CESifo WORKING PAPERS

7145
2018

July 2018

Intermittent Price Changes in Production Plants: Empirical Evidence Using Monthly Data

Øivind A. Nilsen, Magne Vange

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7145*
Category 7: Monetary Policy and International Finance

Intermittent Price Changes in Production Plants: Empirical Evidence Using Monthly Data

Abstract

The price-setting behaviour of manufacturing plants is examined using a large panel of monthly surveyed plant- and product-specific prices. The sample shows a high frequency of zero changes, relatively small price changes, and a strong seasonal price-change pattern. The intermittent feature of price changes is modelled with thresholds which are smaller in January, and a quadratic loss function associated with the distance from the target price. The findings show statistically significant pricing thresholds, which are only two-thirds in January, and partial adjustment parameters implying that 60% of the deviation between the target price and the current price is closed each month.

JEL-Codes: E300, E310, E370.

Keywords: price setting, micro data, simulated method of moments.

*Øivind A. Nilsen**
Norwegian School of Economics
Department of Economics
Norway - 5045 Bergen
oivind.nilsen@nhh.no

Magne Vange
Norwegian School of Economics
Department of Economics
Norway - 5045 Bergen
magnevange@gmail.com

*corresponding author

May 2018
We would like to thank Wilko Letterie, Matthias Meier, Dominik Menno, and Michael Weber, as well as participants at RES 2017 in Bristol, the Nordic Econometric Meeting 2017 in Tartu, EARIE 2017 in Maastricht, SAEe 2017 in Barcelona, and at the 12th Workshop on Macroeconomics and the Business Cycle at Ifo Institute, Dresden Branch. The paper has also benefited from thorough discussions with and the assistance of Eivind Bøe when testing out the ideas using annual data. We are also grateful to Magne K. Asphjell for providing help with preparing the dataset. Part of this work was done while Nilsen visited the Institute for Macroeconomics and Econometrics— University of Bonn, whose hospitality is gratefully acknowledged. The usual disclaimer applies.

1 Introduction

Most modern macroeconomic models assume price stickiness, i.e., that price setters are faced with frictions, without sufficient knowledge about the underlying microeconomic implications. This calls for a further empirical assessment of the theoretical premises in the macroeconomic models we use today.

A common method for analysing price stickiness is to investigate the role of thresholds in the pricing patterns of individual firms. In this literature, the (S, s) rule, proposed by Sheshinski and Weiss (1977), plays an important role. These authors argue that firms kept the price fixed within certain bounds, denoted (S, s). As a result, prices exhibit a pattern of inaction followed by large price changes, so called "zeroes and lumps". The authors argue that this pattern is caused by the fact that changing the price induces a fixed cost for the firm, which is referred to as the menu cost. The (S, s) methodology has been adopted and further extended by many and, thereby, represents a large share of the current price stickiness literature (see, e.g., Caballero and Engel, 1993; Ratfai, 2006; Alvarez et al., 2011; Dhyne et al., 2011). An essential assumption in these models is that adjustment costs are independent of the size of the price change.

One aspect to consider when searching for thresholds in pricing patterns is whether the thresholds are symmetric, i.e., if the magnitudes of the thresholds are the same upwards and downwards. A study on microeconometric evidence from Switzerland by Honoré et al. (2012) finds a smaller upper than lower threshold. According to this study, price changes are more likely to be positive than negative, *ceteris paribus*. The study ignores, however, the magnitude of price changes, because only the frequency and the duration of inaction are accounted for. Loupias and Sevestre (2012), on the other hand, include the magnitude of price changes, and find that when firms face cost variations, they appear to adjust their prices more often and more rapidly upwards than downwards.

The counterpart of the (S, s) methodology in the price-stickiness literature assumes that the adjustment cost is a convex function of the size of the price change, i.e., that greater changes lead to higher costs (Rotemberg, 1982). Whereas the assumption of fixed costs implies that one should observe large and infrequent price changes, the convex-cost assumption implies the opposite: frequent changes of small size. As emphasized by Zbaracki et al. (2004), most of the literature finds evidence supporting the former. However, if there are only fixed and not convex price-adjustment costs, we fail to see why the pricing data show a relatively high proportion of small price changes.[1] Earlier research with (S, s) pricing rules has, in part, failed to include small price changes.

As highlighted by Klenow and Malin (2011), access to good microeconomic data is crucial, and is a common problem in all empirical research related to pricing. The basis of our analysis is monthly collected micro price data for Norwegian manufacturers. Although consumer prices are relevant for the monitoring of inflation by central banks, the prices at the producer level are most often modelled into the macroeconomic policy models (Vermeulen et al., 2012). Accordingly, knowledge about producer price adjustments is essential to improve macroeconomic modelling and central bank policies.

In this paper, we propose a model where the adjustment towards the new price is conditional on both thresholds and partial adjustments. Thus, our model—in contrast to many other models—therefore allows for both inaction and inertia in pricing. The hypothesis is that there are fixed costs associated with setting a new target price. There are also two convex components: one associated with deviation from the new target price and another that increases with the scale of the price change. These

[1] The study of Eichenbaum et al. (2014) on CPI data suggests that the observation of small price changes is largely due to measurement errors and quality adjustments, and should therefore be neglected. However, that study's findings are contradicted by a vast majority of empirical research suggesting that small price changes are relatively common (Klenow and Kryvtsov, 2008; Bhattarai and Schoenle, 2014; Midrigan, 2011; Wulfsberg, 2016).

latter convexities make the firm favour slow adjustments and small price changes, and might be due to convexity of customer and managerial costs. Customer costs because a larger upward price change may lead customers to search for more attractive outside offers. When it comes to managerial coosts, Zbaracki et al. (2004) observe that:

"*The greater the proposed price change, the more people are involved, the more supporting work is done, and the more time and attention is devoted to the price change decisions*" (see their pp. 523-524)

Thus, our model sets out to explain both the occurrence of price adjustments of different sizes, and inaction. The model is tested on a dataset based on survey data behind the commodity price index for the Norwegian industrial sector (PPI). These data include monthly price quotations for a representative sample of Norwegian plants. In contrast to, for instance, Ratfai (2006), whose sample includes eight outlets only, our sample includes more than 350 different producers. The data show a high frequency of price change inaction, by relatively small price changes when changed, and by a much higher occurrence of price changes in the beginning compared with the end of a year. To analyse the intermittent nature of the model, a simulated method of moments is used. This advantages estimations based on, for instance, maximum likelihood methods, which are often based on quite restrictive assumptions. The estimations reveal thresholds such that prices are changed only if the deviations from the underlying frictionless prices are approximately 15%. When changed, the prices are changed rather quickly with only 10% of the initial gap existent after three months. The asymmetry between upward and downward rigidities is minor but statistically significant. Finally, the thresholds in January are approximately two-thirds compared with the other months.

The remainder of the paper is organized as follows. Section 2 describes the data, whereas the model, method, and moments are presented in Section 3. Section 4 reports and discusses the results and Section 5 gives some concluding remarks.

2 Data

The basis for our empirical analysis is the survey data behind the commodity price index for the Norwegian manufacturing industry (PPI) obtained from Statistics Norway (SSB).[2] The data are collected on a monthly basis for a selection of Norwegian plants.[3] Plants with more than 100 employees are included in the sample at all times, and the selection of producers is updated continuously, securing a high level of relevance (SSB, 2015). Plants are repeatedly surveyed, participation is compulsory, and Statistics Norway revises the data regularly to detect measurement errors and nonconformity.[4] Considering this, and that the PPI is an important tool for governing bodies, it is fair to assume that the data are representative for Norwegian producers and of high quality.

The initial dataset contains price observations ranging from year 2002 until 2009. In the construction of the final dataset for this study, plants with observations for less than 24 months have been omitted, as well as plants with less than 10 employees. Furthermore, only years with observations for all months in a given year are included. Due to the implementation of a new sampling procedure at Statistics Norway, there was a clear shift in the reported price change frequency in 2004. We therefore discard the data prior to January 2004. Furthermore, plants related to the energy sector (oil, gas, electricity, etc.), and mining and quarrying have been left out of the sample because they are known to have an abnormally high adjustment frequency. The original dataset contains prices for both domestic and export markets, but to prevent interference by exchange- rate movements and international competition, export market prices are omitted. Additionally, because very large price changes are likely to

[2] See SSB (2015) for more information about the PPI.

[3] In the remainder of the paper, we use the terms plant, firm, producer, and establishment interchangeably.

[4] One plant might be recorded with one or multiple products. It should be noted that for data collection purposes, firms may be targeted for some, but not all of the products they manufacture. If Statistics Norway regards a subset of the products to be important to obtain an accurate estimate of the price index, data will be requested for these ones only.

reflect changes to design or quality of the product rather than to common pricing decisions, price growth observations outside the [0.01, 0.99] interval we consider to be new products. Finally, we focus on single-plant firms only.[5] This leaves us with a final sample of 76,804 observations for 1,676 products over the years 2004–2009 covering 21 two-digit SIC2002 industry codes.

2.1 Descriptives

[Figure 1 "Distribution of Price Change Rates" about here]

Figure 1 shows the proportion of observations in different price change intervals, both for the actual data and for the later preferred simulated model (black and grey columns, respectively). Observations with price changes with absolute values less than 0.005 represent the majority of the dataset (80%).[6] In other words, most observations are characterized as price change inaction (later, we refer to this as the "zero price-change" interval). This indicates the existence of fixed or non-convex price-adjustment costs. At the same time, we observe a substantial proportion of small price-change observations, i.e. $0.5\% < \left|\frac{\triangle p}{p}\right| \leq 5.0\%$. If there is only a fixed cost independent of the magnitude of the price change, one would not expect to see these small price changes.[7] This observation could, however, be an indication of convex adjustment costs, which put a penalty on large adjustments and, thereby, force the producers to adjust gradually. The observation of several periods of inaction, combined with series of small price changes, may tell a story of firms being faced with both non-convex and convex price adjustment costs.

[5] With this choice, we are sure that the price decisions are not made beyond the plant level.

[6] The percentage of exactly zero prices in the actual data is 75%. This is in line with numbers from the Euro area and the USA (see Vermeulen et al. (2012), Table 2). In the simulated data, there are no observations with exactly zero changes because of numerical precision.

