaller number of candidate prices $m_{it} < N_{it}$. This process will continue until all prices in the set of candidates are meeting equation (6) or (7), and then they will be changed. Alternatively, it may be optimal to change no prices at all. Let us assume now that $0 < m_{it} < N_{it}$, and that m_{it} is the actual number of prices to be changed. We know from this that in the previous round of the optimization process, all prices that remain unchanged satisfy

$$-\sqrt{\frac{2\cdot a\cdot c}{(m_{it}+1)\cdot P_{ijt-1}}}-b\leq q_{ijt}\leq\sqrt{\frac{2\cdot a\cdot c}{(m_{it}+1)\cdot P_{ijt-1}}}+b,$$

as in equation (8). Note that the boundaries set on q_{ijt} in this expression are stricter when dividing by $(m_{it}+1)$ rather than by m_{it}. The set of product prices to be changed is given by

$$\left\{k\in\{1,\ldots,N_i\}\wedge\left(q_{ikt}\leq-\sqrt{\frac{2\cdot a\cdot c}{m_{it}\cdot P_{ikt-1}}}-b\vee q_{ikt}\geq\sqrt{\frac{2\cdot a\cdot c}{m_{it}\cdot P_{ikt-1}}}+b\right)\right\}.$$

The analysis of price decisions is summarized in Figure 3 for a firm producing two goods. On the horizontal and vertical axes, the marginal values of a price change for products 1 and 2 are provided, q_1 and q_2, respectively. For $j\in\{1,2\}$, the thresholds determining when q_1 and q_2 are small or large enough to induce price change are given by

$$S_j=\sqrt{\frac{2\cdot a\cdot c}{2\cdot P_{jt-1}}}+b\text{ and }T_j=\sqrt{\frac{2\cdot a\cdot c}{P_{jt-1}}}+b.$$

With $N_{it}=2$, m_{it} can take values $m_{it}\in\{0,1,2\}$. The inaction area I in the middle of Figure 3 is caused by the presence of the fixed menu cost parameter, a, and the linear menu cost component, b. To see this, our equations (6) and (7) state that complete inaction—that is, $m_{it}=0$—requires $j\in\{1,2\}$, $-T_j<q_j<T_j$. The same equations (6) and (7) state that the firm will adjust both prices—that is, $m_{it}=2$—if $q_j\geq S_j$ or $q_j\leq -S_j$ for $j\in\{1,2\}$. This happens in the areas denoted by II and III, bounded by what is referred to as S_j. It is scope advantages that create the difference between the thresholds S_j and T_j. More explicitly, for each product, the threshold that prevents price adjustment decreases when m_{it} goes from 1 to 2. Let us now look at the potential case where $q_2\geq T_2$ while $S_1<q_1<T_1$. In this case, it is clear that the price of product 2 will be changed independent of whether or not the price of

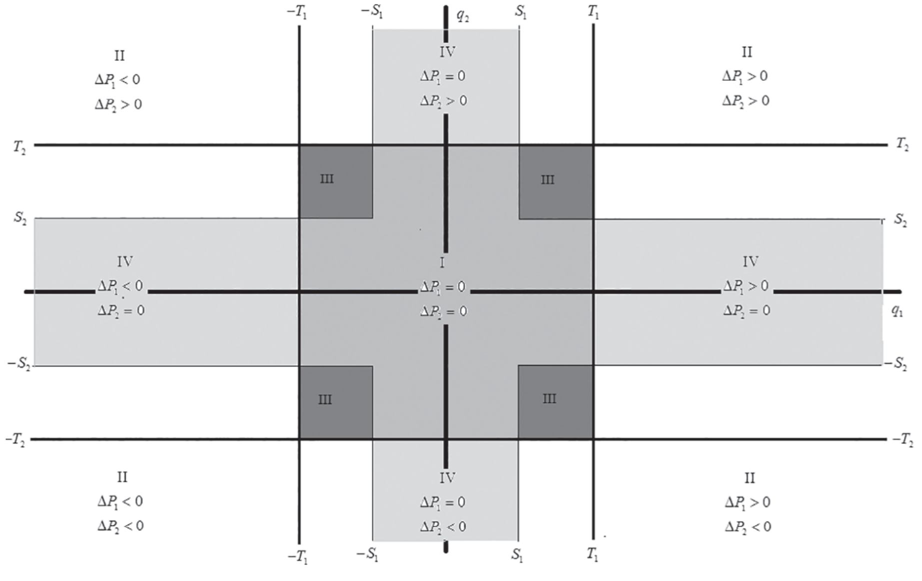

FIGURE 3. Pricing decisions by a two-product firm.

product 1 is changed. We see that if the price of product 1 will be changed too, $m_{it} = 2$, then the relevant threshold will be S_1. Thus we will end up with $m_{it} = 2$ instead of $m_{it} = 1$. Such a logic can also be applied in the case where $S_j < q_j < T_j$ for both products j. If the price adjustments can be coordinated, then the relevant thresholds are S_j, not T_j, which results in $m_{it} = 2$. Therefore in areas III, we see that in this case, scope advantages cause joint price adjustments ($m_{it} = 2$) instead of no price change at all ($m_{it} = 0$). Finally, we have the case where only one product will be changed, for instance, $q_1 > T_1$ and $-S_2 < q_2 < S_2$. These are the areas denoted IV in Figure 3.

We observe that fixed menu costs featuring economies of scope induce price stickiness and hence infrequent adjustments. Sheshinski and Weiss (1992) describe the implications of scope advantages in the adjustment cost technology for a multi-product firm as well. In their model, the firm makes two products, similar to the example depicted in Figure 3. Sheshinski and Weiss (1992) also find that firms have an incentive to synchronize price changes and exhibit infrequent adjustments given that the two prices are strategic complements. Note, however, with our model, where the menu costs are shared among the products whose prices are changed, small price changes are more likely, as economies of scope tend to lower the thresholds determining when price change occurs (see equations (6) and (7)). All in all, economies of scope prompt synchronization of price change, illustrated by the presence of areas III in Figure 3.

III. Estimation

As seen from equation (3), the expression for q_{ijt} is composed of discounted expected values of the marginal profit and the marginal menu cost function, respectively. The first element of q_{ijt}, $\partial\pi(.)/\partial P_{ijt+s}$, reveals that a price change influences marginal profits in future periods. For notational convenience we temporarily abstract from sub-indices for the firm, product and time. Essentially, we follow Alvarez and Lippi (2014), assuming that a monopolist sells N products with additively separable demands.[11] A product is assumed to be produced according to a Cobb–Douglas production technology with a flexible and homogeneous labour input component, L, where w denotes the exogenous wage for a worker. The production is determined by $Q^S(L) = A \cdot L^{\alpha}$, where A captures supply shocks, and $0 < \alpha < 1$. Firms' market power is modelled by assuming an isoelastic demand function given by $Q^D(P) = B \cdot \left(P/P^C\right)^{-\varepsilon}$, where B captures demand shocks, and $\varepsilon > 1$. The price of a firm's product is given by P, and P^C denotes the general price level in the industry. Both P^C and wages w are exogenous to the firm, reflecting that we employ a partial equilibrium model. Abstracting from inventory, the profit is determined by $\pi(.) = P \cdot B \cdot \left(P/P^C\right)^{-\varepsilon} - w \cdot (B/A)^{1/\alpha} \cdot \left(P/P^C\right)^{-\varepsilon/\alpha}$. The first-order derivative of profit $\pi(.)$ with respect to price P, $\partial\pi(.)/\partial P_{ijt+s}$, is a non-linear function of A, B, w and P^C. It is worth noting that with our assumptions concerning the profit structure of the firm, sales volume does not feature explicitly in the marginal profit of the firm. Instead, demand conditions are represented by B and P^C.

