Mixed Frequency Data in a Heterogenous Sticky Price Model

BY Jonas Andersson, Øivind A. Nilsen and Hans J. Skaug

DISCUSSION PAPER

NHH

Institutt for foretaksøkonomi
Department of Business and Management Science

FOR 26/2025
ISSN: 2387-3000
December 2025

Mixed Frequency Data in a Heterogenous Sticky Price Model*

Jonas Andersson†
(Norwegian School of Economics)

Øivind A. Nilsen‡
(Norwegian School of Economics)

Hans J. Skaug§
(University of Bergen)

December 30, 2025

Abstract

We develop a model to estimate price-adjustment behavior when prices are observed more frequently than key explanatory variables such as wage costs. We propose a mixed-frequency stochastic (S, s)-model that accommodates infrequently observed costs and allows for plant-, product-, and season-specific heterogeneity. The model is estimated using a likelihood-based nonlinear state-space approach, enabling estimation as if all variables were observed at the same frequency. Applied to monthly price survey data matched with plant-level annual cost data for manufacturing producers, the model yields precise estimates of both cost pass-through and price-inaction thresholds, and reduces the blurring of intermittent price changes.

Keywords: Latent Variables, Pass-Through, Panel Data.
JEL classification: E31, D43, C34, E37, L16

*Acknowledgement: Thoughtful discussions with Benjamin Born, Erwan Gautier, Bertrand Koebel, Wilko Letterie, Lars Sørgard, Vincent Sterk and Peter Zorn are highly appreciated. We are also grateful for comments during presentations at the Annual International Conference on Computational and Financial Econometrics (CFE), the Annual Conference of the International Association for Applied Econometrics (IAAE), the Annual Conference of the European Association for Research in Industrial Economics (EARIE), University of Strasbourg, Center for Data Science (CEDAS) of the University of Bergen, the Norwegian School of Economics, the Nordic Econometric Meeting - NEM 2024, University of Stavanger, the Annual Meeting for Norwegian Association of Economists, and the 15th and final ifo Conference on Macroeconomics and Survey Data. The authors have no competing interests to declare. The usual disclaimer applies.

†Norwegian School of Economics, Department of Business and Management Science, Helleveien 30, 5045 Bergen, Norway. E-mail: jonas.andersson@nhh.no

‡(Corresponding author.) Norwegian School of Economics, Department of Economics, Helleveien 30, 5045 Bergen, Norway. E-mail: oivind.nilsen@nhh.no.

§Department of Mathematics, Postboks 7803, 5020 Bergen, Norway. E-mail: hans.skaug@uib.no

1 Introduction

Price stickiness may induce nominal shocks to have real effects. Even though most modern macroeconomic models assume price stickiness, these models are often only loosely linked to microeconomic building blocks. Moreover, much of the microeconomic empirical research on price changes relies on outlet- or scanner-level data for consumer goods, often leaving the cost side of pricing decisions unobserved. One reason for this is the lack of possibilities to link price changes to changes in costs at the plant or product level. Our paper takes advantage of survey data of producers' product prices and merges these with register data on their production costs. In this way we can test the relationship between costs shocks and the price change pattern.

A common method for analysing price stickiness is the use of the (S, s) model suggested by Sheshinski and Weiss (1977), and developed further by for instance Caballero and Engel (2007), Ratfai (2006), Alvarez et al. (2011), which assumes that firms keep the price fixed within certain bounds, denoted (S, s), and that the prices are changed only when the underlying optimal (frictionless) prices move outside these bounds. The bounds reflect the presence of a fixed cost incurred when changing prices, commonly referred to as a menu cost. To identify price stickiness empirically, a necessary condition is that time periods where the price is unchanged are observed. Since prices are rarely unchanged from one year to the next, identifying all components of an (S, s) model using annual data is difficult, if not impossible. Thus, identification requires a certain frequency of price observations. A dilemma is, however, that production costs, such as labour, are very seldom available on a monthly basis. Thus, both utilising the mixed frequencies of prices and the production costs seems highly relevant when analysing price stickiness.

When variables are observed at different frequencies, the most common solution in time-series analysis is to aggregate all variables to the lowest available frequency. However, more elaborate methods have been employed. One is the *mixed data sampling* (MIDAS) model, e.g. Ghysels et al. (2004), where the independent variables are observed with a higher frequency than the dependent one. Another is the mixed-frequency vector autoregressive (MF-VAR) model where a state space formulation is used to accommodate the different frequencies (see, Aruoba et al., 2009; Mariano and Murasawa, 2010; Zadrozny, 1990; Schorfheide and Song, 2015; Clements and Galvão, 2008; Guérin and Marcellino, 2013). Particularly related to our study are Foroni et al. (2018) and Xu et al. (2019) who, like us, use low frequency information for predicting high frequency variables in so-called reverse MIDAS models. The approach taken in those models is similar to ours. They assume that the high-frequency variable follows an exogenous explanatory variable which is partly observed. In the reverse MIDAS model, both flow- and stock variables for the explanatory variable are included as special cases. Our explanatory variable, production costs, is a flow variable and can thus be covered. This could, therefore, be a starting block for our modelling efforts. However, we need the (S, s)-model to account for price change thresholds. It is not clear to us how to do this within the reverse MIDAS framework.[1]

Our aim is to estimate price-adjustment behavior and cost pass-through using micro data, while correctly accounting for mixed frequencies and price-adjustment thresholds. Within a non-linear state-space framework, we model several key components as latent variables, including unobserved monthly production costs, firm- and product-specific heterogeneity, and the price-adjustment thresholds in the (S, s)-model. In this sense, our approach can be viewed as an extension of mixed-frequency VAR (MF-VAR) models to a nonlinear, state-dependent pricing environment. The mixed-frequency problem is addressed by specifying producers' monthly prices as functions of latent monthly marginal wage costs, subject to the restriction that these monthly costs aggregate to observed annual wage costs. The latent monthly wage cost is assumed to follow a random walk, reflecting strong persistence in production costs at the monthly level. Estimation of this model is challenging due to the presence of many latent variables and random effects. The starting

[1] An interesting recent work on a nonlinear reverse MIDAS model is Motegi et al. (2025) who develop a Threshold AutoRegression (TAR) model with higher frequencies in the dependent variable. The estimation method they use exploits that the model can be transformed in to a linear model which can be estimated by standard regression methods, something not possible for our (S, s)-model.

point, both in the method used in this paper and in the EM algorithm (Dempster et al. (1977)), is to formulate the log-likelihood as if the latent variables were observed, that is, the so-called full likelihood. In our case, this means formulating the log-likelihood as if wage costs were observed at the monthly frequency and conditional on the random effects. Since the latent variables are not observed, inference requires integrating them out to obtain the likelihood of the observed data. In the EM algorithm, this is achieved by computing expectations of the full log-likelihood conditional on the observed data, which in our setting would be analytically extremely complicated, if not infeasible. We therefore adopt a likelihood-based approach that integrates out latent variables numerically using a Laplace approximation. The practical implementation relies on modern tools for efficient likelihood evaluation and optimization, combining automatic differentiation (Griewank and Walther (2008))[2] with Laplace approximation as implemented in the so-called Template Model Builder (`TMB` (Kristensen et al., 2016)).