[7] Our definition of small price changes (less than 5% in absolute value) is consistent with the assumptions of Klenow and Kryvtsov (2008) and Eichenbaum et al. (2014). Note, however, as recognized by these latter authors: *"The definition of what constitutes a "small" price change is, inevitably, somewhat arbitrary"* (p. 138). They, therefore, study small price changes defined as 1%, 2.5%, and 5%, whereas Midrigan (2011) uses thresholds of 3% and 5%.

To identify lumpy adjustment behaviour, we rank, for each product and each year, the 12 monthly price changes from lowest to highest.[8] Rank 1 thereby represents the largest monthly price change, Rank 2 the second largest price change, and so on. For each rank, we then calculate the average price change over all products and all years. The intuition is that if there is a large gap between the largest (smallest) and the second largest (second smallest) price change compared with the other ranks, this indicates that producers are faced with fixed costs of adjustment and, therefore, change the price quite substantially when first changing it. Otherwise, with normally distributed shocks to the fundamentals, and no adjustment costs, one would expect the mean price change of adjacent observations to be rather similar and, therefore, that there is a downward-sloping linear relationship between the ranks (for more details, see Doms and Dunne 1998).

[Figure 2 "Ranked Price-Change Rates" about here]

Figure 2 shows the ranking of the monthly price changes. As seen from the figure, there is a gap of approximately three percentage points between the first and second ranks, and two percentage points between the eleventh and twelfth ranks. In contrast, the differences between the intermediate ranks are modest. As already pointed out, this is consistent with non-convexities in adjustment costs, even though coexistence of both fixed and convex adjustment costs cannot be excluded. That means, even if fixed adjustment costs are preventing the firms from adjust continuously, when they actually do change their price, convex costs are forcing them to do so gradually.[9]

[Figure 3 "The Occurrence of Price Changes by Months" about here]

Figure 3 shows the average frequency of price-change quotations greater than

[8] Such a measure has been used in the investment and labour demand literature (see, for instance, Doms and Dunne (1998), Nilsen and Schiantarelli (2003), and Varejão and Portugal (2007)).

[9] We also observe that all the ranks are shifted to the left, because only rank five is below zero. This is expected because inflation will cause the producers to have more positive price changes than negative ones.

$|0.005|$ within each month. There is a relatively high price-change frequency in the beginning of the year compared with the remaining months, a pattern also described by Nakamura and Steinsson (2008) and Vermeulen et al. (2012). This seasonality could be explained by the producers' economic environments, for instance, seasonal demand effects. Furthermore, it may be explained by the costs of information acquisition and processing (see Maćkowiak and Wiederholt (2009) and Mankiw and Reis (2002)), and the pricing season effect related to negotiation and of signing of price contracts, described by for instance, Zbaracki et al. (2004). Finally, it is also consistent with the theories focusing on staggered contracts (see, for instance, Taylor (1980, 1999)) with a duration of one year, and that a majority of these contracts start in January.

3 Model, Method, and Moments

As already discussed, several theories have been proposed to explain the intermittent price-adjustment patterns observed in many datasets at both the consumer level and at the producer level. Here, we suggest a simple reduced-form model that describes the price-adjustment behaviour of production plants with the following three features: plants adjust prices infrequently where only 20% of the price observations change from one month to another; there are a lot of small price changes; and there is a seasonal pattern in the incidence of price changes, with most price changes taking place in January.

3.1 Model Specification and Predefined Parameters

Because firms require a degree of monopoly power to be able to set prices, we assume that producers operate in monopolistic competitive markets. Furthermore, it is assumed that each firm is able to continuously observe and monitor its frictionless price without any costs.

We start from the observations of high frequencies of zero price adjustments. This would be observed if there were some menu costs, and if it was costly to continuously adjust the product prices. The firm operates with a target price (in logs) for product i at time t, denoted by $p_{it}^{\#}$, and leaves this unchanged unless the distance from the frictionless price p_{it}^{*} (also in logs) becomes too large. The latter price p_{it}^{*} represents the frictionless equilibrium price if there are no price-change costs. The costs associated with setting a new target price is $F \cdot I\left(p_{it}^{\#} \neq p_{it-1}^{\#}\right)$, where F is the actual cost and $I\left(p_{it}^{\#} \neq p_{it-1}^{\#}\right)$ is an indicator function. The formation of the target price is determined by:

$$p_{it}^{\#} = \begin{cases} p_{it}^{*} & if \ \ | \, p_{it}^{*} - p_{it-1}^{\#} \, |> \tau \\ p_{it-1}^{\#} & otherwise \end{cases} \tag{1}$$

where τ denotes a threshold. Thus, if the shock to the frictionless price is large enough relative to its target value, in absolute value, the firm finds it profitable to set a new target price and to start to adjust its price. The formulation in eq. (1) states that the threshold is symmetric, i.e., that the "band of inaction" is the same whether the price shock is positive or negative. We relax this restriction and allow the thresholds for price increases and price decreases to be different. With this modification, the formation of the target price is determined by:

$$p_{it}^{\#} = \begin{cases} p_{it}^{*} & if \ p_{it}^{*} - p_{it-1}^{\#} > U \ or \ p_{it}^{*} - p_{it-1}^{\#} < L, \\ p_{it-1}^{\#} & otherwise \end{cases} \tag{2}$$

where U denotes the upper threshold and L denotes the lower threshold, i.e., $L \leq 0 \leq U$. It means that the target price is changed only if the frictionless price moves outside the interval determined by L and U. Otherwise, if the frictionless price is greater than L and smaller than U, the producer leaves its target price $p_{it}^{\#}$ unchanged.

Following Alvarez et al. (2011); Nakamura and Steinsson (2008) and others, we let the logarithm of the frictionless nominal price for product i at time t, denoted by p^*_{it}, follow a random walk with drift:

$$p^*_{it} = \alpha + p^*_{it-1} + \varepsilon_{it}, \qquad \textit{where } \varepsilon_{it} \sim N(0, \sigma^2_\varepsilon) \tag{3}$$

The random walk process implies that the frictionless price is adjusted immediately as a consequence of new information in the idiosyncratic shocks ε_{it} with variance σ^2_ε. And unless new information arrives, the frictionless price will follow α the deterministic drift. This latter component is meant to capture trend inflation either in output prices or input factors—or productivity growth. Thus, current frictionless price (plus the trend α) is the best prediction about next periods' frictionless price. If α were not included, the trend inflation would be embedded into the threshold parameters and, therefore, bias the results. If α is set too low (too high) compared with the actual trend inflation, the estimated threshold parameters L and U would be biased downwards (upwards). The deterministic trend parameter α, is set as close to the actual inflation as possible to limit the effect of inflation bias.[10] The idiosyncratic shock parameter, ε_{it}, is meant to reflect any shocks to either demand, cost, or technology excess of the underlying trend captured by the trend parameter α.[11] It is, of course, possible to allow for serial correlation in ε_{it}, and thus, to make the frictionless price less persistent. However, as a richer specification would require more parameters to be estimated and therefore raise additional identification issues, and also make the exposition more complicated, we keep the simpler process.

[10] We could also include α as a parameter to be estimated. A simpler approach, adopted here, is to perform a series of simulations with different values for α. We comment on this further when testing the robustness of our model.

[11] An argument against letting demand shocks from technology and/or costs be treated identically is that firms react quicker to positive than to negative cost shocks, but slower to positive than to negative demand shocks (see, for instance, Dias et al. (2015) and Loupias and Sevestre (2012)). Differencing between types of shocks would require a more sophisticated model than the one presented here. Furthermore, with the available data, it might be difficult to identify different types of shocks. Thus, we choose to include the shocks as an aggregate effect (the Appendix shows how the prices are affected by demand, technology, and cost shocks.).

Even though the firm has decided to change its target price and, therefore, also to change the price of the given product, it does not move directly to the new target price. As discussed by Zbaracki et al. (2004), price-change costs also include convex components such that it is costlier to make one greater price change compared with several smaller ones. Two characteristics of the data support the convexity assumption. First, the descriptive evidence given in Figure 1 indicates that small prices changes are not that uncommon, whereas greater price changes are. Second, a more detailed look at the actual data states that the probability of observing a price change, conditional that a price change also took place during the previous period, is 0.525. Comparing this with the unconditional probability of a price change of 0.200, it is a forceful indication that the price changes in the data are strongly correlated over time. In addition to the convexity of costs related to price changes, there might also be losses for being too far away from the new target price $p_{it}^{\#}$. A formulation that encompasses both these elements, given that the deviation between the new frictionless price, p_{it}^{*}, and the target price, $p_{it}^{\#}$, is large enough to initiate price changes, the "out-of-equilibrium costs" is as follows:

$$AC(p_{it}) = C \cdot \left\{ (1-\theta) \left(p_{it} - p_{it}^{\#} \right)^2 + \theta \left(p_{it} - p_{it-1} \right)^2 \right\} \tag{4}$$

Thus, the formulation consists of a weighted sum (where $0 \leq \theta \leq 1$) of two quadratic terms, which denotes the difference between the new price and the target price, and the difference between the new price and the previous price, respectively. A plant seeks to minimize these "out-of-equilibrium costs" $AC(.)$. The first-order condition of equation (4) with respect to the new current price p_{it} rearranged is therefore:

$$(p_{it} - p_{it-1}) = (1-\theta) \left(p_{it}^{\#} - p_{it-1} \right) \tag{5}$$

Thus, we have the traditional partial adjustment model where the "out-of-equilibrium costs" $AC(.)$ prevent the producer from adjusting immediately to its target price, except that the "usual" frictionless seen in partial adjustment models is exchanged with the target price $p^{\#}_{it}$. An implication is that the producer will close $(1-\theta)$ of the deviation between the target price and the old actual price. For example, $\theta = 0.10$ will indicate that the producer closes 90% of the desired price change in the first period. If the target price remains unchanged in the subsequent period, the producer will close 90% of the remaining price gap. This will continue until the producer decides to set a new target price or when the target price is reached.