The second term of equation (3), involving $C\left(\Delta P_{ijt+s}, P_{ijt+s-1}\right)/\partial P_{ijt+s}$, depicts that a change in price saves menu costs in future periods. One may abstract the convex component from the q expression given that the price changes are rather small.[12] The derivative of the quadratic adjustment cost expression, $\left(\Delta P_{ijt}/P_{ijt-1}\right)^2$, will be negligible in our proxy for q as given by equation (3). This assumption is supported by the descriptives already shown in Figure 1.[13] With the simplifications discussed above, we assume that q is given by

$$q_{ijt} = \gamma_0 + \gamma_1' X_{ijt} + \kappa_{it} - \eta_{ijt}. \tag{9}$$

The zero mean stochastic terms η_{ijt} are assumed to be normally distributed with variance σ_η^2. To proxy the marginal profit of the firm, the vector X_{ijt} contains information reflecting both supply and demand shifters A and B, approximated by a set of year dummies, and a set of 11 monthly dummies to control for seasonal effects. Furthermore, the vector includes two commodity group-specific dummies and a monthly commodity group-specific price index P^C for the relevant product. This index may pick up changes in demand conditions due to competition (i.e. other competitors' actions), but might also say something about the relevant cost level in the industry not accounted for in the simple model to derive marginal profit. However, note that within-industry synchronization of prices has been found to be negligible in practice (Nilsen *et al.* 2021). We follow Rotemberg and Woodford (1999) and Carlsson (2017), and consider that wages are a fundamental cost component driving prices. In the manufacturing industry of Norway, the labour cost as a share of the operating profit is in the range 0.75–0.90 in the years of our study (NOU 2019, p. 6). We incorporate the (log-transformed) wage rate w and its square to capture some of the non-linearity of marginal profit / marginal costs discussed above.[14]This latter variable w is measured annually at the firm level, not the product level.[15] Hence the vector contains wage information for the previous year. This is consistent with an assumption that the firms use an AR(1) process to predict the wage rate. Using information for the previous year also reduces potential endogeneity problems. The monthly dummies may pick up systematic deviation between the annual and monthly variables. They will also control for general inflationary developments in the macro-economy.

One explanation of price synchronization could be that a firm is subject to a demand or supply shock that is common to all of its products driving all prices in the same direction simultaneously. To control for this, we implement a latent class model allowing for a shock that is firm- and time-specific.[16] The latent class approach is implemented by adding a shock κ_{it} to equation (9), where the process generating κ_{it} is characterized by two parameters to be estimated: ψ and κ. With probability ψ, the shock is $\kappa_{it} = \kappa$, and with probability $1-\psi,\ \kappa_{it} = 0$. All products within the firm are subject to this shock, which will be picked up by the latent class parameters. That means that if the observed coordination is due to only these common shocks—and we have controlled for these—then we would expect the fixed menu cost generating coordination to be insignificant as well.

Our estimation strategy is based on a two-step Heckman type selection estimator. First, an ordered response model is developed to estimate the probability of price increases, maintaining the current price, and price reductions—the extensive margins of price changes. The main objective of the first step is to get an estimator for the determinants of q_{ijt}. Second, we estimate the equations determining the level of the price adjustment, using selection correction terms based on the estimates obtained from the ordered response model.[17] Note that the number of products that the firm changes, m_{it}, is endogeneous. We discuss the sensitivity of our results to endogeneity issues in Section IV.

Extensive margin

Using equations (3), (4), (5) and (9), we show in the Online Appendix that the log-likelihood function is

$$(10)\qquad \ln L = \sum_{t=1}^{T} \sum_{\Delta P_{ijt}>0} \ln E\left(\Phi\left[\widetilde{\gamma}_1' X_{ijt} + \kappa_{it} + \left(\widetilde{\gamma}_0 - \widetilde{b}\right) - \sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{m_{it}\cdot P_{ijt-1}}}\right]\right)$$

$$+\sum_{t=1}^{T}\sum_{\Delta P_{ijt}<0}\ln E\left(\Phi\left[\widetilde{\gamma}_1' X_{ijt}+\kappa_{it}+\left(\widetilde{\gamma}_0+\widetilde{b}\right)-\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{m_{it}\cdot P_{ijt-1}}}\right]\right)$$

$$+\sum_{t=1}^{T}\sum_{\Delta P_{ijt}=0}\ln E\left(\begin{array}{l}\Phi\left[\widetilde{\gamma}_1' X_{ijt}+\kappa_{it}+\left(\widetilde{\gamma}_0+\widetilde{b}\right)-\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{(m_{it}+1)\cdot P_{ijt-1}}}\right]\\-\Phi\left[\widetilde{\gamma}_1' X_{ijt}+\kappa_{it}+\left(\widetilde{\gamma}_0-\widetilde{b}\right)-\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{(m_{it}+1)\cdot P_{ijt-1}}}\right]\end{array}\right)$$

where the operator $E(.)$ takes expectations with respect to the shock κ_{it}, and $\Phi(\cdot)$ denotes a standard normal cumulative distribution function. A large number of the structural parameters in the model can be estimated. Nevertheless, the variance of the error term remains unknown, as is common in probit type models. As a consequence, the variance σ_η^2 of the error term in equation (9), must be set equal to 1.[18] Hence all structural parameter estimates have to be understood as relative to the standard deviation σ_η. This is not very harmful in terms of interpretation. For instance, if our estimate for the convex cost of price changes is $\widetilde{c}=c/\sigma_\eta$, then according to equations (6) and (7), its inverse measures how much of a one standard deviation shock is transmitted into a price change. Likewise, the scaled parameters $\widetilde{a}=a/\sigma_\eta$ and $\widetilde{b}=b/\sigma_\eta$ measure how important the original parameters are in determining the decision of whether or not to change price relative to a one standard deviation shock. From now on a $\sim$ on top of a parameter indicates that the original parameter is divided by the standard deviation σ_η. Maximizing the log-likelihood in equation (10) allows us to acquire estimates of $\widetilde{\gamma}_0$, $\widetilde{\gamma}_1$, $\widetilde{b}$, $\widetilde{a}\cdot\widetilde{c}$, $\widetilde{\kappa}$ and ψ. To construct a proxy for q, the estimates for $\widetilde{\gamma}_0$ and $\widetilde{\gamma}_1$ can be used.