Our study contributes to the literature in several ways. First, we formulate an empirical model allowing for different observational frequencies of the prices and the cost variable(s) in an (S, s)-model. To our best knowledge, our paper is the first one exploiting data with different frequencies on the dependent and independent variables in an (S, s)-model, and we validate the performance of the proposed estimator through Monte Carlo simulations. Second, we extend the models of both Fougère et al. (2010) and Dhyne et al. (2011) by utilising plant-specific wage costs instead of just sector-specific variables. This possibility is quite unique in this literature. Third, we study the occurrence of price changes and stickiness using combined survey and register data from the manufacturing industry in Norway. In most macro-economic models, imperfect competition occurs at the producer level, while the prices in the consumer markets are decided as in competitive markets. Thus, our paper extends our knowledge about price change behaviour for the units at the level which are essential for macroeconomic modelling and central bank policies. More

[2]Automatic differentiation is a computationally efficient and numerically accurate method for evaluating derivatives in complex models. While widely used in other fields, it has so far seen relatively limited use in applied economic research.

specifically, our data are based on survey data used to construct the producer price index (PPI) for Norway. They cover monthly price quotations for a representative sample of Norwegian producers' products, and the selected firms are repeatedly surveyed. The price data are linked to plant level data such as wage costs, the most important cost factor in most plants. We estimate the pass-through of labour costs with our suggested mixed frequency approach. Compared to the use of annual data, (see, for instance, Carlsson and Skans (2012)), the monthly frequency of our price data relative to annual frequency is a huge advantage, as the time aggregation might have important implications for our understanding of sticky prices. Fourth, our model builds on the approaches by Fougère et al. (2010) and Dhyne et al. (2011), which allow for stochastic bounds. These stochastic bounds are important as they account for substantial heterogeneity in the frequency and magnitude of price changes observed in the data. In addition to capturing heterogeneous price-setting behaviour across sectors, plants, and products (see, for instance, the overview by Klenow and Malin (2010)), stochastic bounds also allow for periods of price inaction, while in other periods - for the same product, prices might change every month, even when the underlying driving forces remain the same across these periods. The stochastic bounds, together with product heterogeneity and firm- and product-specific shocks, make it possible to estimate a single *one-size-fits-all* model that accommodates multiple industries and products simultaneously, rather than requiring separate estimations for each industry or product category.

The estimates reveal economically meaningful adjustment thresholds, with substantial variation across products and over time. Consistent with findings in other studies, the price change frequency is higher in the beginning of the year, relative to the other months, indicating that the menu costs are smaller in January and February, at least for positive price changes. The immediate pass-through of marginal costs to prices is approximately 0.40 which seems reasonable with monopolistic competition and standard curvature of the production function. Finally, the analyses also show that the intermittency is blurred when aggregating across time and/or plant/products.

The remainder of the paper is organized as follows. Section 2 describes the model. Section 3 presents the estimation methodology together with an appendix which presents simulation results, showing the performance of the estimator. Section 4 describes the data used for the illustration, followed by Section 5 which reports and discusses the results. The final Section 6 gives some concluding remarks and closes the paper.

2 The model

2.1 The theoretical model

We start by looking at a single-product firm facing the following iso-elastic demand function,

$$D(P_t) = B_t \cdot P_t{}^{-e}, \tag{1}$$

where B_t $(=B_0 \exp(\xi_t^D))$ is a shift parameter which includes demand shocks ξ_t^D, P_t is the price, and e denotes the demand elasticity ($e > 1$). A standard Cobb-Douglas production function is assumed

$$Q_t = A_t \cdot L_t{}^a, \tag{2}$$

where A_t $(=A_0 \exp(\xi_t^S))$ denotes a stochastic productivity shock, and a is a constant ($0 < a < 1$). Assuming wages F_t are set in the market, that demand equals supply, and using the two preceding eqs., the profit for a firm, π_t, can be expressed as follows

$$\pi_t = P_t \cdot (B_t \cdot P_t^{-e}) - F_t \cdot \left[\frac{B_t}{A_t}\right]^{\frac{1}{a}} \cdot P_t^{-\frac{e}{a}}. \tag{3}$$

Taking the first-order condition with respect to P_t, and applying a log-transformation, we get

$$\ln P_t = h + \beta \cdot \ln F_t + v_t, \tag{4}$$

Here, h is a time-invariant, non-linear function of the supply- and demand shock parameters A_0 and B_0, the production function curvature parameter a, and the demand elasticity parameter e. The parameter in front of $\ln F_t$, $\beta = \dfrac{a}{e(1-a(1-\frac{1}{e}))}$, is the pass-through rate, the portion of the marginal cost change which will be passed through to prices.[3] Typical values used for calibration of pricing models in the literature, are $a = 2/3$, and $e = 3$, which would give a pass-through parameter of $\beta = 0.40$.[4] Furthermore, a demand elasticity where $e = 3$ corresponds to a markup of $\dfrac{e}{e-1} = 1.5$, a value that for instance De Loecker et al. (2020) find for the period 2000-2014 (see their Figure 1). The residual v_t might be seen as a linearization of the idiosyncratic parts of the two shock parameters, ξ_t^D and ξ_t^S.