To avoid the restriction that the weights in equation (4), $(1-\theta)$ and θ, are common to price increases and price decreases, we allow for asymmetric inertia in addition to the already-discussed asymmetric thresholds. This reflects asymmetric adjustment costs discussed and analysed in the microeconomic literature (e.g., Peltzman, 2000; Yang and Ye, 2008; Lewis, 2011; Loy et al., 2016).[12] If the price is increasing, θ_{up} is supposed to capture upward inertia and, conversely, if the price is decreasing, θ_{down} is supposed to capture downward inertia. Thus, we let the logarithm of the nominal price of product i at time t be given by:

$$(p_{it} - p_{it-1}) = \begin{cases} (1-\theta_{up})(p^{\#}_{it} - p_{it-1}) & if\ p^{\#}_{it} - p_{it-1} > 0.005, \\ 0 & if\ |p^{\#}_{it} - p_{it-1}| \leq 0.005, \\ (1-\theta_{down})(p^{\#}_{it} - p_{it-1}) & if\ p^{\#}_{it} - p_{it-1} < -0.005, \end{cases} \tag{6}$$

As seen from (6), we associate price deviations relative to the target price within the $[-0.005, 0.005]$ interval with zero price changes because such minor deviations are likely to be of little economic importance.[13] Furthermore, the numerical simulations described in a subsequent section, make it necessary to define very small actual or desired price changes as inaction.

[12]See also Laxton et al. (1999); Dolado et al. (2005); Dobrynskaya (2008).

[13]Managers are likely to abstain from closing minor price gaps also because of the uncertainty related to the manager-calculated target prices. Such uncertainty exists because the target prices are functions of current and future general market conditions, and of static, dynamic, and strategic (price, quantity, technology, and input prices) considerations of both the firm itself and its competitors (see also Mankiw and Reis (2010), p. 190 and their FN 6).

It should be mentioned that if $\theta_{up} \neq 0$, $\theta_{down} \neq 0$, $U = L = 0$, the model specification reduces to a partial adjustment model (because then $p^*_{it} = p^{\#}_{it}$ and eqs. (1) and (2) would be irrelevant). Conversely, if $\theta_{up} = \theta_{down} = 0$, $U \neq 0$, $L \neq 0$, the model reduces to a (S, s) pricing model.[14] Note also, that if the target price $p^{\#}_{it}$ had not been introduced explicitly, the (symmetric) threshold specification would be as follows:

$$p_{it} = \begin{cases} p^*_{it} & if \ \mid p^*_{it} - p_{it-1} \mid > \tau \\ p_{it-1} & otherwise \end{cases} \tag{7}$$

Furthermore, the (symmetric) partial adjustment expression would then be:

$$(p_{it} - p_{it-1}) = \begin{cases} (1-\theta)(p^*_{it} - p_{it-1}), & if \ |p^*_{it} - p_{it-1}| > 0.005, \\ 0 & if \ |p^*_{it} - p_{it-1}| \leq 0.005, \end{cases} \tag{8}$$

Note however, that the price-adjustment process would stop when $p^*_{it} - p_{it-1}$ reaches the threshold τ. Thus, without the target price $p^{\#}_{it}$ in the model, we would not observe many small price changes and p^*_{it} would never be fully reached.[15]

[Figure 4: Illustration of Price-Change Process - about here]

In Figure 4, we illustrate how our model works. Starting with the evolvement of the frictionless price, p^*, we see clearly the upward trend, but with an interim period with sudden price decreases. The thresholds have a constant distance relative to the actual price (bold line). We see that, in period $t = t^A$, the frictionless price has evolved such that it is greater than the upper threshold U and, consequently, the target price and the actual price both changed. We see, however, that the actual price is moving slowly towards the new target price. This is caused by the inertia parameter(s) θ. In period $t = t^B$, a sudden negative shock occurs, pushing the target price below the lower threshold L. Subsequently, the price reaches this new target

[14] Eq. (5) then states that, conditional on changing, one immediately goes to the new target price.

[15] Dhyne et al. (2011) have such a model, but with asymmetries. To be able to incorporate the existence of small price changes, they let the threshold be stochastic.

price. Thus, we see that price changes can be caused by accumulated small shocks, or one large shock to the underlying frictionless price. We also see intermittence, and small price changes consistent with the descriptive statistics. Finally, the figure shows that the thresholds change across time.

The values of the descriptive statistics show that the incidence of price changes is 31% in January, whereas the average over the other eleven months is 19%. To control for this seasonal effect, we include a January-specific parameter, defined as $0 \leq y \leq 1$, which is multiplied with the threshold parameters U and L if the current month is January. This decreases the thresholds in the beginning of each year and, thereby, increases the probability of a price change. Furthermore, this might also reflect the potential existence of staggered contracts starting in January and with 12 months' duration.[16]

This leaves us with the following parameters to be estimated:

$$\begin{aligned} \text{Upper threshold: } & U \\ \text{Lower threshold: } & L \\ \text{Inertia upwards: } & \theta_{up} \\ \text{Inertia downwards: } & \theta_{down} \\ \text{January-specific scalar: } & y \\ \text{Standard deviation of idiosyncratic shocks: } & \sigma_{\varepsilon} \end{aligned}$$

In our main estimates, we set α =0.0025, which gives an annual inflation equal to 0.03, close to the average annual inflation rate of the producer price index (PPI) between the years 2004 and 2009. We have also made the mean and standard deviation of the initial (log) frictionless price, p^*_{i0}, corresponding to the distribution of the actual prices in June 2006 (after having taken into account the underlying annual inflation).

[16]Nilsen et al. (2016) show, using the same data used in this paper, a flat price change hazard with a peak after 12 months.

3.2 Estimation Method

Given that the empirical model includes the thresholds, the model does not have an analytical closed-form solution. This again prevents us from using "standard" regression techniques. We therefore use a simulated method of moments (SMM). In short, SMM seeks to minimize the distance between two sets of moments—the moment vector generated conditional on a vector of parameters to be estimated β, and the corresponding moment vector in the actual data, i.e., to find the vector of l unknown parameters β that minimizes the following quadratic form $J(\beta)$:

$$J(\beta) = [\Phi^A - \frac{1}{\kappa}\sum_{j=1}^{\kappa}\Phi^S(\beta)]' W [\Phi^A - \frac{1}{\kappa}\sum_{j=1}^{\kappa}\Phi^S(\beta)] \tag{9}$$

where Φ^A and $\Phi^S(\beta)$ denote the vectors of m actual moments and simulated counterparts, respectively; W denotes an optimal weighting matrix; and κ denotes the number of panels with the same size as the actual data. The distance between two sets of moments $J(\beta)$, has a χ^2 distribution with $m - l$ degrees of freedom, where m is the number of moments.[17]

3.3 Selection of Moments

The model should explain both inaction and small price changes at the same time. Thus, the proportion of observations within the following intervals are included:

$$\begin{aligned}
-0.050 \leq p_{it} - p_{it-1} &< -0.025 \\
-0.025 \leq p_{it} - p_{it-1} &< -0.005 \\
-0.005 \leq p_{it} - p_{it-1} &\leq 0.005 \\
0.005 < p_{it} - p_{it-1} &\leq 0.025 \\
0.025 < p_{it} - p_{it-1} &\leq 0.050
\end{aligned} \tag{10}$$

[17] See the Appendix for more details about the simulated method of moments approach.

These moments should contribute to identifying all the parameters, especially the threshold parameters and inertia parameters: non-zero U and L will cause zero-inflated price changes, and positive θ_{up} and θ_{down} will cause small price changes.

Larger inertia parameters, θ_{up} and θ_{down}, will make plants smooth their adjustments over time, which indicates that there will be several consecutive periods of small price changes. A consequence of this gradual adjustment is serial correlation in price changes. We therefore choose to include the following correlation coefficient moments:

$$Corr[p_{it} - p_{it-1}, p_{it-1} - p_{it-2}] \ if \ p_{it} - p_{it-1} > 0.005 \tag{11}$$

$$Corr[p_{it} - p_{it-1}, p_{it-1} - p_{it-2}] \ if \ p_{it} - p_{it-1} < -0.005 \tag{12}$$

On the other hand, the threshold parameters U and L will also be affected by these moments, as larger $|U|$ and larger $|L|$ will lead to more inaction and smaller serial correlation. The asymmetry is such that the moment in (11) should identify θ_{up}, whereas the moment in (12) should identify θ_{down}.

The standard deviation of the shocks to the frictionless price, σ_ε, is likely to be directly related to the standard deviation of price changes, $sd(p_{it} - p_{it-1})$. We therefore choose to include the standard deviation of price changes as a moment. The standard deviation of price changes is also likely to be affected by the friction parameters: as already pointed out, larger $|U|$ and $|L|$ lead to more inaction and, thereby, smaller variance of the observed price changes. Equation (5) shows that greater values of θ lead to price changes of more similar size which, again, will reduce the variance of price changes. Thus, the standard deviation of price changes will not only identify σ_ε, but will also contribute to the identification of U, L, θ_{up} and θ_{down}.

The January-specific scalar, y, is supposed to capture the abnormally high adjustment frequency in the beginning of the year. As a primary identifier, the following moment is therefore included:

$$\frac{\text{Number of price quotations with } |p_{it} - p_{it-1}| > 0.005 \text{ in January}}{\text{Total number of price quotations in January}}$$

As previously mentioned, ranked price changes can be a good indicator of lumpy adjustment behaviour. We include the first two and the last two ranks as moments. These are meant to be the primary identifiers of the threshold parameters U and L. The ranks are likely to be affected by σ_ε and the inertia parameters as well: more variation in the frictionless price will cause more variation in the ranks and greater inertia parameters will bring the ranks closer to each other. Hence, the rank moments will also affect θ_{up}, θ_{down} and σ_ε.[18]

One might think that the use of ranks is just another way of describing the seasonal effects and, thus, that there is not much gain in adding ranks for identification. When holding for each month the share of the highest-ranked price changes (rank 1 observations), the evolvement of these shares mimics very much the frequency of price changes by months. Formal testing on monthly aggregates shows a correlation coefficient of 0.99 (and z-value = 21.7). On the other hand, the correlation of the share of the lowest- ranked price changes, and the frequency of price changes by month is small (0.11) and statistically insignificant. Thus, there are likely benefits with regard to identification from including both the January effect and information about the highest and lowest ranks.