Obviously, firms in our sample exhibit a large degree of variation. Altogether, our estimation strategy accounts for heterogeneity among firms in our sample in a number of ways. First, our q proxy takes care of differences between firms by controlling for various commodity groups and by a commodity group-specific price index P^C for the relevant product. Next, the thresholds determining whether or not prices are changed vary with the number of products made by the firm and by the prices set. Finally, we allow for firm-specific shocks.

Intensive margin

Once the estimates are obtained by maximizing the log = likelihood function, equations (6) and (7) can be used to determine a model for the size of the price change, driven by $\widehat{q}_{ijt}$. The hats above some parameters denote that estimated values based on the first-stage extensive margin have been used. This model needs to account for selection. We estimate

(11) $$\frac{\Delta P_{ijt}}{P_{ijt-1}}=\frac{\widehat{\widetilde{\gamma}}_0-\widehat{\widetilde{b}}}{\widetilde{c}}+\frac{\widehat{\widetilde{\gamma}}_1 X_{ijt}+\widehat{\psi}\cdot\widehat{\widetilde{\kappa}}+\widehat{\lambda}_{ijt}^{+}}{\widetilde{c}}+\vartheta_{ijt}^{+}$$

and

(12) $$\frac{\Delta P_{ijt}}{P_{ijt-1}}=\frac{\widehat{\widetilde{\gamma}}_0+\widehat{\widetilde{b}}}{\widetilde{c}}+\frac{\widehat{\widetilde{\gamma}}_1 X_{ijt}+\widehat{\psi}\cdot\widehat{\widetilde{\kappa}}-\widehat{\lambda}_{ijt}^{-}}{\widetilde{c}}+\vartheta_{ijt}^{-}$$

for price increases, and price reductions, respectively. Thus our estimation technique relies on both the extensive and intensive margin empirical moments. Equations (11) and (12) allow us to identify the parameter $\widetilde{c}$ representing the quadratic adjustment cost component. With this estimate, and those obtained in the first step, it is then also possible to obtain the parameters of the fixed cost term, $\widetilde{a}$. The terms ϑ_{ijt}^{+} and ϑ_{ijt}^{-} denote zero mean errors, while the expressions λ_{ijt}^{+} and λ_{ijt}^{-} are inverse Mills ratios. For more detail on the estimation strategy, see the Online Appendix.

As shown in this section, the interdependency between the price changes in our model—economics of scope—is easily incorporated in the q framework. One may also employ simulated method of moments (SMM) to estimate the structural model outlined above. However, as prices cannot be regarded as independent, in an SMM routine this expands the state space considerably. Firms in our sample on average report about 5 product prices (some firms report as many as 20 different prices). Assuming, for each of these 5 product prices, that 100 points are used in a grid, one would already have a state space with at least $100^5 = 10^{10}$ points, as in this calculation stochastic processes expanding the dimensionality of the state space have not been accounted for yet.[19] In spite of necessary simplifying assumptions used when approximating the marginal value of a unitary price change (i.e. q), we prefer the ML routine to the SMM due to computational feasibility. Another advantage of our approach is that we do not have to assume that firms are similar in the number of products that they manufacture. Very often, SMM approaches can account for only a limited number of types of firms. In fact, the ML approach that we use can fully exploit the richness of our data and incorporate each firm's number of products that we observe.

IV. Results

The estimation results are reported in Table 3. In columns (1) and (2), we allow all the three adjustment costs components to take values different from zero; in columns (3) and (4), we abstract from the latent class approach. Next, we reintroduce the latent class approach, but in columns (5) and (6), we set $\widetilde{a} = 0$, and in columns (7) and (8), $\widetilde{b} = 0$. The first observation that we make, before getting into the details, is that there is a concave relationship between q and the wage rate. A second result worth noticing is that the bootstrapped 95% confidence intervals for all the estimated adjustment costs parameters ($\widetilde{a}$, $\widetilde{b}$ and $\widetilde{c}$) show that these parameter estimates all are significantly different from zero. In particular, that $\widetilde{c} > 0$, together with the strong significance of the $\widetilde{a} \cdot \widetilde{c}$ term in the first stage of the estimation procedure, is comforting, as our econometric model depends strongly on the assumption that $c > 0$. We also find evidence supporting the use of the latent class model. A second class exists with a probability of about 4.5%. The size of the shock $\widetilde{\kappa}$ is positive and attains a value of about 2.6. Given that it is scaled by the standard deviation of a normal distribution—that is, $\widetilde{\kappa} = \kappa/\sigma_\eta$—its size is quite large. This implies that once such a shock hits a firm, which is the case about once every 2 years—that is, $1/\psi = 1/0.047 = 21$ months—price changes tend to be synchronized within the firm. Due to the large size of the shock, at least partly, synchronization is explained by firm-specific shocks. However, it is not the sole explanation of the probability of a price change. In fact, Table 1 shows that full synchronization happens relatively often with a frequency of about 18%. Hence the shock process that we identify does not explain entirely the synchronization observed in the data. We conclude from this that idiosyncratic shocks likely play an important role and may lead to synchronized price change as well. We have made an attempt to estimate a model including an additional latent class. However, in that case the estimation routine indicated a flat likelihood surface. In the context of a latent class model, this is associated with over-parametrization of the model, that is, too many latent classes (Cameron and Trivedi 2005). We interpret this as two classes already capturing the existing shock process quite well.

Starting with columns (1) and (2) of Table 3, we observe the existence of significant linear menu costs, $\widetilde{b}$. Estimating equations (11) and (12) by OLS reveals that the convex cost

TABLE 3
ESTIMATION RESULTS

	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)	Coeff. (5)	SE (6)	Coeff. (7)	SE (8)
Maximum likelihood results								
$\ln w_{t-1}$	1.662	0.326	1.080	0.300	2.0045	0.311	1.033	0.327
$(\ln w_{t-1})^2$	−0.547	0.113	−0.389	0.104	−0.678	0.108	−0.291	0.113
$\widetilde{a} \cdot \widetilde{c}$	17.134	0.944	14.632	0.769	—	—	316.031	3.369
$\widetilde{b}$	1.003	0.010	0.924	0.009	1.303	0.006	—	—
$\widetilde{\kappa}$	2.568	0.057	—	—	2.578	0.054	3.094	0.057
ψ	0.047	0.003	—	—	0.044	0.002	0.051	0.003
Log-likelihood	−25,217.9		−26,373.1		−26,042.1		−32,388.6	
Observations	39,082		39,082		39,082		39,082	
OLS with selection correction								
$1/\widetilde{c}$	0.050	0.001	0.052	0.001	0.051	0.001	0.020	0.003
Bootstrap confidence intervals								
$\widetilde{a}$	0.856	[0.455, 1.771]	0.754	[0.336, 1.398]	—	—	6.326	[4.730; 9.463]
$\widetilde{b}$	1.003	[0.836, 1.239]	0.924	[0.723, 1.086]	1.303	[1.237, 1.503]	—	—
$\widetilde{c}$	20.016	[16.429, 25.837]	19.393	[15.533, 23.875]	19.449	[15.099, 20.646]	49.956	[45.110, 60.312]