The formulation described above assumes that prices adjust immediately to any kind of shock, a rather heroic assumption. Nevertheless, the description might serve as a picture of the frictionless price, the price which represents the equilibrium price if there are no price-adjustment costs, and not as a model of actual prices. To link the unobserved frictionless price to observed price change patterns with infrequent changes, we adopt an (S, s) price formulation. We start by noting a high frequency of zero price adjustments.[5] It is hard to believe that such a pattern is driven by demand- and supply-shocks only. On the other hand, this would be observed if some menu costs are present, and thus, if it is costly to continuously adjust the product prices. The (S, s) formulation captures that the frictionless target price (in logs) for a product time t, denoted by P_t^*, will not affect the actual (and observed) prices unless the "gap" between the frictionless price and existing price, P_t, is

[3]Here we have assumed a constant markup. We could have made the theoretical model richer by also having a conduct parameter, θ, which would describe the intensity of competition, equal to 0 if there is perfect competition and 1 under monopoly (see enlightened discussions about the cost pass-through by Amiti et al. (2019), Weyl and Fabinger (2013), and Genakos and Pagliero (2022)). A reason for this simplification is identification of the deep parameters, a, e, and potentially θ, and that we are mainly interested in the net effect of these parameters, i.e. the pass-through.

[4]Klette (1999) finds that the average firm in most Norwegian manufacturing industries face constant or moderately decreasing return to scale. Given that we include only labour as input factor, an $a < 1$ seems not unreasonable.

[5]Descriptives show a mass-point of zero price changes (70% of the observations) consistent with findings in the established literature.

larger than a threshold value S_t. The formation of the target price is determined by

$$\ln P_t = \begin{cases} \ln P_{t-1} \text{ if } |\ln P_t^* - \ln P_{t-1}| \leq S_t \\ \ln P_t^* \quad \text{otherwise,} \end{cases} \tag{5}$$

where S_t denotes a threshold. Thus, if the magnitude of the shock to the frictionless price is large enough relative to its target value, the firm finds it profitable to set a new price. For the ease of exposition, equation (5) presents the threshold as symmetric, that is, that the 'band of inaction' is the same whether the "price-gap" is positive or negative. Note however, as there is an upward trend in the wage costs, F_t, and if at the same time $\beta > 0$, which is expected, there will be an upward trend in the underlying optimal/frictionless prices despite the symmetry assumption of the threshold.[6] Following from this, we will then observe a higher share of positive price changes compared to negative price changes.

2.2 The empirical model

Based on the theoretical model presented above we specify a combined state-space and panel data regression model where the unobserved, monthly, frictionless price at time t for product i, produced by producer j, P_{ijt}^*, is determined by the unobserved, monthly, labour costs per worker, F_{jt}, through the relationship

$$\ln P_{ijt}^* = \alpha_{ij} + \beta \ln F_{jt} + (\varepsilon_{jt}^{\text{pl}} + \varepsilon_{ijt}^{\text{id}}). \tag{6}$$

where again, β, denotes the costs-to-price pass-through rate. The product specific time-invariant effect α_{ij} picks up for instance relative price difference to other products (and also common industry effects). The plant-specific $\varepsilon_{jt}^{\text{pl}}$ is meant to capture unobserved variation common to all products within plant j at time t. These could either be demand shocks, including seasonal demand effects, or supply shocks including prices of inputs other than

[6]The symmetry assumption will be relaxed further out in the paper.

labour.[7] The last term, $\varepsilon^{\text{id}}_{ijt}$, is idiosyncratic supply- and demand shocks.[8]

The various terms are distributed as follows; α_{ij} is $N(\mu_\alpha, \sigma^2_\alpha)$-distributed and independent over i and j, $\varepsilon^{\text{id}}_{ijt}$ is $N(0, \sigma^2_{\text{id}})$-distributed and independent across i, j, t, and ε_{jt} is $N(0, \sigma^2_{\text{pl}})$-distributed and independent across j and t. The ranges of the indices are $i = 1, \ldots, n_j$, $j = 1, \ldots, J$ and $t = 1, \ldots, T$, where n_j is the number of products for plant j. Identification of the latent monthly wage costs, F_{jt} (and $\varepsilon^{\text{pl}}_{jt}$), comes mainly from autocorrelation in the wages, monthly and yearly, and the independence of $\varepsilon^{\text{pl}}_{jt}$. In addition it is helped by the fact that for each month t the $\ln F_{jt}$ is identical for all products of a multi-product plant, while the price movements are not.

As one plant might produce one or several products, the multi-product aspect of a plant has to be addressed. In our model, this feature is encompassed in the following ways (in addition to through the plant-specific shocks $\varepsilon^{\text{pl}}_{jt}$ in the price equation). The demand-side part is included by assuming that each plant's products have additively separate demands, in line with the assumptions made by Midrigan (2011) and Alvarez and Lippi (2014) in their multiproduct analyses.[9] On the supply-side, the multiproduct feature is included by letting the monthly latent wage costs, F_{jt}, in the $\ln P^*_{ijt}$-equation to be common to all products within the same plant. This is also motivated by the fact that we have no information about how to split the wage costs between the various products. However, by assuming the cost shares of each product ij by a multiproduct producer j, b_{ij}, to be constant over the sample period, such that $F_{ijt} = b_{ij} F_{jt}$, our model expressed in (6) therefore also encompasses the fact that the cost-shares might not be the same across the products.[10]

[7] We follow Rotemberg and Woodford (1999) and Carlsson (2017) and consider wages as the 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 of 0.75-0.90 in the years of our study NOU (0196).

[8] In this equation, inflationary pressure arises only through the cost variable F.

[9] In 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 make that the estimates of the fixed menu cost will be smaller to capture the benefits of coordination due to market dependency.

[10] More formally, if one log-transforms $F_{ijt} = b_{ij} F_{jt}$ and substitute this into an equation where $\ln P^*_{ijt}$ is a linear function of $\ln F_{ijt}$, the α, and the two ε shocks, one gets $\ln P^*_{ijt} = (... + \beta \ln b_{ij}) + \beta \ln F_{jt} + (\varepsilon^{\text{pl}}_{jt} + \varepsilon^{\text{id}}_{ijt})$, i.e., almost identical to our equation (6).