[18]The moments in our data (Table 1 Column 8) are quite close to those found in other relevant studies. The frequency of zero-price changes is 75.0% in our actual data, which is quite similar to corresponding numbers from other studies using PPI data. For example, Vermeulen et al. (2007, Table 2) report 79.2% in the Euro area, Nakamura and Steinsson (2008, Supplementary Material Table 11) and Bhattarai and Schoenle (2014, Figure 1) report 71% to 81% for the US. Regarding the timing of price changes, in our data 31.5% (19.0%) of all price changes take place in January (February to December). The seasonal variation is supported by the studies on the Euro area and the US. Vermeulen et al. (2007) report 31.2% (20.2%) frequency of changes in January (February to December), while Nakamura and Steinsson (2008, Supplementary material, page 10) state that "producer prices are more than twice as likely to change in January than on average in other months of the year".

4 Results

Table 1 shows the parameter estimates for the various model specifications by columns. Standard errors are presented in parentheses. The parameter estimates and standard errors are presented in the upper part of the table. The corresponding moments of the various model specifications are presented in the lower part of the table. We estimate all specifications against the 13 moments already described, i.e., distribution of $\triangle p/p$, January effect, serial correlations, standard error of $\triangle p/p$, and rank moments.

[Table 1 "Parameter Estimates and Moments" about here]

Starting with a broad look at this table, we see that all the estimated parameters are statistically significant (with the exception of those in Column (5)). In Column (1), we report the results of the full model which include both thresholds and partial adjustment parameters, and that their magnitudes depend on whether prices are increasing/decreasing relative to the previous month.

The estimates of $U = 0.140$ means that the distance between the frictionless price and the existing target price, $p^*_{it} - p^{\#}_{it-1}$, has to be 14% before a price-adjustment process is initiated. Then, the actual price changes are decided by the partial adjustment model. The value of $\theta_{up} = 0.370$ is interpreted as meaning that the producer will close 63% (= 1−0.370) of the desired price change, $p^{\#}_{it} - p_{it-1}(= p^*_{it} - p_{it-1})$, in the same period as the firm decides to reset the target price and start a price increase process. The parameter estimates of $U = 0.140$ and $\theta_{up} = 0.370$ together state that the initial price increase will be at least 0.088 of the current price.[19] Thus, the new prices are reached quite quickly.[20] The estimate of the lower threshold $L = -0.170$ means that the (absolute) distance between the frictionless price and the existing target price, $\left|p^*_{it} - p^{\#}_{it-1}\right|$, has to be almost 17% before a price decrease process is initiated. The

[19]The initial price increase is found by multiplying U with $(1 - \theta_{up})$: 0.140×(1−0.370)≈ 0.088.

[20]The negative autocorrelation moments imply that an above (below) average positive (negative) price adjustment is likely to come subsequent to a below (above) average price adjustment.

producer will again quite quickly adjust to the new smaller price, seen from the θ_{down} $= 0.409$.[21]

The values of $\theta_{up} = 0.370$ and $\theta_{down} = 0.409$ mean that the producers find it more important to close the gap relative to the new target price, $\left(p_{it} - p_{it}^{\#}\right)$, than reducing the implied convex adjustment costs associated with the period-to-period adjustment, $(p_{it} - p_{it-1})$.

The findings suggest that adjustments are faced with two different forms of friction. First, the effect of the threshold is that it must be desired to change the price by at least the size of the threshold before the firm decides to adjust. Second, the effect of the inertia is that the initial price change will be equal to $(1 - \theta)$ of the target price gap $p_{it}^{*} - p_{it-1}$, whereas subsequent adjustments will be smaller. The January effect, meant to capture the fact that the incidence of price changes is higher in January compared with the other months, means that the thresholds U and L are two-thirds in January compared with the other months. This is consistent with theories focusing on the costs of information acquisition and processing, i.e pricing season effects, and seasonal demands.

The model performs relatively well, as seen from the J-statistic in the last row.[22] This reflects that the empirical moments, reported in the lower part of the table, are matched quite well with the empirical ones reported in the last column of the table, Column (8)

In Column (1) are the magnitude of the parameter estimates for the pairs U and L, and θ_{up} and θ_{down}, which are very similar, even though their significance clearly

[21] The initial price decrease will be at least 0.099 of the current price, and there will subsequently be several smaller adjustments downwards until the firm reaches the target price or decides to set a new one. Thus, one may think that price changes of 8.5% and 10.0% (initial price changes for positive and negative price adjustments, respectively) are definitely not ignorable.

[22] While the J-statistics reported in Table 1 are low compared to related studies, the numbers imply that all specifications are rejected. This is not surprising given that the moments are very precise (see our Table A1) and, consequently, the weighting matrix has very large values. Furthermore, we are admittedly conflating the parameter estimates across all different types of products, which also affects the J-statistic.

states that they are statistically different ($|L| \neq U$, and $\theta_{up} \neq \theta_{down}$). Still, when forcing the parameters within each pair to be the same (in absolute values), reported in Column (2), the J-statistic indicates a worse model fit.[23]

We have also tested a frictionless model, reported in Column (3). What we see is a very bad model fit, both measured by the J-statistic, and by comparing the individual moments of the simulated model with the empirical ones. Thus, a model without any price-adjustment frictions is inferior compared with the two former models. The frictionless model in Column (3) states that the estimated standard error $\widehat{\sigma}_{\varepsilon}^{frictionless} = 0.0012$ and that the frequency of zero price changes is 0.981. In our model, we have $\triangle p^*_{it} = p^*_{it} - p^*_{it-1} = 0.0025 + \varepsilon_{it}$ where $\varepsilon_{it} \sim N(0, \sigma^2_{\varepsilon})$ (see eq. (3)). Note also that a frictionless model has $p_{it} = p^*_{it}$, i.e., the new price is equal to the frictionless price. Simple calculation shows that $\sigma_{\varepsilon} = 0.0029$ is necessary in a frictionless model to get 80% of the observations of $\triangle p^*_{it}$ (and therefore also $\triangle p_{it}$) within the interval [-0.005, 0.005], the "zero price changes"-interval moment. When our frictionless model gives the best fit with a frequency of zero price changes of 0.981, it is due to the other moments which force the standard error $\widehat{\sigma}_{\varepsilon}^{frictionless} = 0.0012$ for the best possible weighted match with all of the empirical moments.[24] We have also tested specifications where only some of the friction parameters are present. In Column (4), we include only the threshold parameters U and L, whereas in Column (5) we have estimated a pure partial adjustment model. The results of the latter are identical to those reported in Column (3).[25] The overall finding, based on the results reported in

[23] Clearly, the difference in J-statistic = 59.7 (=235.5-175.8), $df = 2$, indicates that this restriction largely distorts the performance of the model.

[24] The J-statistic is 2546.3 for the estimated frictionless model, whereas it is $J = 16805.0$ when we set $\sigma_{\varepsilon} = 0.0029$. It turns out that when we force the magnitude of σ_{ε} going from the $\widehat{\sigma}_{\varepsilon}^{frictionless} = 0.0012$ up to $\sigma_{\varepsilon} = 0.0029$, we get too many small positive price adjustments $\langle 0.5\%, 2.5\%]$ very quickly. This indicates that the moments describing the distribution are good for identification of the variance of the underlying process of the frictionless nominal prices.

[25] The similarity of the two sets of results is as expected. The two serial correlation moments, meant to capture inertia, are based on the autocovariance of $\triangle p^*_{it} = p^*_{it} - p^*_{it-1}$. When the process of p^*_{it} is a random walk with drift, the two autocovariances—upwards and downwards—should both be zero. Thus, when $U = L = 0$, there is nothing in the model nor moments that help distinguish the frictionless model from the partial adjustment model. The results should therefore be the same

Columns (1)–(5), is that a full model with both asymmetric price thresholds and an asymmetric partial adjustment process, is the preferred specification.

As an additional way to show the model fit, in addition to the J- statistics and the individual moments compared with their empirical counterpart, we report the whole distribution of price changes, the share of price changes in a given month, and the mean price change for all ranks for the simulated data. Going back to Figures 1–3, we show the simulated moments for the full model (grey bars) reported in Table 1, Column (1). Starting with the predictions of average share of price changes, Figure 3 shows that the seasonal pattern is reasonably good. For the distribution of price changes and the ranks (Figures 1 and 2, respectively) in particular, the fit is also very good for the moments not used when estimating the model. We interpret this as supporting evidence for our model formulation.

4.1 Robustness Checks and Discussion

In our underlying frictionless price, the predetermined trend parameter $\alpha = 0.0025$, which corresponds to an annual price increase of 3%. Two alternative simulations are done where α is set equal to 0.0016 and 0.0035. These are denoted as α^{low} and α^{high} and correspond to underlying annual price increases of 2% and 4%. The results of these two alternative trends, and using the full specification, are reported in Table 1, Columns (6) and (7). The model fit, measured by the J-statistic, is worse when using α^{low}. U and L are both somewhat smaller compared with the full model when $\alpha = 0.0025$. The largest, but moderate, difference is for the θ_{up} (0.370 and 0.298, in Columns (1) and (6), respectively). Turning to the α^{high} results, we see that both the J-statistic and the estimated coefficients are very close to the initial model reported in Column (1). The conclusion we draw from this robustness check, is that the initial guess of $\alpha = 0.0025$ is not too wrong and, if any changes should be made, the α

for the two model specifications, as we obtain.

should be set somewhat greater rather than smaller compared with 0.0025.[26] We have considered estimating α, instead of predetermining it, as in this paper. The small deviation between the estimates based on different values of α tells us that the gain from doing this is limited.

So far, we have assumed that the friction parameters U, L, θ_{up} and θ_{down} are all independent of product characteristics. There is a broad agreement in the literature that price setting is heterogeneous across sectors, firms, and products (see, for instance, the overview by Klenow and Malin (2011)). Two exercises are performed to address heterogeneity. First, for each product, the *share of price-change observations* (outside the zero-price-change interval [-0.005, 0.005]) relative to the total number of observations for that product is computed. The distribution (based on all products) of this *share of price-change observations* per product is right-skewed and has the following properties: mean = 0.207, median = 0.087, standard deviation = 0.258. The median is very close to 0.083 (= 1/12) and indicates that almost half of the product prices are changed only once a year. The right-skewness implies that some of the products have a much higher frequency of price changes (see Figure A1 for a density plot of the *share of price-change observations* outside the zero-price-change interval). Thus, there is heterogeneity in our data.[27]

A further exercise to address heterogeneity involves estimating the model for five

[26] A regression model where the dependent variable is log-transformed product prices and where a time trend is used, together with product-specific dummies, month-specific dummies, and year-specific dummies, gives a time trend $\alpha = 0.0029$, which corresponds to a 3.5% annual increase.