Notes
Commodity-specific price indices, commodity-type dummies, year dummies and monthly dummies are included in the first-stage equations. All the parameters except for ψ should be thought of as normalized by the standard deviation σ_η. In square brackets, 95% confidence intervals are provided, obtained by bootstrapping. This works as follows. From the dataset that we use to estimate the model, we draw N observations with replacement, where N is the number of firms analysed for the initial estimations. This means that we cluster around the producers. The ordered probit model is estimated first to obtain estimates $\widetilde{\gamma}_0$, $\widetilde{\gamma}_1$, $\widetilde{b}$, $\widetilde{a} \cdot \widetilde{c}$, $\widetilde{\kappa}$ and ψ, for each new bootstrap sample. Next, we estimate the price change equations. This step is replicated 200 times. After 200 replications, we have obtained a distribution for each parameter of interest. The 95% confidence interval for these parameters is based on the limits of the 2.5% and 97.5% quantiles.

parameter $\widetilde{c}$ is significantly different from zero. Bootstrapping yields that $\widetilde{a}$ is different from zero as well, according to common statistical conventions. These findings are in line with our descriptive statistics. They revealed a large amount of zeros. Inactivity can be explained by both linear and fixed menu costs. As we control for common shocks to products within the firm, coordination of prices is also explained by economies of scope in menu costs.

In columns (3) and (4) of Table 3, we set the parameters related to the latent class approach, $\widetilde{\kappa}$ and ψ, equal to zero. Due to this, the performance of the model, measured by the log-likelihood, is reduced from -25217.9 to -26373.1. However, controlling for common shocks does not affect the main conclusions obtained from the model, and the coordination results are therefore likely to stem from the shape of the menu cost function, not the common shocks. When we turn to columns (5) and (6), we reintroduce the latent class approach but set $\widetilde{a}=0$.[20] Now the $\widetilde{b}$ parameter is approximately 30% larger relative than the one in columns (1) and (2). The reason is that there is no help from the square root in the threshold $|\sqrt{(2\cdot\widetilde{a}\cdot\widetilde{c})/(m\cdot P)}+\widetilde{b}|$ given that $\widetilde{a}=0$. Thus, to ensure enough inaction, the $\widetilde{b}$ parameter has to increase. In columns (7) and (8), we set $\widetilde{b}=0$. Looking at the threshold for (in)action, which is $|\sqrt{(2\cdot\widetilde{a}\cdot\widetilde{c})/(m\cdot P)}+\widetilde{b}|$, it is clear that when $\widetilde{b}=0$, the product $\widetilde{a}\cdot\widetilde{c}$ has to be larger to induce inaction. Both parameters $\widetilde{a}$ and $\widetilde{c}$ increase in columns (7) and (8). An indicator hinting at misspecification is the log-likelihood of the first-stage estimations. We find these to be -25217.9, -26373.1, -26042.1 and -32388.6, respectively). Thus the full model reported in columns (1) and (2) outperforms all other models statistically when using conventional likelihood ratio tests, and is therefore our preferred specification.

We have also made an attempt to estimate a model without the assumption of economies of scope by assuming that the fixed menu cost is given by $\widetilde{a}$ rather than by a/m_{it}.[21] For such a model, the maximum likelihood routine is driving the $\widetilde{a}\cdot\widetilde{b}$ term in $\sqrt{(\widetilde{a}\cdot\widetilde{c})/P_{ijt-1}}$ towards zero. Thus a model without scope economies becomes equivalent observationally to the one presented in columns (5) and (6) of Table 3, where $\widetilde{a}=0$. We observe that this specification is outperformed in terms of the value of the log-likelihood function by the full model in columns (1) and (2). This is clear evidence for the importance of menu costs subject to scope economies.

Economic importance

To obtain some insight into the economic importance of the various menu cost components, we conduct some exercises based on the results presented in columns (1) and (2) of Table 3. Abstracting from fixed costs (i.e. $\widetilde{a}$), we see that convex costs are larger than linear costs when $\Delta P/P$ exceeds 0.100 ($=2*1.003/20.016$).[22] This happens in about 2% of the observations. Focusing on non-convex costs, we find that the linear costs are largest when $\Delta P/P\geq\widetilde{a}/(\widetilde{b}\cdot m\cdot P)$.[23] Setting $m=1.06$, the average number of simultaneous product price changes, and $p=1531$, the average price, and using the parameter estimates for $\widetilde{a}$ and $\widetilde{b}$ reported in columns (1) and (2) of Table 3, i.e. $\widetilde{a}=0.856$ and $\widetilde{b}=1.003$, we find that linear costs are largest when $\Delta P/P\geq 0.856/(1.003\cdot 1.06\cdot 1531)\approx 0$. This means that at the intensive margin, linear costs are relatively important.

In Table 4, we use the predicted q values from the full model as presented in columns (1) and (2) of Table 3, but calculate the alternative price adjustment probabilities after setting either $\widetilde{a}$ or $\widetilde{b}$ equal to zero in the thresholds.[24] We think of these as counterfactual analyses. The actual frequencies and kurtosis from the actual data are presented in column (1), while the corresponding results based on the extensive margin of the full menu costs model (columns (1) and (2) of Table 3), are reported in column (2). The full model generates probabilities that come very close to the observed frequencies in the data. The kurtosis for the standardized

TABLE 4
COUNTERFACTUAL ANALYSES BASED ON FULL-MODEL SPECIFICATION

%	Data (1)	Full model (2)	Full model and $\widetilde{a}=0$ (3)	Full model and $\widetilde{b}=0$ (4)
Price increase	14.8	14.7	22.6	45.1
Inaction	76.5	75.5	63.7	23.7
Price decrease	8.7	9.4	13.7	30.3
Kurtosis	6.0	4.5	2.0	3.3

Notes
The results are based on the 'Full model' results in columns (1) and (2) of Table 3. Contrary to Table 1 where the observation unit is firm, the observation unit here is product.

price changes (see Alvarez *et al.* 2016) of the actual data in column (1) is 6.0. The predicted kurtosis drops to 4.5 in column (2).[25] Turning to column (3), where the fixed cost parameter is $\widetilde{a}=0$ and the other parameters of the full model remain at the values in columns (1) and (2) of Table 3, the average probability of inaction decreases by more than 10 percentage points, and the average action probabilities increase correspondingly. We also observe a 55% decrease of the kurtosis. In column (4), we see that setting the linear menu costs at $\widetilde{b}=0$, while all other parameter values of the full model are used from columns (1) and (2) of Table 3, strongly deteriorates the match between the probabilities and the figures presented in columns (1) and (2) of Table 4. Thus, in particular, linear menu costs in the full model are important to understand key data moments.

One explanation is that in our data, we observe quite some variation in the price level (see also Table A1 in the Appendix). Equations (6), (7) and (8) show that the price level enters the thresholds, which determine when a price change occurs or not. Firms for which the price level is high will have a low threshold. In those cases, the specification in columns (1) and (2) of Table 3 is likely to over-predict the occurrence of a price change in the case when the linear adjustment cost component is equal to zero.