In datasets with only yearly observations of prices, almost all prices will be changed over the 12 months period. The alternative is to keep monthly observed prices and distribute the yearly wage costs evenly over the year. The latter solution is not satisfactory either since there is no variation in the wage costs over the year, implying that when an actual price change occurs, a wage cost change will not be observed unless the price change happens in January. Our solution to this is to add a new component to the model. We assume that the log-transformed monthly wage costs, $\ln F_{jt}$, follow a random walk (RW) process with a drift ϕ_0. This reflects that we do not have information indicating at what time wages change significantly, for instance, due to the outcome of central or local wage negotiations. At the same time, we are aware that there is a long term trend in aggregate wages. Thus, a random walk with a drift is therefore our best guess for the development of the latent monthly wages. The process for wages costs F_{jt} is therefore implemented through

$$\ln F_{jt} = \phi_0 + \ln F_{j,t-1} + \eta_{jt}, \tag{7}$$

where η_{jt} is $N(0, \sigma_\eta^2)$-distributed and independent over j and t.[11]

Even though monthly observations for wages, F_{jt}, are unavailable, and therefore a latent variable in our model, we do observe wage costs on an annual basis, denoted W. Thus, we have the following aggregation of the latent, monthly wages

$$W_{j,12s} = F_{j,12s-11} + F_{j,12s-10} + ... + F_{j,12s}, \tag{8}$$

for $s = 1, 2, ..., T/12$. The index s denotes year and the observed wage cost for year s has the index $12s$ signifying that the yearly wage is observed every 12th month.[12]

[11]We also estimated our model using an AR(1) process with a trend for the latent wages $ln(F_{jt})$. These results strongly point towards a random walk with drift process.

[12]By recursive substitution, it is evident that $\ln F_{jt}$ depends on the initial plant-specific value $\ln F_{j0}$ (and consequently on W_{j0}). Therefore, equation (7) implicitly accounts for firm-specific heterogeneity.

The observed price, P_{ijt}, is related to the frictionless target price P^*_{ijt} through

$$\ln P_{ijt} = \begin{cases} \ln P_{ij,t-1} \text{ if } -S^-_{ijt} \leq (\ln P^*_{ijt} - \ln P_{ij,t-1}) \leq S^+_{ijt} \\ \ln P^*_{ijt} \qquad \text{otherwise,} \end{cases} \tag{9}$$

where S^-_{ijt} and S^+_{ijt} are assumed to follow truncated normal distributions $TN(\mu^s_{ijt}, \sigma^2_s, 0, \infty)$ and to be independent over i, j and t, i.e., independent of both $\varepsilon^{\text{pl}}_{jt}$ and $\varepsilon^{\text{id}}_{ijt}$ also. The random and asymmetric thresholds S^-_{ijt} and S^+_{ijt} reflect that there might be heterogeneity between the various products and their thresholds over time, and allows the thresholds to be different depending on the sign of the price change. Such a formulation also allows for the existence of both large and small price changes over time even for the same product. Thus, there might periods in which firms refrain from adjusting despite large deviations from target prices. As a result, average threshold estimates from stochastic models can be smaller while still implying substantial price rigidity.[13] The adjustment thresholds S^-_{ijt} and S^+_{ijt} are identified from the joint behavior of prices and costs over time. Observed variation in plant-level wage costs induces movements in latent monthly costs F, and through the pass-through parameter β in the P^*-equation. Periods of price inaction following such cost-induced movements provide information about the width of the inaction region. Repeated observations of cost changes and price responses across plants and time therefore jointly identify the threshold parameters and the cost pass-through.

Also Fougère et al. (2010) and Dhyne et al. (2011) have a random threshold in their models, but where the threshold is assumed to be normally distributed, $S \sim N(.,.)$. Different from these authors, our truncated normal distribution $S \sim TN(\mu^s_{ijt}, \sigma^2_s, 0, \infty)$ ensures $S_{ijt} \geq 0$, which is intuitively more appealing as the threshold is meant to mirror menu costs and that negative menu costs are hard to justify.[14]

[13]Since price adjustment decisions depend only on the difference, $\ln P^*_{ijt} - \ln P_{ij,t-1}$, relative to the random thresholds, S^+_{ijt} and S^-_{ijt}, a permanent fixed component in the $\ln P^*$-equation instead of the chosen random α_{ij} would not be separately identified from the thresholds, and therefore, not be meaningful.

[14]Following Dotsey et al. (1999), Gautier and Le Bihan (2011) use instead a Beta distribution to ensure non-zero menu costs. Our use of a truncated normal distribution is inspired by Vincent (2019).

Descriptive analysis shows that the sampled products have excess price change frequency in January and February. Such a pattern is also described by Nakamura and Steinsson (2008), Vermeulen et al. (2012), and for Norway by Nilsen and Vange (2019). This seasonality could be explained by the producers' economic environments, seasonal demand effects, or seasonal pricing effects related to negotiation and of signing of price contracts, described by for instance, Zbaracki et al. (2004). A more detailed look at the monthly price change frequencies indicates that the described seasonal pattern is mainly present for positive price changes and only a smaller degree for negative price changes. Thus, going stepwise forward, we end up with a model where the thresholds take different values throughout the year; one value for January-February for positive price changes, one for positive price changes rest of the year, one value for January-February for negative price changes, and finally one for negative price changes for March-December. This is accomplished by including covariates, determining the distribution of the thresholds for season and positive vs. negative changes, in the model,

$$\mu_{ijt}^{s} = a_s + b_s D(\text{JanFeb}_t) + g_s D(\text{neg}_{ijt}), \tag{10}$$

where $D(\text{JanFeb}_t)$ and $D(\text{neg}_{ijt})$ are dummy variables for January-February and negative changes, respectively.[15]

Taken together, equations (6)–(10) define our mixed-frequency, state-dependent pricing model in which both cost pass-through and adjustment thresholds are identified from observed price and cost dynamics. The mixed-frequency model can be seen as a non-linear state-space model. In this interpretation, the state variables are the frictionless price P_{ijt}^{*}, the latent monthly wage F_{jt}, the random effects α_{ij}, $\varepsilon_{jt}^{\text{pl}}$ and the threshold S_{ijt}. The observed variables are the actual prices P_{ijt} and yearly wages $W_{j,12t}$.

[15]It would not be a problem to include monthly dummies in the threshold (see for instance Fougère et al. (2010)). However, as already pointed out, January and February are standing out with higher frequencies of (positive) price changes, and thus, our model are allowing for the most important seasonal pattern of price changes.