[27] We have done a simulation where we randomly assign a product to be either high- or low-spread shock type, with equal probability of being assigned to one of the two types. Having included the st.dev. of *share of price-change observations* as an additional moment, we find the st.dev. of the idiosyncratic shocks for the low-spread shock type is only 0.433 (st.error = 0.033) out of the st.dev. of the high-spread shock type ($\sigma_\varepsilon^{low-spread} = 0.433\sigma_\varepsilon^{high-spread}$). As a consequence, about 80 percent (20 percent) of price changes occur in high-spread shock type products (low-spread shock type products). More importantly, the parameters and simulated moments are of the same order of magnitude as for our benchmark model. Thus, one potential source of heterogeneity might be related to business environment and demand fluctuations. When randomly assigning the threshold parameters U and L either high or low values to each product (to create a wide or narrow zero price-change interval), we have problems in getting convergence. The conclusion we draw from the investigation of heterogeneity, both based on our own simulations and seen in the literature, is that which model fits best the “stylized facts” - and in particular heterogeneity - remains an open issue.

different product groups. These results are reported in Table 2.

[Table 2: Empirical Results by Product Groups - about here]

We see a very good model fit, seen from the low J-statistics, much better than the J-statistics found when estimating the model with the sample of all product types (Table 1, Column 1). This shows that there are differences in the pricing patterns among products, both when it comes to the threshold parameters U and L, the inertia parameters θ_{up} and θ_{down}, and the January-specific parameter y. In particular, "Capital goods" stand out with much greater thresholds compared with our previous results and a significantly different January effect. The overall pattern shown in Table 2 implies that there are huge differences in price-change patterns across product groups.

[Table 3: Counterfactual Analyses Results - about here]

Even though heterogeneity is important, we are also interested in obtaining "average" parameters for macro simulations. Thus, we go back to the model where heterogeneity is ignored to analyse the importance of the respective price-adjustment parameters in explaining the main characteristics of observed price changes. To shed some light on this, we simulate the preferred asymmetric model under exactly the same circumstances as the estimated preferred model (Table 1, Column 1), but set different price-adjustment cost parameters to zero, and measure the impact on the set of moments used for identification. Table 3 shows the results of this exercise. The first thing to notice is the huge increases in the J-statistics, meaning that overall model fits are much worse when some of the friction parameters are ignored. A more detailed look, starting with Column (1) where the two θs are set equal to zero, we see an increased share of zero price changes, no small price changes, and also an increase in the observations outside the [-5%, 5%] interval. The two inertia parameters, θs, are, therefore, very important for creating small price changes, $0.5\% < \left|\frac{\triangle p}{p}\right| \leq 5.0\%$.

Omitting them from a model, leads to the conclusion that the dynamics of price changes are best described as zeroes and lumps driven by a fixed- costs model. Turning to Column (2) where the thresholds U and L are ignored, the model fit is extremely bad, especially in producing a large enough share of inaction. Thus, both sets of friction parameters, inertia and thresholds, are important for understanding the dynamics of price changes. [28]

As already mentioned, each plant may produce one or several products.[29] In our model, we have treated each product as independent from each other. This means that we assume that the products within a plant's product portfolio are sufficiently differentiated and we, therefore, abstract from strategic complementarity and substitution between the various products.[30] We have analysed the various components of the variance of price changes in the dataset by using a multi-level mixed-effects linear regression model (see Baltagi et al., 2001). The analysis shows that the plant-specific share of the overall price variation is approximately 40%, and that 60% of the variation is related to product-specific shocks. One modification to control for this effect would be to introduce two variance components in our idiosyncratic shocks σ_ε^2, such that $\sigma_\varepsilon^2 = \sigma_p^2 + \sigma_u^2$, where σ_p^2 denotes variance of plant-specific shocks and σ_u^2 denotes idiosyncratic product shocks. One could also consider modelling frictions at the level of the plant rather than at the level of products. Our simplification to treat each product independently is likely to affect the magnitude of the threshold parameters, because they will pick up the effect of both price-change costs and, to some degree, shocks that are common to all products of a producer. Nevertheless, we leave these

[28]We have also estimated a model where we look at annual data—i.e., the price changes from June one year to June the subsequent year— and estimated with a different moment vector. Unsurprisingly, these results indicate much smaller thresholds, asymmetric but still statistically significant. The downward inertia parameter θ_{down} is significant, whereas $\theta_{up} = 0$. Thus, time aggregation blurs the price-changing picture compared with using a model that is able to take advantage of the monthly frequencies.

[29]The mean number of products per producer is 3.8, whereas the maximum is 20.

[30]See also Woodford (2003) and Gertler and Leahy (2008) for discussions about strategic complementarity and "real rigidity" but then, in relation to firms' competitors.

potential extensions for a later paper.

Our estimates show that price-adjustment frictions are important for understanding the intermittent price-adjustment pattern seen in the data, and that ignoring frictions biases our results. Furthermore, our findings, of both thresholds and inertia together, indicate that different forms of rigidities exist, which are only partly consistent with the assumptions of most existing macroeconomic pricing models. For instance, models such as those in Golosov and Lucas (2007) and Gertler and Leahy (2008) explain patterns of inaction, followed by large price changes, by assuming thresholds; however, these models seem to neglect small price adjustments. While the threshold parameters in our model enable inaction, the inertia parameters implicate that a large initial price change is followed by smaller adjustments. Accordingly, the results imply that our model is able to account for periods of inaction, as well as both large and small price changes. The existing literature, discussed in the earlier sections, is not conclusive when it comes to whether nominal price rigidities are symmetric or not. Our findings indicate that prices are almost similarly flexible upwards and downwards. The seasonal effect, picked up by our January parameter, y, may have implications for the effectiveness of monetary policy interventions depending on the month of the year in which the intervention takes place (for related findings, see Olivei and Tenreyro (2007)).

5 Concluding Remarks

In this paper, we specify and estimate a model that describes production plants' price-adjustment behaviour. The model includes thresholds that are smaller in January than in the other months, together with a quadratic loss function associated with the distance from a target price. The simplistic reduced form model is tested on a sample based on repeated monthly plant- and product-specific survey data from the Norwegian manufacturing industry. The model is meant to reproduce the following features of the data. First, plants adjust prices infrequently as only 20% of the price observations change from one month to another. Second, there is a seasonal pattern in the incidence of price changes, with most price changes taking place in January. Third, there are also many small price changes.

The simulated method-of-moment estimates reveal thresholds such that prices are only changed if the deviation from the target price to the underlying frictionless price is greater than approximately 15%. However, if the shocks are such that the prices should be changed, the gap between the current price and the new target price is reduced quite quickly and only 10% of the initial gap exists after three months. There are statistically significant cost differences whether the prices move upwards or downwards. However, the magnitudes of these differences are very moderate. Furthermore, the January-specific effect, indicating that the thresholds are only two-thirds for this month, is consistent with theories focusing on the costs of information acquisition and processing, and seasonality in the signing of price contracts.

Several checks are applied to test the robustness of the model and our findings. First, the preferred specification outperforms a frictionless model or models with only some parts of the price-adjustment friction parameters present. A counterfactual analysis, where some of the friction parameters are set equal to zero, shows that the moment fit becomes much worse compared with the preferred model specification. Furthermore, the model seems to be fairly robust to changes in the underlying deterministic trend, as our approximation of the trend gives a better a fit than alternative approximations.

While our evidence implies both large and small price changes, many model contributions in the literature are only able to account for one of these two characteristics. However, a few of the models that assume thresholds in the price setting are able to explain small price changes. These models assume either stochastic thresholds or economies of scope in price setting, and represent an increasingly sophisticated group of pricing models in which more micro evidence is incorporated. However, our model is rather simple and transparent, and computationally easy.

There are a few issues we have not addressed that need to be explored in future work. The model is admittedly a reduced-form model. A structural model would be more informative. This would call for a full dynamic specification and optimization. However, with the current dataset, there is no information about quantities, even though annual revenues and costs are available at plant level. Furthermore, the only information available at product level with monthly frequency are prices themselves. Thus, a structural model would partly require unverifiable assumptions about inputs and outputs. Still, our findings strongly indicate that such a model needs to include both convex and non-convex price-adjustment costs. The analyses also point in the direction for taking into account and controlling for product- (and plant-) specific heterogeneity. Note however, the mixed frequency of price information, and other plant- or firm-specific information (monthly versus annually), present some econometric challenges. However, the evidence provided in this paper, based on a simple and transparent simulation model, shows the importance and potential fruitfulness of using model formulations and estimation techniques that can take into account the non-convexities in the price-adjustment costs function.

References

Alvarez, F. E., Lippi, F., and Paciello, L. (2011). Optimal price setting with observation and menu costs. *The Quarterly Journal of Economics*, 126(4):1909–1960.

Baltagi, B. H., Song, S. H., and Jung, B. C. (2001). The unbalanced nested error component regression model. *Journal of Econometrics*, 101:357–381.

Bhattarai, S. and Schoenle, R. (2014). Multiproduct firms and price-setting: Theory and evidence from U.S. producer prices. *Journal of Monetary Economics*, 66:178–192.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Caballero, R. J. and Engel, E. M. (1993). Microeconomic rigidities and aggregate price dynamics. *European Economic Review*, 37(4):697–711.

Dhyne, E., Fuss, C., Pesaran, M. H., and Sevestre, P. (2011). Lumpy price adjustments: A microeconometric analysis. *Journal of Business & Economic Statistics*, 29(4):529–540.

Dias, D. A., Marques, C. R., Martins, F., and Santos Silva, J. M. C. (2015). Understanding price stickiness: Firm-level evidence on price adjustment lags and their asymmetries. *Oxford Bulletin of Economics and Statistics*, 77(5):701–718.

Dobrynskaya, V. V. (2008). Asymmetric price rigidity and the optimal interest rate defense of the exchange rate: Some evidence for the US. *Journal of Policy Modeling*, 30(5):713–724.

Dolado, J. J., Maria-Dolores, R., and Naveira, M. (2005). Are monetary-policy reaction functions asymmetric? the role of nonlinearity in the phillips curve. *European Economic Review*, 49(2):485–503.