Robustness

We have investigated whether endogeneity of m_{it} is driving our main conclusions by employing a two-step control function approach inspired by Rivers and Vuong (1988), which is discussed further by, for instance, Wooldridge (2014, 2015). We first estimate a model for the fraction of prices changed at the firm, that is, m_{it}/N_{it}. Here we have used an interval regression model that is a generalization of censored regression, since the degree of coordination is such that $0 \leq m_{it}/N_{it} \leq 1$. An exclusion restriction in our control function approach is not essential due to the highly non-linear nature of the first-step model (Altonji *et al.* 2005; Card and Giuliano 2013). The first stage delivers an estimation error measuring the difference between the realization of m_{it}/N_{it} and its predicted value: $\widehat{v}_{it}=m_{it}/N_{it}-\widehat{m}_{it}/N_{it}$. Like the latent term κ_{it}, this estimation error $\widehat{v}_{it}$ denotes a firm-level shock. The estimation error is to be included in equation (10).[26] The bottom line of this additional exercise controlling for endogeneity of m is that endogeneity is not driving our results.[27]

The estimation results are robust to initiating the estimation algorithm from different sets of starting values. Thus the parameter estimates reported in Table 3 seem to correspond to a global maximum. We have also performed two additional analyses to see whether our results are driven by unobserved heterogeneity (not reported, but available from the authors

TABLE 5
FULL AND PARTIAL PRICE SYNCHRONIZATION—SUB-INDUSTRIES

	Intermediate (1)	Durables (2)	Non-durables (3)
No price change at all	70.5	79.8	64.1
Partial synchronization	11.7	9.6	17.8
All prices change	17.9	10.6	18.1

Notes
The reported numbers are based on the individual price observations aggregated up to the firm level.

on request). First, in one version of the latent class model we have replaced the shock κ_{it} by a term κ_i, which is hence only firm-specific but time invariant. Hence equation (9) becomes $q_{ijt} = \gamma_0 + \gamma_1' X_{ijt} + \kappa_i - \eta_{ijt}$. Second, we have also estimated the model for two different groups of firms in terms of the number of products they make, that is, $N_{it} \leq 4$ and $N_{it} \geq 5$. The estimates for these two approaches to control for unobserved heterogeneity do not alter our conclusions.

There is ample evidence in the literature that there is huge heterogeneity in pricing behaviour across sectors, firms and products (see, for instance, Klenow and Malin 2010). We have therefore estimated the 'full model' as reported in columns (1) and (2) of Table 3, but now split the sample according to (sub-)industries Intermediate goods, Durables and Non-durables. These additional results are reported in Tables 5 and 6.[28] When it comes to full and partial synchronization, we see in Table 5 that the Intermediate goods sector is very similar to the already reported results for the complete sample, which is obvious given that most observations are from this industry (23,716 out of 39,082). The most prominent deviation from the full sample characteristics is perhaps found for Durables. Here we observe a larger frequency of 'No price change at all', and corresponding lower frequency of 'All prices change'. This might indicate that the fixed price change component is more important in this industry.

Turning to Table 6, we find the same pattern. Durables seems to deviate most from the two other industries. In column (2) of Table 5, we see that the fixed component is much larger for this industry, which probably reflects the specific nature of the products in that industry. Such a result is very interesting, as Barsky *et al.* (2007) argue that sticky prices in the durable sector can make the entire economy behave like a sticky price economy, even if the non-durable sector has flexible prices. For Non-durables we see that the convex price adjustment parameter is larger. Note that altogether, we find statistically significant evidence for all types of adjustment cost components in the three industries investigated, apart from the linear cost parameter estimated for Durables.

V. SIMULATION

To illustrate the importance of coordination in price changes further, two versions of a partial equilibrium model are simulated: a 'scope' economy version where the fixed part of a firm's price change cost a is divided among only the products for which the prices are changed, m, and a 'no scope' version where coordination of price changes yields no cost advantages and prices are set independently as scope economies are absent. The two versions differ only by the price change technology, that is, the fixed costs represented as a/m or a/N, where m is the number of product prices that are changed, and N represents the number of products made by the firm.

TABLE 6
ESTIMATION RESULTS—SUB-INDUSTRIES

	Intermediate		Durables		Non-durables	
	Coeff.	SE	Coeff.	SE	Coeff.	SE
Maximum likelihood results						
$\ln w_{t-1}$	0.786	0.447	1.894	1.230	3.178	0.593
$(\ln w_{t-1})^2$	−0.236	0.153	−0.817	0.451	−1.051	0.207
$\widetilde{a} \cdot \widetilde{c}$	18.969	1.329	646.991	88.721	26.082	1.934
$\widetilde{b}$	0.999	0.012	0.840	0.056	0.718	0.019
$\widetilde{\kappa}$	2.695	0.093	3.133	0.141	2.290	0.079
ψ	0.044	0.003	0.068	0.008	0.063	0.007
Log-likelihood	−15,416.3		−1818.79		−7064.9	
Observations	23,716		5018		10,348	
OLS with selection correction						
$1/\widetilde{c}$	0.054	0.001	0.039	0.003	0.033	0.001
Bootstrap confidence intervals						
$\widetilde{a}$	1.026	[0.324, 2.644]	24.913	[1.053, 111.375]	0.849	[0.270, 2.172]
$\widetilde{b}$	0.999	[0.804, 1.319]	0.840	[−0.290, 1.921]	0.718	[0.225, 1.302]
$\widetilde{c}$	18.482	[13.979, 24.198]	25.970	[11.129, 48.727]	30.710	[20.228, 48.759]

Notes
See Table 3.

We approximate q with an AR(1) process, which is in contrast to the model used in our ML analysis, where q is static. The idiosyncratic shocks of the q-process are drawn from a normal distribution $N\left(0, \sigma_v^2\right)$ and are autocorrelated, with correlation coefficient ρ. With the simulated q-series at hand, we build up the series of price changes and prices using equations (6), (7) and (8) with $a = 0.856$, $b = 1.003$ and $c = 20.016$. These values correspond to the parameter estimates for the full model in columns (1) and (2) of Table 3.[29]

Table 7 presents the results of the simulations with scope economies in column (2). The first issue to notice is that the actual moments, in column (1), and the simulated ones for the 'Scope' version in column (2) are very similar. Thus our model describes the behaviour in the actual data quite well.

We simulate the 'No scope' version of the model using the identical q-values as for the first version. In columns (2) and (3) of Table 7, the moments of 'Scope' show a better than or equal fit than those of 'No scope'. In particular, the model with 'Scope' economies produces a lower degree of 'Partial synchronization' at the firm level than those of 'No scope' advantages. This means that the incentive to change a single price increases moving from the 'Scope' to 'No scope' columns. In the 'Scope' model, a single price will be adjusted while other prices remain the same in case the marginal benefits of doing so for this individual price are high, while those for the other prices are low. A single price will be adjusted only, in which case $m = 1$, if q exceeds the threshold with $m = 1$. In this case, the fixed cost is equal to a. This makes it less likely to observe single price changes in the model with scope advantages compared to the 'No scope' model, where fixed costs are equal to a/N, unless a large shock hits a product, but not the entire firm.