3 Estimation and methodology

In the models by Dhyne et al. (2011) and Fougère et al. (2010), the explanatory variables are directly observable. We build on their approaches, and address the mixed-frequency characteristic of our data by treating the monthly labour costs as a latent variable following a random walk with drift (7). This latent variable is defined so that the twelve monthly labour costs always sum up to the observed yearly labour cost (8). The vector of all model parameters, excluding latent variables, to be estimated by maximum likelihood is $\boldsymbol{\theta} = (\mu_\alpha, \sigma_\alpha, \beta, \sigma_{\text{pl}}, \sigma_{\text{id}}, \phi_0, \sigma_\eta, a_s, b_s, g_s, \sigma_s)'$.

The conditional log-likelihood for the *full data*, given explicitly in Appendix A, is built upon the one in Dhyne et al. (2011). In order to compute the likelihood for the *observed data* we need to integrate the joint density of data and the latent variables, α_{ij} for $i = 1, ..., n_j$, $j = 1, ..., J$, and both $\varepsilon^{\text{pl}}_{jt}$ and F_{jt} for $j = 1, ..., J$, $t = 1, ..., T$:

$$L(\boldsymbol{\theta}) = \int\int\int L(\boldsymbol{\theta}|\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})p(\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})d\boldsymbol{F}d\boldsymbol{\alpha}d\boldsymbol{\varepsilon^{pl}}, \tag{11}$$

where $p(\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})$ represents the density of the latent variables as defined in Section 2. Since this integration must be performed for each iteration of the numerical optimization of the likelihood, computational efficiency is essential. In our case this is achieved by using the R-package `TMB` (Kristensen et al., 2016) which automatically evaluates the Laplace approximation of the high dimensional integral (11), given `C++` code[16] for the conditional log-likelihood $\ln L(\boldsymbol{\theta}|\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})$ and the prior density $\ln p(\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})$. We work exclusively with the Laplace approximation of (11), and details of this approximation and its implementation are provided in Appendix B. `TMB` also evaluates the derivative of the Laplace approximation of $L(\boldsymbol{\theta})$ using automatic differentiation, ensuring efficient and accurate numerical optimization.

The log-likelihood function includes a bivariate cumulative distribution function (CDF),

[16]Recently, Kristensen (2023) has developed the R package `RTMB` which makes it possible to use `TMB` without writing C++ code. However, currently only a subset of `TMB` is supported.

denoted $\Phi_2(,,)$. This function has posed some practical problems as the bivariate CDF is not a native function in the TMB-package (in contrast to the univariate normal CDF). We therefore employ the approximation proposed by Cox and Wermuth (1991), which provides a good compromise between numerical accuracy and computational efficiency. To assess the finite-sample performance of the estimator and the reliability of the reported standard errors, we conduct a Monte Carlo study based on the estimated model. The results, reported in Appendix C, indicate that the estimator performs well in capturing the key properties of the model.

4 Data

The empirical analysis is based on two different datasets, both provided by Statistics Norway (SSB). The most important dataset for this study is the survey data behind the commodity price index for the Norwegian manufacturing industry (PPI). SSB collects the data monthly for a selection of Norwegian plants. 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).[17] Plants are repeatedly surveyed, participation is compulsory, the required information is collected through electronic reporting (mainly), and Statistics Norway revises the data regularly to detect measurement errors and nonconformity.[18] Furthermore, the gathered data is subject to several controls aiming at identifying extreme values and mistyping. Considering this, and that the PPI is an important tool for governing bodies, it is fair to assume that the data are of high quality

[17]In the remainder of the paper use the terms firm and producer interchangeably. Note that the selection routines by Statistics Norway are such that the size of the firms sampled for the PPI is in general larger than the mean size of a firm in the manufacturing sector. Furthermore, an underlying assumption in our approach is that that the pricing decisions are likely to be delegated to the firm level. Note that before we balance the data sample (see further details later on) the dataset contains a yearly average of 339 firms and 361 plants. Thus, the reality of the pricing dataset in isolation is that there is a 1-to-1 relationship between most firms and their associated plants in the price data.

[18]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.

and representative for Norwegian producers.[19] The price data cover the years 2005-2016.

The monthly prices are merged with annual plant level information, using the plant identifiers in the price survey data. The plant level information includes records on the number of employees, net revenues, and wage costs. Thus, these additional data allow us to investigate and control for attributes such as plant size (number of employees) and cost (wage) shocks in our analyses. Our plant-level data are collected for public registers and have universal coverage. Furthermore, they are scrutinized by auditing firms and the tax authorities before being made available for aggregate public statistics and research. Hence, also the plant-level data are of a high quality.

In our data cleaning process, we exclude observations where prices are reduced by more than 49%, or increased by more than 99%.[20] Furthermore, missing or incomplete records are deleted. We exclude the SIC sectors representing mining and quarrying, plants related to the energy sector (oil, gas, electricity, etc.), water supply and sewage, and wholesale and retail. These sectors are of little relevance when analysing producer pricing behaviour. In the construction of the final data set for this study, plants with less than 10 employees are excluded. Only products for which we have price information all twelve months in a given year are included. The original data set contains prices for 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.

Since our (S, s)-model is a typical state-dependent model, we try to remove products whose price-change pattern is better explained by time-dependency.[21] We use the HS code to categorize the products.[22] For each 4-digit product HS group the number of price changes in a given year is divided by the number of price observations, i.e., the average

[19]A comparison of the data to the European reference literature (summarized by Vermeulen et al. (2012)) shows that Norwegian firms' pricing pattern is more or less in line with what is observed for the rest of Europe (see Letterie and Nilsen (2022)).

[20]A price change is calculated as $\frac{P_{ijk,\ t} - P_{ijk,\ t-1}}{P_{ijk,\ t-1}}$, where i denotes product, j plant in industry k in period t.

[21]See for instance Klenow and Kryvtsov (2008), Midrigan (2010), Dias et al. (2013), and more recently Dixon and Grimme (2022) for interesting discussions about state-dependent versus time-dependent models' ability to describe moments of various price datasets.

[22]The Harmonized System, HS, is an international customs and statistical nomenclature SSB (2020).

share of price changes per year, per product category. We name the HS 4-digit product groups into the behavioural categories “frequent” and “infrequent” price changes, such that if a HS 4-digit product group has an average share of price changes below 0.09 per year, i.e. slightly above one change per year, it is defined to be a group with ”infrequent” price changes, and “frequent” otherwise. This procedure categorizes approximately one third of the distinct products as ”infrequent” price changers, and the rest as ”frequent” changers. In the remainder of the paper we focus on the frequent changers as these are more likely to be products who’s behaviour is better described as state-dependent model while the ”infrequent” ones are more likely to be products who’s behaviour is better described as time-dependent.