Doms, M. and Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, 1(2):409–429.

Eichenbaum, M., Jaimovich, N., Rebelo, S., and Smith, J. (2014). How frequent are small price changes? *American Economic Journal: Macroeconomics*, 6(2):137–155.

Gertler, M. and Leahy, J. (2008). A phillips curve with an ss foundation. *Journal of Political Economy*, 116(3):533–572.

Golosov, M. and Lucas, R. E. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, 115:171–199.

Honoré, B. E., Kaufmann, D., and Lein, S. (2012). Asymmetries in price-setting behavior: New microeconometric evidence from Switzerland. *Journal of Money, Credit and Banking*, 44(2):211–236.

Klenow, P. and Malin, B. (2011). Microeconomic evidence on price-setting. *Handbook of Monetary Economics*, 3:231–284.

Klenow, P. J. and Kryvtsov, O. (2008). State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? *Quarterly Journal of Economics*, 123(3):863–903.

Laxton, D., Rose, D., and Tambakis, D. (1999). The U.S. Phillips curve: The case for asymmetry. *Journal of Economic Dynamics and Control*, 23(9-10):1459–1485.

Lee, B. S. and Ingram, B. F. (1991). Simulation estimation of time-series models. *Journal of Econometrics*, 47(2-3):197–205.

Lewis, M. S. (2011). Asymmetric price adjustment and consumer search: An examination of the retail gasoline market. *Journal of Economics and Management Strategy*, 20(2):409–449.

Loupias, C. and Sevestre, P. (2012). Costs, demand, and producer price changes. *Review of Economics and Statistics*, 95(1):315–327.

Loy, J. P., Weiss, C. R., and Glauben, T. (2016). Asymmetric cost pass-through? Empirical evidence on the role of market power, search and menu costs. *Journal of Economic Behavior and Organization*, 123:184–192.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *The American Economic Review*, 99(3):769–803.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, 117(4):1295–1328.

Mankiw, N. G. and Reis, R. (2010). Imperfect information and aggregate supply. *Handbook of Monetary Economics*, 3A:183–229.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026.

Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, 79(4):1139–1180.

Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*, 123(4):1415–1464.

Nilsen, Ø. A., Bratlie, J., and Pettersen, P. M. (2016). Time-dependency in producers price adjustments: Evidence from micro panel data. mimeo, Norwegian School of Economics.

Nilsen, Ø. A. and Schiantarelli, F. (2003). Zeros and lumps in investment: Empirical evidence on irreversibilities and nonconvexities. *Review of Economics and Statistics*, 85(4):1021–1037.

Olivei, G. and Tenreyro, S. (2007). The timing of monetary policy shocks. *American Economic Review*, 97:636–663.

Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057.

Peltzman, S. (2000). Prices rise faster than they fall. *Journal of Political Economy*, 108(3):466.

Ratfai, A. (2006). Linking individual and aggregate price changes. *Journal of Money, Credit and Banking*, 38(8):2199–2224.

Rotemberg, J. (1982). Sticky prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.

Sheshinski, E. and Weiss, Y. (1977). Inflation and costs of price adjustment. *The Review of Economic Studies*, 44(2):287–303.

SSB (2015). Commodity price index for the industrial sector. http://www.ssb.no/en/vppi.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *The Journal of Political Economy*, pages 1–23.

Taylor, J. B. (1999). Staggered price and wage setting in macroeconomics. *Handbook of macroeconomics*, 1:1009–1050.

Varejão, J. and Portugal, P. (2007). Employment dynamics and the structure of labor adjustment costs. *Journal of Labor Economics*, 25(1):137–165.

Vermeulen, P., Dias, D. A., Dossche, M., Gautier, E., Hernando, I., Sabbatini, R., and Stahl, H. (2012). Price setting in the Euro area: Some stylized facts from individual producer price data. *Journal of Money, Credit and Banking*, 44(8):1631–1650.

Woodford, M. (2003). Optimal interest-rate smoothing. *Review of Economic Studies*, 70:861–886.

Wulfsberg, F. (2016). Inflation and price adjustments: Micro evidence from norwegian consumer prices 1975-2004. *American Economic Journal: Macroeconomics*, 8(3):175–194.

Yang, H. and Ye, L. (2008). Search with learning: Understanding asymmetric price adjustments. *RAND Journal of Economics*, 39(2):547–564.

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(October 2015):514–533.

Figures

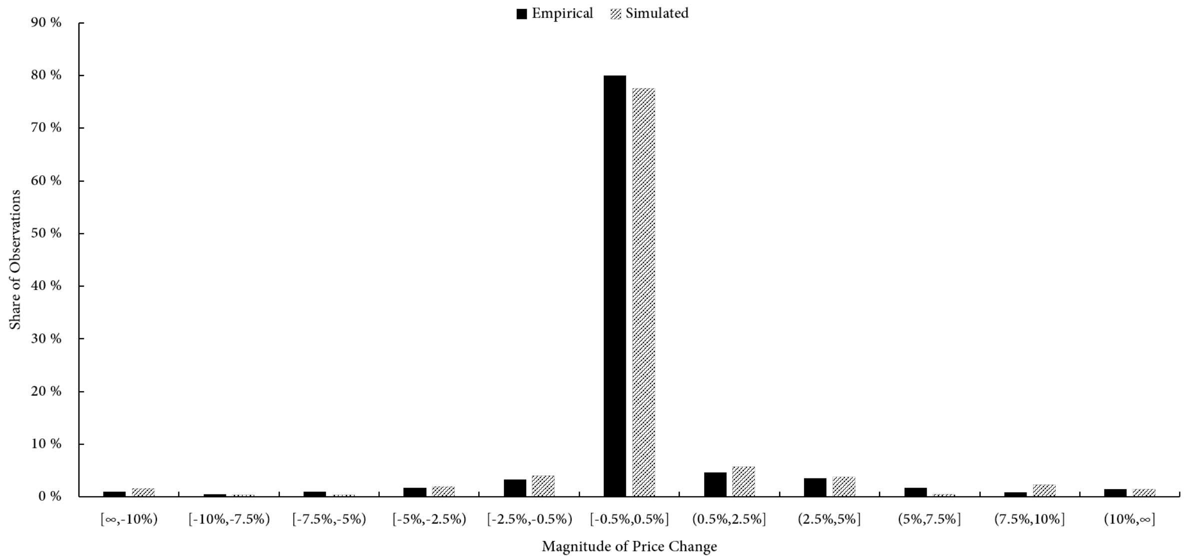

Note: *The figure shows the proportion of observations in different price change intervals. The price changes in this figure are calculated using the following logarithmic approximation*: $\ln(p_{it}) - \ln(p_{it-1}) \approx \frac{p_{it}-p_{it-1}}{p_{it-1}}$, *where* p_{it} *denotes price. Because this is a differenced variable, we loose one observation for every product.*

Figure 1: Distribution of Price Change Rates, Empirical and Simulated

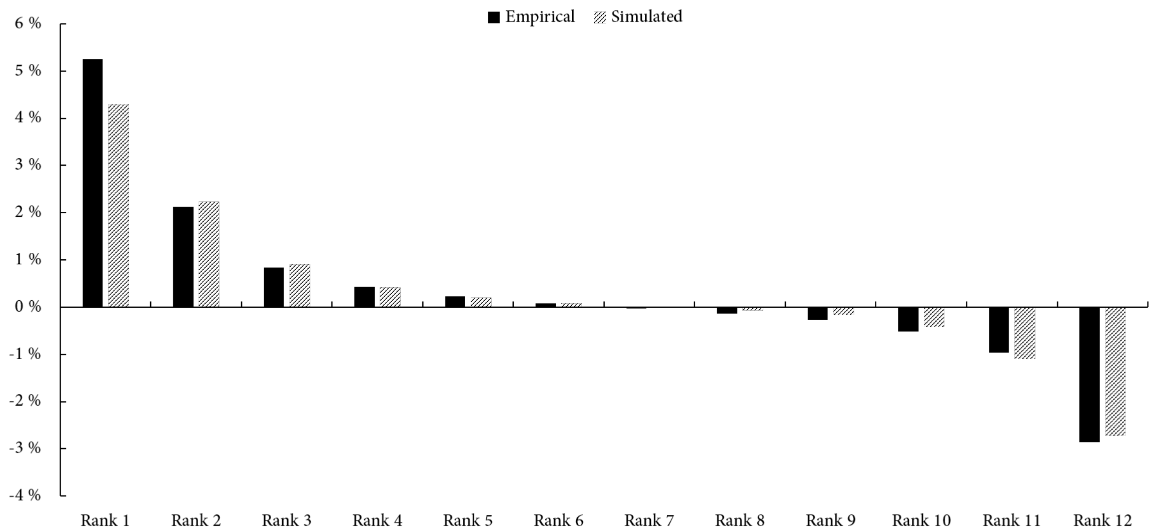

Note: *Ranks are first calculated on a yearly basis for every product before they are averaged across products and years. Rank 1 represents the average highest monthly price change and rank 12 represents the average lowest monthly price change.*

Figure 2: Ranked Price Change Rates, Empirical and Simulated

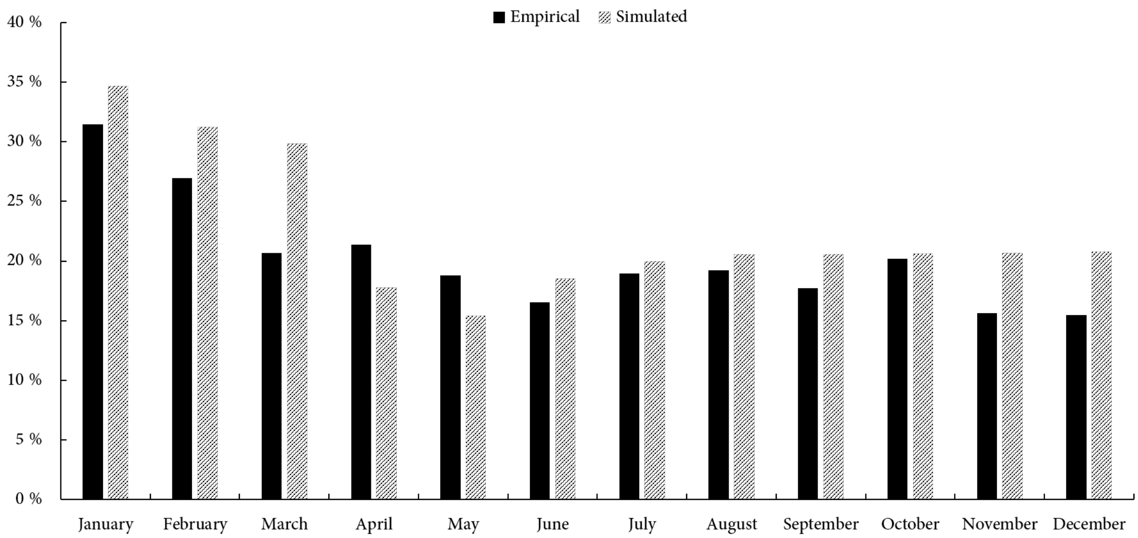

Note: *The figure shows the average frequency of producer price changes in the years 2004 to 2009 by calendar month. Shares are given as the number of price changes larger than 0.5% within each month divided by the total number of price quotations in the month.*