An important implication of this finding is that firms are less inclined to self-select into price changes in the 'Scope' case. Self-selection, as discussed by Golosov and Lucas (2007), Midrigan (2011) and Carlsson (2017), occurs when price changes are driven by a cost–benefit analysis. Non-neutrality of monetary policy is determined by the extent to which such selection effects exist. Our model with scope economies produces a lower degree of state dependency. This can be seen as follows. Recall that by going from the 'Scope' results to the 'No scope' results in Table 7, we observe a higher frequency of price changes. This is caused by the probability of an individual price change also depending on the need to change other prices set by the firm in the economies of scope version of the model. This is due to fixed adjustment costs being given by a/m, where m is the number of product prices that are changed. In the 'No scope' model, the fixed adjustment costs are always given by a/N

TABLE 7
SIMULATED MODEL RESULTS

%	Data (1)	Full model: Scope (2)	Full model: No scope (3)
No price change at all	69.1	69.6	66.5
Partial synchronization	12.5	12.6	15.8
All prices change	18.4	17.8	17.8
Price increase	14.8	14.7	15.5
Inaction	76.5	75.7	74.2
Price decrease	8.7	9.6	10.3
Annual inflation	3.9	5.0	5.0

Notes
Contrary to Table 1 where the observation unit is firm, the observation unit here is product.

and hence are independent of the number of product changes. Due to scope economies, the selection effect is lower as firms are less likely to respond to a money shock as the adjustment threshold depends on other prices to be changed in the firm as well. Especially in the case when product prices in a firm are subject to idiosyncratic shocks, selection effects tend to be low. Heterogeneity in price shocks may induce fundamentals of certain product prices within the firm to be far away from the thresholds to trigger price change, while for other products, price change is desirable. In such circumstances, the thresholds are high for prices that are featuring high benefits of price change, which signals that then an individual price change may be too costly altogether. The opposite holds in the case when firms' product prices are subject to common firm-level shocks. When such shocks arrive, many product prices will benefit from price adjustment. Then the thresholds will be lower due to scope advantages, signifying that it is possibly not too costly to adjust individual prices. Obviously, the ratio of the variance of common shocks to the variance of idiosyncratic shocks will determine the importance of selection effects. As a consequence, in Table 7 our 'Scope' model featuring lower selection effects should produce a higher response of real output to a monetary shock than a state-dependent model with a pure fixed menu cost where selection effects are stronger. This real response of the economy to a money shock dampens when the adjustment becomes more state-dependent, relative to a standard Calvo model (Calvo 1983).

VI. Concluding Remarks

Knowledge concerning pricing behaviour is crucial for constructing models facilitating the analysis of macro-economics dynamics in general, and especially how monetary shocks affect the real economy. Most of the existing studies in the literature on pricing behaviour focus on price changes in single-product firms, but the number of studies focusing on multi-product firms is increasing. Furthermore, an analysis of the manufacturing industry is of great importance as most macro-economic models assume that the price rigidities stem from this sector. Hence our analysis contributes to advancing knowledge concerning non-neutrality of monetary shocks. Furthermore, as our data are based on what the Norwegian authorities use to track and scrutinize the producer price index, they are of high quality.

Both infrequent and many small price changes together with a high degree of within-firm synchronization are observed in the data. We show that such patterns are driven by a combination of both product-specific and firm-specific shocks, and by price adjustment costs of multiple functional forms: convex, linear and fixed menu costs subject to scope economies. A simulation of our model shows that scope economies in the pricing technology decrease the responsiveness of a single-product price to a product-specific shock, because this single-product pricing decision is now also dependent on the benefits associated with changing prices of other products made by the firm. The economies of scope in the pricing technology reduce the extent to which pricing behaviour is state-dependent. As a result, it also influences non-neutrality of monetary policy. Altogether, the results reveal the potential benefits of deviating from traditional menu cost models in which only fixed or convex costs are included.

Acknowledgments

We thank two reviewers for insightful and constructive comments, and Fred Schroyen for constructive discussions. We would also like to thank participants for valuable comments during several presentations at Mannheim University, the Cologne Workshop on Macroeconomics, the Norwegian School of Economics, Maastricht University, the Annual Norwegian Business Economics and Finance Conference (FIBE), the Meeting of the Norwegian Association for Economists, the Royal Economic Society

Conference, Statistics Norway, the Nordic Econometric Meeting, CESifo Area Conference on Macro, Money & International Finance, Catholic University of Eichstätt-Ingolstadt, Oslo Business School, the EARIE conference and the International Association for Applied Econometrics conference. The usual disclaimer applies.

ENDNOTES

1. For a somewhat earlier study on this, see, for instance, Carlson (1992).
2. This means that the number of a firm's product prices and the number of price changes observed provide lower bounds on the actual numbers.
3. In the original dataset, multi-plant firms account for approximately 10% of the observations. Excluding these tends to bias the units in our estimation sample towards smaller entities. Furthermore, Norwegian manufacturing firms are rather small in an international setting (see Table A1 in the Appendix for some descriptives). By focusing on single-plant firms only, we are sure that the price decisions are not made beyond the firm level. Thus the distinction between plants and firms, or the location of the pricing authority within a firm, is no issue in our model. Note also that in the remainder of the paper, we use the terms firm, producer and establishment interchangeably.
4. The data are representative for the Norwegian economy, and important moments resemble those of the Euro area as described by Vermeulen *et al.* (2012). During the sample period, inflation in Norway was low and stable; the annualized monthly consumer price index (CPI) had mean 1.8 and standard deviation 1.4. See also Wulfsberg (2016) for a study of Norwegian CPI data.
5. As control variables, we have included the same covariates as in the more structural model reported later. These are as follows: the natural logarithm of the wage rate, w, and its square for the previous year; year and monthly dummies; two commodity-group-specific dummies; and a monthly commodity-group-specific price index P_c for the relevant product. Note that these findings should be interpreted only as indicative evidence since the shares of both negative and positive price changes are relatively low as the sample is dominated by price inaction observations. Therefore the estimated effects on the predicted probabilities are driven by small local variations in the shares of either negative or positive price changes.
6. Eichenbaum *et al.* (2014) explain small price changes by measurement errors, while Nakamura and Steinsson (2010) use what they call a Calvo-Plus model to incorporate smaller price changes, and Dhyne *et al.* (2011) assume that the menu costs are stochastic to explain the same phenomenon.
7. We do not specify a full dynamic stochastic general equilibrium model. This is done in order to focus on firms' pricing decisions and not let the analysis be affected by possible misspecifications or problems in other parts of the macro-economy.
8. As mentioned in note 2, we do not observe all product prices. This means that the fixed menu cost a should in fact be divided by a higher number N_{it}. As a consequence, a downward bias is expected for our estimate of the parameter a. For that reason, our findings with respect to fixed menu costs should be interpreted as conservative. Furthermore, we abstract from asymmetry in the menu cost function. In the data, firms have price increases and decreases simultaneously. With asymmetric costs, a firm then incurs a fixed menu cost for both. As we focus on synchronization, where the total fixed costs of price changes are shared across price changes, we disregard this issue.
9. Note that the first-order conditions hold exactly in continuous time. We write the model in discrete time to facilitate bringing it to the monthly data.
10. Note that this price adjustment cost function is not twice differentiable at $\Delta P_{ijt} = 0$. In fact, in our case we have that $b = \lim_{\Delta P_{ijt} \downarrow 0} \frac{\partial C\left(\Delta P_{ijt}, P_{ijt-1}\right)}{\partial \Delta P_{ijt}} \neq \lim_{\Delta P_{ijt} \uparrow 0} \frac{\partial C\left(\Delta P_{ijt}, P_{ijt-1}\right)}{\partial \Delta P_{ijt}} = -b$. This does not constrain the possibility of finding a solution. The presence of linear price adjustment cost components implies that the adjustment cost function is not twice differentiable at $\Delta P_{ijt} = 0$. If $b = 0$, then it is twice differentiable everywhere.
11. In the case of substitution possibilities between the products, a product-specific shock induces internal price coordination to avoid cannibalization of the firm's own products. Product-specific shocks affect demand for a complementary product as well. Internal product market dependency may induce within-firm price coordination. Coordination of price change may be caused by menu costs providing scope advantages too. So disregarding a benefit of price coordination in the firm's profit function will mean that the estimates of the fixed menu cost will be smaller, to capture the benefits of coordination due to market dependency. Thus our estimates for the fixed menu costs, a, that cause coordination in our model are likely to be biased downwards.
12. In empirical factor demand models with quadratic adjustment cost components, it has been a standard assumption to abstract from these future adjustment cost savings in the q expression (see Abel and Blanchard 1986).
13. Setting prices may induce adjustment cost savings in next periods. We have performed an *ad hoc* test to see whether disregarding such cost savings in proxying for q is harmful. We included a dummy that takes value 1 if there has been a price change for the product in one of the two previous months. The included dummy might pick up future menu cost savings associated with the non-convex menu costs. The results of these exercises indicate statistical insignificance of the dummy.