Finally, we include only products for which we have observations for all months in all years. This balancing is due to convenience only when programming and testing the model empirically. This leaves us with a final dataset covering the manufacturing industry containing 36 432 price observations for 253 products, coming from 105 plants over the years 2005–16. Among these we find 43 single product plants and 62 multiproduct plants. Thus, a majority of the observations are from multiproduct plants.

5 Empirical results

The estimation results, coefficients and the standard errors, are given in Table 1. The first observation we make, before getting into details, is that all of the estimated parameters have small standard errors.

In the first column of Table 1, denoted Model 1, we include both plant-specific and a product-specific shocks, $\varepsilon^{\mathrm{pl}}_{jt}$ and $\varepsilon^{\mathrm{id}}_{ijt}$, while we at the same time allow for asymmetric thresholds but ignores seasonality. As reported in Table 1, column 1, there is a quite high variation in prices among the various products, with a $\sigma_\alpha = 2.690$. The estimate of the pass-through parameter is $\beta = 0.385$. This is much smaller than 1 and says that a change in the marginal costs is not fully reflected in the prices charged. Our β-estimate is a function

	Model 1		Model 2		Model 3		Model 4	
	Coef	SE	Coef	SE	Coef	SE	Coef	SE
Price equation								
μ_α	4.102	0.171	4.149	0.171	4.166	0.172	4.104	0.169
σ_α	2.690	0.120	2.697	0.120	2.697	0.120	2.689	0.117
β	0.385	0.008	0.378	0.007	0.380	0.008	0.386	0.007
σ_{pl}	0.119	0.001			0.108	0.001	0.118	0.000
σ_{id}	0.083	0.001	0.130	0.001	0.081	0.001	0.083	0.000
Wage evolvement								
ϕ_0	0.003	0.000	0.003	0.000	0.003	0.000	0.003	0.000
σ_η	0.019	0.000	0.019	0.000	0.019	0.000	0.019	0.000
Price thresholds								
$\mu_S^{\text{JanDec+}}$	0.009	0.002	0.006	0.005				
$\mu_S^{\text{JanDec-}}$	0.214	0.003	0.205	0.008				
μ_S^{JanFeb}					0.053	0.006		
μ_S^{MarDec}					0.113	0.002		
$\mu_S^{\text{JanFeb+}}$							-0.009	0.001
$\mu_S^{\text{MarDec+}}$							0.014	0.002
$\mu_S^{\text{JanFeb-}}$							0.196	0.002
$\mu_S^{\text{MarDec-}}$							0.218	0.002
σ_S	0.169	0.001	0.273	0.003	0.202	0.002	0.169	0.001
plant-specific shocks	Yes		No		Yes		Yes	
asym. thresholds	Yes		Yes		No		Yes	
seasonal threshold	No		No		Yes		Yes	

Table 1: Estimation results - the mixed frequency model.
Notes: Parameter estimates and st. errors (SE). The number of plants is 105, total number of products 253 and number of observations 36432. For models 1 and 2 we set $b_s = 0$ in equation (10) while we in model 3 set $g_s = 0$ and in model 4 we allow all parameters to be free.

of the production function curvature parameter a, and the demand elasticity parameter e. As previously discussed, typical values for a and e are 2/3 and 3, respectively, resulting in a pass-through parameter of $\beta = 0.40$. Thus, our estimate of the pass-through is therefore comparable to already existing estimates. It also interesting to compare our estimate of the β to the one by Carlsson and Skans (2012) who use Swedish annual manufacturing industry data. They find a β in the range of 0.25-0.50. Amiti et al. (2019), using annual data from the Belgian manufacturing industry, have β in the range of 0.59-0.65 (see their Table 1). Our β-estimate appears to align well with these findings.[23, 24]

Focusing on the two shock parameters, σ_{pl} and σ_{id}, we see that both of them are statistically significant, where the magnitude of the first is somewhat larger than the second. Thus, taking into account both types of shocks seems to be highly relevant when analysing multiproduct pricing behaviour. With such plant-specific shocks, desired prices of goods within the plant become more likely to move in the same direction.[25] It is also essential to include the plant-specific shock, $\varepsilon_{jt}^{\text{pl}}$, given that our marginal costs or wage-variable, $\ln F_{jt}$, is plant-specific, and exclusion might induce omitted-variable bias in the estimation of the costs pass-through parameter β.

Turning to the wage equation, $\phi_0 = 0.003$ indicates an upward drift corresponding

[23]In the literature, the pass-through parameter, β, is usually interpreted as how changes in marginal costs affect actual prices (p), rather than frictionless prices (p^*) as discussed here. Therefore, careful consideration should be given when interpreting the actual magnitude of the β parameter. Additionally, our analysis only presents the contemporaneous pass-through, represented by the β parameter. Since we assume that monthly ln-transformed wage costs, $\ln F_{jt}$, follow a random walk with drift, it becomes evident that the long-run pass-through is a non-linear function of β, ϕ_0, and past shocks ($\eta_{ij,t-k}$). Furthermore, the extent to which these past shocks will contribute incrementally to the long-run pass-through depends on the characteristics of the stochastic threshold S_{ijt}. Consequently, the analytical solution for the evolution of the long-run pass-through becomes a stochastic function. Given this complexity, conducting simulations may be the most straightforward approach to gaining a clear understanding of the long-run pass-through dynamics. However, we defer such an exercise to a future follow-up study.

[24]Our analysis is conducted under the assumption of constant elasticity of substitution (CES) demand, which implies constant markups. While this assumption delivers a convenient log-linear pricing rule, alternative demand functions — such as Kimball (1995) demand with variable elasticity—would introduce endogenous markup adjustments in response to cost shocks. Importantly, these alternative specifications are often observationally equivalent to CES or decreasing returns to scale in reduced-form price equations. As a result, the estimated pass-through and adjustment thresholds can be interpreted as capturing the combined effects of cost shocks and demand curvature, consistent with findings in studies such as Carlsson and Skans (2012), and Amiti et al. (2019).