Figure 3: The Occurrence of Price Changes by Months, Empirical and Simulated

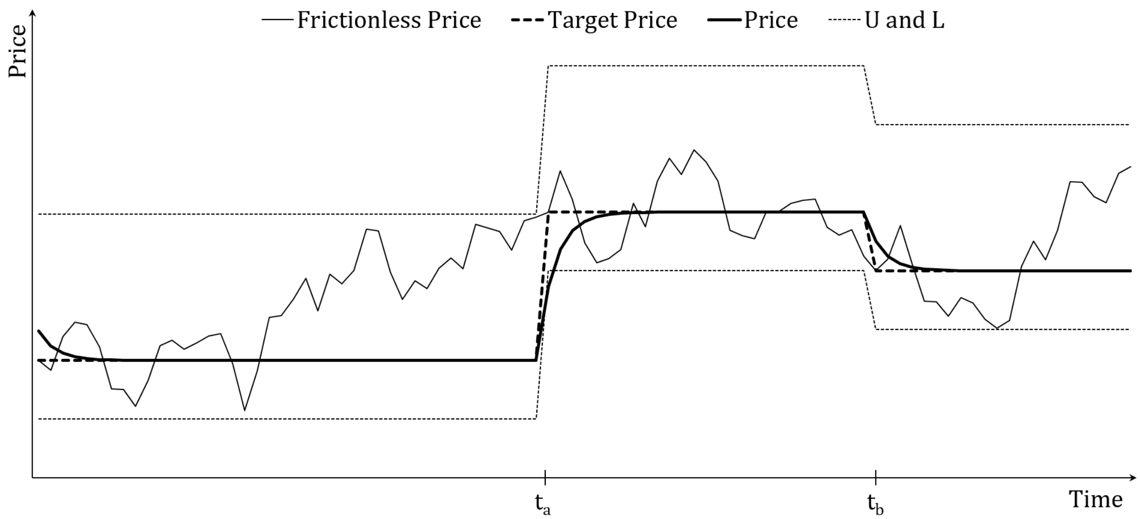

Note: The figure illustrates the dynamics of the full model with thresholds and inertia. In order to simplify the exposition, a manipulated price series is used.

Figure 4: Illustration of the Price Changing Process

Tables

Specifications:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
U	0.140	0.144	-	0.044	-	0.118	0.128	-
	(0.003)	(0.003)	-	(0.001)	-	(0.002)	(0.003)	-
L	-0.167	-0.144	-	-0.129	-	-0.180	-0.162	-
	(0.004)	(0.003)	-	(0.025)	-	(0.003)	(0.004)	-
θ_{up}	0.370	0.379	-	-	0.000	0.298	0.358	-
	(0.003)	(0.003)	-	-	(0.000)	(0.007)	(0.005)	-
θ_{down}	0.409	0.379	-	-	0.000	0.372	0.386	-
	(0.010)	(0.003)	-	-	(0.000)	(0.007)	(0.010)	-
y	0.712	0.732	-	0.6346	-	0.681	0.693	-
	(0.017)	(0.016)	-	(0.0221)	-	(0.015)	(0.017)	-
σ_ε	0.042	0.040	0.001	0.016	0.001	0.041	0.041	-
	(0.0007)	(0.001)	(1.98E-05)	(0.001)	(1.98E-05)	(0.001)	(0.001)	-
Moments								
$[-5.0\%, -2.5\%\rangle$	0.019	0.023	0.000	0.000	0.000	0.019	0.019	0.017
$[-2.5\%, -0.5\%\rangle$	0.041	0.039	0.000	0.000	0.000	0.038	0.037	0.033
$[-0.5\%, 0.5\%]$	0.779	0.775	0.981	0.939	0.981	0.790	0.774	0.800
$\langle 0.5\%, 2.5\%]$	0.059	0.065	0.019	0.000	0.019	0.050	0.058	0.047
$\langle 2.5\%, 5.0\%]$	0.038	0.037	0.000	0.028	0.000	0.035	0.043	0.036
Chgs in Jan	0.333	0.332	0.019	0.200	0.019	0.335	0.348	0.315
Serial corr (up)	-0.364	-0.309	0.011	-0.059	0.011	-0.488	-0.445	-0.136
Serial corr (down)	-0.264	-0.321	0.000	0.000	0.000	-0.292	-0.291	-0.385
$sd(p_{it} - p_{it-1})$	0.028	0.027	0.001	0.014	0.001	0.028	0.028	0.041
Rank 12	-0.026	-0.026	0.001	-0.005	0.001	-0.030	-0.026	-0.029
Rank 11	-0.011	-0.010	0.001	0.000	0.001	-0.011	-0.010	-0.010
Rank 2	0.022	0.021	0.004	0.010	0.004	0.020	0.024	0.021
Rank 1	0.043	0.042	0.005	0.027	0.005	0.042	0.046	0.053
J:	175.8	235.5	2546.3	819.2	2546.3	347.7	188.8	-

The column numbers represent the following specifications: (1): Full model, (2): Full model with symmetric friction parameters, (3): Frictionless, (4): No inertia parameters, (5): No threshold parameters, (6): Full model with $\alpha = 0.0017$, (7): Full model with $\alpha = 0.0033$, (8): Empirical moments.

Table 1: Parameter Estimates and Moments

	(1)	(2)	(3)	(4)	(5)
Specifications:	Capital goods	Durables	Intermediate goods	Non-durables, food	Non-durables, non-food
U	0.282	0.119	0.156	0.153	0.170
	(0.061)	(0.006)	(0.004)	(0.004)	(0.015)
L	-0.265	-0.240	-0.212	-0.162	-0.206
	(0.070)	(0.031)	(0.008)	(0.006)	(0.161)
θ_{up}	0.305	0.299	0.356	0.407	0.026
	(0.023)	(0.025)	(0.004)	(0.018)	(0.002)
θ_{down}	0.252	0.443	0.478	0.376	0.003
	(0.034)	(0.020)	(0.014)	(0.017)	(0.171)
y	0.418	0.820	0.712	0.883	0.616
	(0.086)	(0.087)	(0.026)	(0.039)	(0.065)
σ_e	0.030	0.041	0.051	0.048	0.019
	(0.002)	(0.004)	(0.002)	(0.002)	(0.004)
Moments					
$[-5.0\%, -2.5\%\rangle$	0.005	0.011	0.022	0.024	0.000
$[-2.5\%, -0.5\%\rangle$	0.011	0.023	0.043	0.052	0.000
$[-0.5\%, 0.5\%]$	0.913	0.826	0.750	0.731	0.981
$\langle 0.5\%, 2.5\%]$	0.024	0.041	0.062	0.079	0.001
$\langle 2.5\%, 5.0\%]$	0.017	0.036	0.038	0.036	0.000
Chgs in Jan	0.340	0.230	0.354	0.307	0.178
Serial corr(up)	-0.358	-0.503	-0.340	-0.259	0.000
Serial corr(down)	-0.401	-0.119	-0.106	-0.277	0.000
$sd(p_{it} - p_{it-1})$	0.022	0.027	0.033	0.032	0.020
Rank 12	-0.011	-0.021	-0.031	-0.036	-0.003
Rank 11	-0.003	-0.009	-0.015	-0.014	0.000
Rank 2	0.010	0.019	0.026	0.026	0.003
Rank 1	0.033	0.040	0.050	0.048	0.032
N	243	134	807	359	133
J	93.7	75.9	85.7	40.5	75.7

Note: standard errors in parentheses, N denotes number of products, J denotes the criterion value.

Table 2: Empirical Results by Product Groups

	(1)	(2)	(3)
Specifications:	No inertia	No thresholds	No friction parameters
U	0.140	-	-
	(0.003)	-	-
L	-0.167	-	-
	(0.006)	-	-
θ_{up}	-	0.370	-
	-	(0.020)	-
θ_{down}	-	0.409	-
	-	(0.019)	-
y	0.712	-	-
	(0.014)	-	-
σ_{ε}	0.042	0.042	0.042
	(0.001)	(0.001)	(0.0001)
$[-5\%, -2.5\%)$	0.000	0.133	0.151
$[-2.5\%, -0.5\%)$	0.000	0.241	0.173
$[-0.5\%, 0.5\%]$	0.940	0.143	0.096
$(0.5\%, 2.5\%]$	0.000	0.248	0.181
$(2.5\%, 5.0\%]$	0.000	0.163	0.168
Chgs in Jan	0.189	0.857	0.905
Serial corr (up)	-0.019	0.233	0.001
Serial corr (down)	0.000	0.234	0.001
$sd(p_{it} - p_{it-1})$	0.041	0.028	0.042
Rank 12	-0.042	-0.039	-0.065
Rank 11	-0.003	-0.026	-0.044
Rank 2	0.018	0.034	0.052
Rank 1	0.066	0.046	0.071
J	3 213.8	28 850.6	27 037.6