14. Because the distribution of w_{it-1} is highly skewed, we had difficulty interpreting coefficients on the level of wages. Nevertheless, our menu cost estimates are hardly affected by the choice between taking the log or level of the wage rate.
15. Prices are the only variable at product or firm level reported on a monthly basis. All other variables are reported on an annual basis.
16. Latent class models are also referred to as semiparametric heterogeneity models and finite mixture models (Cameron and Trivedi 2005).
17. The use of two-stage estimation methods is recommended in more complicated models in which maximum likelihood is computationally burdensome (Maddala 1983, ch. 8). See also Nilsen *et al.* (2007) for a similar estimation procedure to analyse firm behaviour.
18. The identification problem related to the variance σ_{η}^2 is known for many limited dependent variable models. As the first-stage estimates are necessary for the second-stage estimation, we are therefore facing the normalization problem for the equations describing the intensive margins, that is, also in the second stage.
19. Note that when calibrating his model, Midrigan (2011) assumes only two products per retailer. Bhattarai and Schoenle (2014) consider a maximum of three products for firms in manufacturing industries.
20. Note that if $\widetilde{a} = 0$, then we have no exclusion restriction in the selection error correction term employed in the second step of the estimation procedure. So it is identified only by the functional form.
21. Not reported, but available from the authors on request.
22. This calculation is based on the linear and convex elements of the menu costs; $\widetilde{b}\Delta P \leq (\widetilde{c}/2)\Delta P^2/P$, so $\Delta P/P \geq 2\widetilde{b}/\widetilde{c}$.
23. This holds when $\widetilde{b} \cdot \Delta P \geq \widetilde{a}/m$.
24. For each product price regime, we calculate the probability at a given point in time based on parameter estimates of the ordered probit model. The probability is the unweighted average of these probabilities across product price, for each month.
25. When we observe the ability of the full model (columns (1) and (2) of Table 3) to predict small price changes, defined as greater than zero and less than or equal to 0.05 in absolute value as a share of non-zero price changes, this number is 90.0%, compared to 74.9% for the actual data. Thus our model has somewhat limited ability to predict the intensive margin of price changes.
26. Note that in equation (10) we replace $\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{m_{it}\cdot P_{ijt-1}}}$ by $\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{m_{it}\cdot P_{ijt-1}}} + \alpha \cdot \widehat{v}_{it}$ and $\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{(m_{it}+1)\cdot P_{ijt-1}}}$ by $\sqrt{\frac{2\cdot\widetilde{a}\cdot\widetilde{c}}{(m_{it}+1)\cdot P_{ijt-1}}} + \alpha \cdot \widehat{v}_{it}$.
27. Employing the latent class approach, we have accounted for the endogeneity issue to some extent already. An advantage of this approach is that potentially it reduces biasedness of parameters caused by the endogeneity of m_{it}. The reason is that m_{it} is also largely driven by a firm-level shock process. By including the latent class model generating firm-level shocks, this endogeneity issue is largely circumvented. When we estimate the model, with both the latent variable κ_{it} and the additional estimation error $\widehat{v}_{it}$ to account for endogeneity issues, the results point in the direction of minor deviations from the structural parameter estimates presented in columns (1) and (2) of Table 3. Such a finding is to be expected if the latent class approach has already reduced the potential endogeneity problem. As an additional exercise to address the endogeneity problem, we estimate a model corresponding to the model reported in columns (3) and (4) of Table 3, that is, without the latent variable κ_{it}, but where we now include the estimation error $\widehat{v}_{it}$ only. Again, estimation results are of the order of magnitude of the already reported estimates in columns (1) and (2) of Table 3.
28. Looking at the mean values of the prices in these industries, we find them to be NOK 1854, NOK 1932 and NOK 610 (Intermediate, Durables and Non-durables, respectively).
29. In addition, and in line with how we have modelled the q-process in our ML analyses, we induce a firm-specific shock to q that takes the value $\kappa = 2.57$ with probability $\psi = 0.047$; otherwise, it is zero. This firm-specific shock has only a contemporaneous effect, since it is added to the q-values of the underlying AR(1) q-process. The model is calibrated by searching for parameter values of the intercept and the AR(1) coefficient of the q-process (σ_v and ρ) that minimize the criterion function, the sum of squared deviations between the simulated moments and the same moments found in the data. These moments include the share of firm-level observations where no prices in a firm are adjusted, the share of observations where some of the product prices are changed, and the share of observations where all the product prices are changed. In addition, we add to the moment list the overall share of negative price changes, inaction, and positive price changes. Finally, the annual mean price increase for all products (annual inflation) is added. A course grid search gives that the constant term in the q-equation equals 0.0005, the AR(1) coefficient is 0.325, $\rho = 0.965$ and $\sigma_v = 0.99$. The latter coefficient is in line with the normalization that we have also used when estimating the ML model, $\sigma_v = 1.0$. The parameter ρ suggests that firm-specific shocks are important, and the AR(1) coefficient that some persistency is present.

REFERENCES

ABEL, A. B. and BLANCHARD, O. (1986). The present value of profits and cyclical movements in investment. *Econometrica*, **54**(2), 249–73.

ALTONJI, J. G., ELDER, T. E. and TABER, C. R. (2005). An evaluation of instrumental variable strategies for estimating the effects of Catholic schooling. *Journal of Human Resources*, **40**(4), 791–821.