[25]Bhattarai and Schoenle (2014), together with Midrigan (2011), rather model the within-firm product-shock dependency by allowing for a correlation of 0.65 among the good-specific productivity shocks.

to an annual nominal wage growth of 3.7%, which aligns with the observed aggregate wage indices over the sample period. This upward trend, combined with the positive β-parameter, implies that there will be more upward than downward price changes (if we were to use a model formulation with symmetric thresholds).

The parameters of the thresholds are stemming from a $TN(\mu^s, \sigma_s^2, 0, \infty)$ - distribution.[26] We see first of all, that the parameter estimate of $\mu_S^{\text{JanDec+}}$ is much larger than $\mu_S^{\text{JanDec-}}$, i.e., the modes of these two distributions are indeed different and indicate that the threshold is lower for price increase compared to price decreases. The conditional mean of the threshold is a function of μ^s as well as σ_s and is estimated to be 0.138 for increases and 0.247 for decreases.

When we compare our estimates to the ones by Dhyne et al. (2011), who use a model with stochastic thresholds assumed to be normally distributed, we see that our estimates are somewhat smaller compared to the ones reported in their Table 2.[27] Using Belgium and French CPI-data, they report their mean thresholds for broad product categories to be in the interval [0.22, 0.52] if we keep their results for energy and services aside. One reason for the difference, might be consumer goods versus intermediate goods. The estimate, $\hat{\sigma}_s = 0.169$, together with the estimated μ_s's, give us the conditional standard deviations of the threshold 0.103 for increases and 0.142 for decreases, while Dhyne et al. (2011) get [0.12, 0.25], again ignoring their estimates for energy and services. Thus, our estimates of the threshold parameters definitely seem to be reasonable and in line with existing estimates. Moreover, the substantial variation in thresholds across products and over time highlights the importance of allowing for heterogeneous and time-varying adjustment costs when studying price rigidity.[28]

In the second column of Table 1, denoted Model 2, we exclude the plant-specific shocks,

[26] In the truncated normal distribution (truncated at zero), the parameter μ^s may take a negative value, which implies that the mode of the truncated distribution is located at zero and that the density decreases monotonically as the threshold values increase.

[27] Since our standard errors are surprisingly small in comparison with many similar studies we performed a Monte Carlo study, presented in the Appendix C. The results there shows that the standard errors are good estimators of the actual standard deviations of the parameter estimators.

[28] We will show graphically the distribution of the estimated thresholds later.

$\varepsilon_{jt}^{\text{pl}}$. Doing this seems not to make the estimates related to the price equation and the wage equation too different from the results of Model 1 reported in column (1). We also see that σ_{id} has increased but still corresponds to the square root of the sum of the variance of the two shocks in column (1), $\left(= \sqrt{\sigma_{\text{pl}}^2 + \sigma_{\text{id}}^2}\right)$, (0.130 and 0.145, column (2) and column (1) respectively). However, ignoring the plant-specific shocks, $\varepsilon_{jt}^{\text{pl}}$, affects the spread of the thresholds quite clearly. Thus, when $\varepsilon_{jt}^{\text{pl}}$ is included in our model and this shock is large enough to induce a price change, it will necessarily involve several products. In this situation μ_S must therefore be somewhat smaller (compared to a model that includes $\varepsilon_{ijt}^{\text{id}}$ only) to avoid too many price changes, i.e., to match the number of price changes in the actual data. Nevertheless, and once more, taking into account both types of shocks seems to be highly relevant when analysing multiproduct pricing behaviour.

The results of Model 3, where we instead of modeling the price threshold as a function of the sign of the price changes, concentrate on the seasonal pattern with higher frequencies of price changes in January and February compared to the rest of the year. Again we see that the results for the price equation and the wage equations are rather similar to the ones reported for Model 1 and Model 2. The magnitude of the μ_S^{JanFeb} and μ_S^{MarDec} are in between the threshold estimates for Model 1 and Model 2. The same applies for the estimate of σ_S.

In Model 4, we control for both asymmetry and seasonal effects by including different intercepts, $\mu_S^{\text{JanFeb+}}$, $\mu_S^{\text{MarDec+}}$, $\mu_S^{\text{JanFeb-}}$, and $\mu_S^{\text{MarDec-}}$, motivated by the corresponding frequencies of price changes in the actual sample used. For simplicity, it is assumed that the seasonal difference is the same for positive and for negative price changes, i.e., $b_s = (\mu_S^{\text{JanFeb+}} - \mu_S^{\text{MarDec+}}) = (\mu_S^{\text{JanFeb-}} - \mu_S^{\text{MarDec-}})$. From these threshold estimates, we see that the weighted averages of $\mu_S^{\text{JanFeb+}}$ and $\mu_S^{\text{MarDec+}}$, and $\mu_S^{\text{JanFeb-}}$ and $\mu_S^{\text{MarDec-}}$, are similar to $\mu_S^{\text{JanDec+}}$ and $\mu_S^{\text{JanDec-}}$ for Model 1. The same is true for the spread, measured by σ_S. We also observe that all the standard errors indicate rather precise estimates, i.e., that both asymmetry and seasonal effects are relevant.

Given the estimates of Model 4, we calculate the conditional means and standard deviations for the thresholds of Model 4 in Table 2, while we in Figure 1 give a more complete picture of the estimated thresholds S.[29] These estimates are based on the estimates of the parameters μ_S and σ_S in the truncated normal distribution.

	mean	st.dev
JanFeb+	0.132	0.100
MarDec+	0.140	0.104
JanFeb-	0.235	0.139
FebDec-	0.251	0.143

Table 2: Conditional means and standard deviations for the thresholds of Model 4

The table shows significant difference between positive and negative price changes, while the seasonal effect for positive price changes is marginal. Figure 1 (at the end of this paper) shows that the distribution of the estimated thresholds S for positive and negative price changes are substantially different, as already seen in Table 2, with the mode of S close zero for price increases while approximately 0.20 for price decreases. We also observe that the curves that describe the distribution of $S^{\text{JanFeb+}}$ and $S^{\text{MarDec+}}$ overlap to a great extent, and also $S^{\text{JanFeb-}}$ and $S^{\text{MarDec-}}$.[30]

5.1 Overall fit of the mixed frequency model

Of significant interest is the model's ability to replicate the patterns observed in the actual data. Given the non-linearity of our model and the heterogeneity incorporated across several equations, the effects and dynamics of shocks are complex. To assess the goodness of

[29] Gautier and Le Bihan (2011) point out that in stochastic (S, s) models the expected value of the threshold cannot be interpreted as the typical change in the frictionless price required to trigger a price adjustment - unlike deterministic (S, s) models, where such an interpretation is valid.