Table 3: Counterfactual Analysis

A Appendices

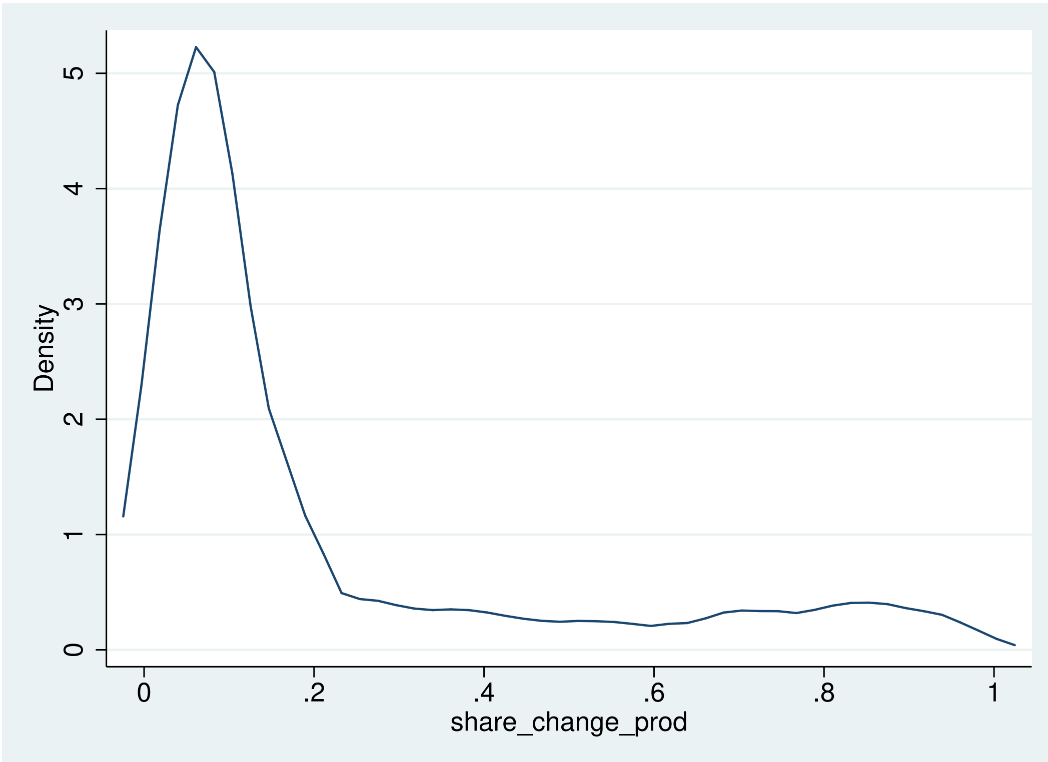

Note: For each product the share of price-change observations outside the zero-price-change-interval [-0.005, 0.005] relative to the total number of observations per product is computed. The distribution of this share of price-change observations per product has the following properties; mean = 0.207, median = 0.087, st.dev. = 0.258.

Figure A1: The share of price-change observations outside the zero-price-change interval per product

	Mean	Std. dev.	Z-value	95% Conf. Interval	
$[-5\%, -2.5\%)$	0.017	0.001	19.000	0.015	0.019
$[-2.5\%, -0.5\%)$	0.033	0.002	19.040	0.030	0.037
$[-0.5\%, 0.5\%]$	0.800	0.006	128.670	0.788	0.812
$(0.5\%, 2.5\%]$	0.047	0.002	21.810	0.043	0.051
$(2.5\%, 5.0\%]$	0.036	0.001	36.040	0.034	0.038
Chgs in Jan	0.315	0.009	33.370	0.296	0.333
Serial corr (up)	-0.136	0.028	-4.770	-0.192	-0.080
Serial corr (down)	-0.385	0.134	-2.870	-0.648	-0.122
$sd(p_{it} - p_{it-1})$	0.0413	0.003	15.020	0.036	0.047
Rank 12	-0.029	0.001	-22.910	-0.031	-0.026
Rank 11	-0.010	0.001	-16.720	-0.011	-0.009
Rank 2	0.021	0.001	28.690	0.020	0.023
Rank 1	0.053	0.001	38.020	0.050	0.055

Note: The first five rows represent the total shares of observations within the given intervals and the following row represents the share of price changes in January.

Table A1: Bootstrapped Moments with Std. Errors

A.1 A Simple Model for Frictionless Price

Assume a Cobb–Douglas production technology with a flexible input factor, K. The costs of this factor are exogenous to the plant and denoted r. Assume also that the plants have some market power and demand is given by an isoelastic function. Also, assume that goods are sufficiently differentiated to abstract from substitution within a multi-product firm's portfolio of products. For notational convenience, we abstract from subindices for the plant, product, and time. Then, production is determined by $Q^S(K) = A \cdot K^a$ where $0 < a < 1$ and the isoelastic demand function is given by $Q^D(P) = B \cdot \left(P/P^C\right)^{-\epsilon}$ where $\epsilon > 1$. The price of a plant's product is given by P, and P^C denotes the general price level in the industry. The price level P^C is exogenous to the plant which implies that we employ a partial equilibrium model. Abstracting from inventory, profit for a single product is then given by

$$\pi(A, B, P^C, r) = P \cdot B \cdot \left(P/P^C\right)^{-\epsilon} - r \cdot \left(\frac{B}{A}\right)^{1/a} \cdot \left(P/P^C\right)^{-\epsilon/a},$$

where A captures supply shocks, and B captures demand shocks. With these assumptions, the first-order derivative of profit $\pi(.)$ with respect to price P can be expressed as follows:

$$P^* = \underbrace{\left[\frac{\epsilon}{a(\epsilon-1)} \; B^{\frac{1-a}{a}} A^{-\frac{1}{a}}\right]^{\frac{a}{\epsilon(1-\gamma)}}}_{(1)} \times \underbrace{r^{\frac{a}{\epsilon(1-\gamma)}}}_{(2)} \times \underbrace{(P^C)^{\frac{1-a}{\epsilon(1-\gamma)}}}_{(3)}, \quad where\ \gamma = a(1-\frac{1}{\epsilon})$$

This expression is a nonlinear function of the state of supply A, the state of demand B, the input costs r, and the general price in the industry P^C. Given that $a < 1$, we see that a positive supply shock, $A\uparrow$, will implicate a smaller price, as expression (1) will get a smaller value. This could be, for example, because the producer obtains better technology that increases productivity. We also see that a positive demand shock, $B\uparrow$, will implicate a higher price, because the net effect on expression (1) will be positive. Furthermore, if producers are faced with a positive cost shock, $r\uparrow$, the frictionless equilibrium price will increase, as expression (2) will be more positive. Note however that the degree of pass-through, the share of the cost increase that will be borne by costumers, depends on the parameter values. Higher competitors' prices, $P^C\uparrow$, also induce a price increase. In our model setup, presented in the main text, part of this latter effect—that the general price level increases—is picked up by the trend parameter α. The remaining supply, demand, and cost shocks, together with the nondeterministic part of a competitor's prices, will all be picked up by the idiosyncratic shocks, ε_{it}, in the model presented in the main text. Finally, we see that the marginal effects of the various shocks affect the price differently.

A.2 Simulated Method of Moments

In the simulated method of moments (SMM) approach, κ simulated datasets are generated for N panels and $96 + T$ time periods. N and T are set equal to the number of panels and time periods in the empirical data.[31] To limit the impact of initial conditions, the first 96 time periods (eight years) are discarded when calculating the simulated moments, leaving only T time periods.[32]

If we let the vector of l unknown parameters be denoted by the vector β, the optimal vector of unknown parameters, $\hat{\beta}$, is given by:

$$\hat{\beta} = \underset{\beta}{argmin}\,[\Phi^A - \frac{1}{\kappa}\sum_{j=1}^{\kappa}\Phi^S(\beta)]'W[\Phi^A - \frac{1}{\kappa}\sum_{j=1}^{\kappa}\Phi^S(\beta)] \tag{13}$$

where W denotes the optimal weighting matrix, and Φ^A and $\Phi^S(\beta)$ denote the vector of m actual moments and the vector of m simulated counterparts, respectively. The weighting matrix W is given by the inverse of the variance–covariance matrix of $[\Phi^A - \frac{1}{\kappa}\sum_{j=1}^{\kappa}\Phi^S(\beta)]$, which is best estimated using the following matrix (see Lee and Ingram (1991)):

$$W = [(1 + \frac{1}{\kappa})\Omega]^{-1} \tag{14}$$

Here, Ω denotes the variance–covariance matrix of the empirical moments, Φ^A. Ω is obtained by a block bootstrap with replacement of empirical data. An implication of using this weighting matrix is that moments with a large variation are given less weight than moments with a small variation. The distance between two sets of moments $\Phi^A$and $\frac{1}{\kappa}\sum_{j=1}^{\kappa}\Phi^S(\beta)$, $J(\beta)$, has a χ^2 distribution with $m - l$ degrees of freedom, where m is the number of moments.

When searching for values of β, an annealing cooling algorithm is used. On the basis of starting values for the estimated parameters, this routine takes random jumps in a predefined parameter space. The routine accepts worse solutions with a decreasing probability, which ensures that the global optimum is found. As the solution is somewhat sensitive to initial values, we do several computations with different starting values to ensure that we find the global maxima.

The standard errors of the parameters are given by the square roots of the diagonals of the variance–covariance matrix for $\hat{\beta}$, which is given by:

$$Q_s(W) = (1 + \frac{1}{\kappa})[\frac{\partial\Phi^S(\hat{\beta})}{\partial\beta}' W \frac{\partial\Phi^S(\hat{\beta})}{\partial\beta}]^{-1} \tag{15}$$

Here, $\frac{\partial\Phi^S(\hat{\beta})}{\partial\beta}$ is the Jacobian $m \times l$ matrix of the moment vector with respect to the parameter vector β, evaluated at $\hat{\beta}$. Given the lack of an analytical solution for the components of this matrix, numerical derivatives are used. More specifically, we use the symmetric difference quotient which is given by:

$$f'(x) \approx \frac{f(x+h) - f(x-h)}{2h} \tag{16}$$

[31] See, for instance, McFadden (1989); Pakes and Pollard (1989) for more details regarding the approach.

[32] In our estimations, we use $\kappa = 10$, and have $N = 1676$ and $T = 60$.

In the expression (16), x denotes the components of $\hat{\beta}$, $f(x)$ denotes the components of $\Phi^S(\hat{\beta})$, and h is a small positive number. A problem with this approach is that the approximation depends on the size of h. We therefore follow Bloom (2009) and calculate four values of the numerical derivative with steps of 0.1%, 1%, 2.5%, and 5% from $\hat{\beta}$, and use the median value of these numerical derivatives.