ALVAREZ, F., LE BIHAN, H. and LIPPI, F. (2016). The real effects of monetary shocks in sticky price models: a sufficient statistic approach. *American Economic Review*, **106**(10), 2817–51.

ALVAREZ, F. and LIPPI, F. (2014). Price setting with menu cost for multiproduct firms. *Econometrica*, **82**(1), 89–135.

BARSKY, R. B., HOUSE, C. L. and KIMBAL, M. S. (2007). Sticky-price models and durable goods. *American Economic Review*, **97**(3), 984–98.

BHATTARAI, S. and SCHOENLE, R. (2014). Multiproduct firms and price-setting: theory and evidence from U.S. producer prices. *Journal of Monetary Economics*, **66**, 178–92.

BLINDER A. S., CANETTI, E. R. D., LEBOW, D. E. and RUDD, J. B. (1998). *Costs of adjusting prices. In* Asking About Prices: A New Approach to Understanding Price Stickiness. New York: Russell Sage Foundation.

BONOMO, M., CARVALHO, C., KRYVTSOV, O., RIBON, S. and RIGATO, R. (2020). Multi-product pricing: theory and evidence from large retailers in Israel. Working Paper no. 2020-12, Economic and Financial Research Department, Bank of Canada, Ottawa, Ontario, Canada.

CALVO, G. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, **12**, 383–98.

CAMERON, A. C. and TRIVEDI, P. K. (2005). *Microeconometrics, Methods and Applications*. Cambridge: Cambridge University Press.

CARD, D. and GIULIANO, L. (2013). Peer effects and multiple equilibria in the risky behavior of friends. *Review of Economics and Statistics*, **95**(4), 1130–49.

CARLSON, J. A. (1992). Some evidence on lump sum versus convex costs of changing prices. *Economic Inquiry*, **30**(2), 322–31.

CARLSSON, M. (2017). Microdata evidence on the empirical importance of selection effects in menu-cost models. *Journal of Money Credit and Banking*, **49**, 1803–30.

DHYNE, E., FUSS, C., PESARAN, M. H. and SEVESTRE, P. (2011). Lumpy price adjustments: a microeconometric analysis. *Journal of Business & Economic Statistics*, **29**(4), 529–40.

DIXIT, A. (1991). Analytical approximations in models of hysteresis. *Review of Economic Studies*, **58**(1), 141–51.

EICHENBAUM, M., JAIMOVICH, N., REBELO, S. and SMITH, J. (2014). How frequent are small price changes? *American Economic Journal: Macroeconomics*, **6**, 137–55.

GOLOSOV, M. and LUCAS, R. Jr (2007). Menu costs and Phillips curves. *Journal of Political Economy*, **115**, 171–99.

KLENOW, P. J. and MALIN, B. A. (2010). Microeconomic evidence on price-setting. In B. M. Friedman and F. H. Hahn (eds), *Handbook of Monetary Economics*, Vol. **3**. Amsterdam: Elsevier, pp. 231–84.

LEVY, D., BERGEN, M., DUTTA, S. and VENABLE, R. (1997). The magnitude of menu costs: direct evidence from large U.S. supermarket chains. *Quarterly Journal of Economics*, **112**(3), 791–825.

MADDALA, G. S. (1983). *Limited Dependent and Qualitative Variables in Econometrics*. Cambridge: Cambridge University Press.

MIDRIGAN, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, **79**(4), 1139–80.

NAKAMURA, E. and STEINSSON, J. (2010). Monetary non-neutrality in a multisector menu cost model. *Quarterly Journal of Economics*, **125**, 961–1013.

NILSEN, Ø. A., SALVANES, K. G. and SCHIANTARELLI, F. (2007). Employment adjustment, the structure of adjustment costs, and plant size. *European Economic Review*, **51**(3), 577–98.

NILSEN, Ø. A., SKUTERUD, H. and WEBSTER, I. M.-K. (2021). Price change synchronization within and between firms. *Economics Letters*, **208**, article 110067.

NILSEN, Ø. A., and VANGE, M. (2019). Intermittent price changes in production plants: empirical evidence using monthly data. *Oxford Bulletin of Economics and Statistics*, **81**(1), 98–122.

NOU (2019). *Norges Offentlige Utredninger—NOU 2019:6 Grunnlaget for inntektsoppgjørene (Official Norwegian Reports 2019:6 The basis for wage settlements)*. Ministry of Labour and Social Affairs, Oslo, Norway.

RIVERS, D. and VUONG, Q. H. (1988). Limited information estimators and exogeneity tests for simultaneous probit models. *Journal of Econometrics*, **39**(3), 347–66.

ROTEMBERG, J. and WOODFORD, M. (1999). The cyclical behavior of prices and costs. In J. Taylor and M. Woodford (eds), *Handbook of Macroeconomics*, Vol. **1B**. Amsterdam: Elsevier, pp. 1051–135.

SHESHINSKI, E. and WEISS, Y. (1977). Inflation and the costs of price adjustment. *Review of Economic Studies*, **44**, 287–303.

——— (1983). Optimum pricing policy under stochastic inflation. *Review of Economic Studies*, **50**, 513–29.

——— (1992). Staggered and synchronized price policies under inflation: the multiproduct monopoly case. *Review of Economic Studies*, **59**(2), 331–59.

SSB (2013). Commodity price index for the industrial sector; available online at https://www.ssb.no/en/priser-og-prisindekser/statistikker/vppi?fane=om (accessed 16 February 2022).

STELLA, A. (2020). Economies of scope in price setting: a moment inequalities estimation. *Journal of Monetary Economics*, **110**, 50–61.

TSOUKIS, C., KAPETANIOS, G. and PEARLMA, J. (2011). Elusive persistence: wage and price rigidities, the new Keynesian Phillips curve and inflation dynamics. *Journal of Economic Surveys*, **25**(4), 737–68.

Vermeulen, P., Dias, D., Dossche, M., Gautier, E., Hernando, I., Sabbatini, R. and Stahl, H. (2012). Price-setting in the euro area: some stylised facts from individual producer price data. *Journal of Money, Credit and Banking*, **44**(8), 1631–50.

Wolman, A. L. (2007). The frequency and costs of individual price adjustment. *Managerial and Decision Economics*, **28**(6), 531–52.

Wooldridge, J. M. (2014). Quasi-maximum likelihood estimation and testing for nonlinear models with endogenous explanatory variables. *Journal of Econometrics*, **182**, 226–34.

——— (2015). Control function methods in applied econometrics. *Journal of Human Resources*, **50**(2), 420–45.

Wulfsberg, F. (2016). Inflation and price adjustments: evidence from Norwegian consumer price data 1975–2004. *American Economic Journal: Macroeconomics*, **8**(3), 175–94.

Yang, C. (2019). *Rational inattention, menu costs, and multi-product firms: micro evidence and aggregate implications*. Mimeo, University of Texas at Austin.

Zbaracki, M. J., Ritson, M., Levy, D., Dutta, S. and Bergen, M. (2004). Managerial and customer costs of price adjustment: direct evidence from industrial markets. *Review of Economics and Statistics*, **86**(2), 514–33.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Supporting Information