[30] As both our pricing and wage equations include plant-specific explanatory variables and shocks, one could argue that also the threshold equation should include a plant-specific variable. Initially, we wanted to include the number of products produced, motivated by the findings of Buckle and Carlson (2000), and Bhattarai and Schoenle (2014). However, our dataset does not include prices for all products for all plants - only for the products that Statistics Norway uses to construct their PPI index. We therefore include the size of the plant, measured as `log(number of employees)`, as a proxy for the number of products produced. However, the null hypothesis that the threshold is the same for small and large plants (defined as having 20 and 100 employees, respectively) cannot be rejected at any reasonable significance level. More importantly, the estimation results are robust to whether the plant-specific variable is included in the threshold or not. These additional results are available on request.

fit, we evaluate the model's capability to match certain data properties through simulations. We start by retrieving the maximum a posteriori estimates for all values of $\ln F_{jt}$ from the TMB routine. There are only 11 values of this latent variable for each plant and year, so the 12th value is estimated based on Equation (8), using the sample values of $W_{j,12s}$ to ensure that the estimated F_{jt} values sum up to the observed annual wage cost. The values of α_{ij} and $\varepsilon_{jt}^{\text{pl}}$ are also obtained from the maximum a posteriori estimates from the TMB routine. This, along with the seasonal effects $\mu_S^{\text{JanFeb+}}$, $\mu_S^{\text{MarDec+}}$, $\mu_S^{\text{JanFeb-}}$, and $\mu_S^{\text{MarDec-}}$, as well as the other estimated parameter values from Model 4 in Table 1 enable us to simulate from the estimated model. Shocks are drawn randomly and independently from the normal distribution of $\varepsilon_{ijt}^{\text{id}}$ and the truncated normal distribution of S_{ijt}. We simulate 105 plants producing 253 products, mirroring our estimation sample.[31] We track these simulated plants over 12 years, making 100 draws and averaging over these to report the moments.

In Table 3, we report the actual number of price changes together with the outcomes generated from our simulations using the parameter estimates for Model 4 in Table 1. As starting value for the unobserved monthly costs for each firm j, $F_{j,1}$, one twelfth of the first observed yearly costs $W_{j,1}/12$ is used. The overall picture is that we achieve quite good predictions, with the moments of the simulated data closely aligning with the actual ones. We see that the simulated frequencies for MarDec+, JanFeb- and MarDec- closely match the actual data, whereas the simulated frequency for JanFeb+ is not matching the actual value quite as well. While the seasonal effect is assumed to be the same, in absolute terms, for negative as for positive values, the consequence in terms of seasonality in the frequency of price changes is larger for positive changes. This is because the threshold is smaller for positive changes implying that the threshold, relatively speaking, increases more for positive values. However, as can be seen in Table 3, there is a smaller seasonality for positive changes in the simulated than in the actual data but a slightly larger seasonality for negative changes. This indicates that the seasonality for positive and negative changes

[31] The distribution of products per plant is the same during the simulations as in the estimation sample.

changes average each other out to some extent. Furthermore, this average seasonality for positive and negative changes is somewhat too small relative to what we see in the actual data, which might be due to the bias when it comes to the seasonality parameter b_s discussed in Appendix C. It is also interesting to observe that the model capture the overall price inflation, shown by the fact that the mean of the differences of the logarithmic prices are similar in the actual and the simulated data, very well, 2.6% and 2.9%, respectively. This is despite the fact that this is only modelled by a drift parameter, ϕ_0, in the wage cost equation (7).

	actual	simulated
JanFeb+	23 %	18 %
MarDec+	16 %	15 %
JanFeb-	13 %	13 %
MarDec-	12 %	11 %
Inflation	2.6 %	2.9 %

Table 3: Monthly frequencies of price changes - actual and simulated

It is worth emphasizing that the numbers in Table 3 show a good fit despite the fact that we fit a "one-size-fits-all" model to a data set containing many industry sectors and type of goods. The plant-time random effect $\varepsilon_{jt}^{\mathrm{pl}}$ helps to capture this substantial heterogeneity. This "one-size-fits-all" is quite unique compared to the existing literature. For instance, Fougère et al. (2010) focus exclusively on the relatively homogeneous restaurant industry. Dhyne et al. (2011) estimate their model separately for 182 product categories in different industries, while Honore et al. (2012) divide the data into 70 distinct industries and fit their model individually to each.

The estimates highlight the importance of plant-specific wage costs for identifying the adjustment thresholds. With an estimated pass-through parameter $\beta \approx 0.4$, a 20% increase in wage costs implies an increase in the frictionless target price of roughly 8 percent. For a firm that has recently adjusted its price, so that the current price is close to the target price, a lack of subsequent price adjustment following such a cost shock indicates that the relevant adjustment threshold must exceed that deviation. In addition, a simulation exercise reported in the appendix as "The importance of latent F vs observed $W/12$", illustrates how the treatment of wage costs data in the mixed frequency setting in our model affects threshold identification. When annual cost data are available for the same producers as the price data, explicitly modelling latent monthly costs F can mitigate biases in the estimated thresholds relative to simpler approaches that mechanically distribute annual costs across months, $W/12$. This highlights that properly accounting for the temporal structure of cost data can be important for identifying price rigidities.

In the menu-cost literature, key structural objects such as cost pass-through and adjustment thresholds are often inferred indirectly from moments of the price-change distribution, sometimes under stylized assumptions about cost-shock processes (e.g. Alvarez et al. (2016)). By contrast, our approach exploits linked price and cost data, allowing these objects to be estimated jointly while accommodating persistent cost dynamics. Aggregate inflation in such models depends on the fraction of firms whose frictionless target prices cross their adjustment thresholds, implying that the effective slope of the Phillips curve varies with the size and direction of shocks (Gertler and Leahy (2008); Blanco et al. (2024)). The asymmetry we estimate—smaller thresholds for price increases than for price decreases—suggests that inflationary pressures may translate more rapidly into price changes than disinflationary pressures, providing a micro-founded channel for asymmetric and nonlinear Phillips-curve behavior (Karadi et al. (2024). In addition, the seasonal variation in adjustment thresholds suggests that the transmission of monetary policy shocks depends on their timing within the calendar year, consistent with Olivei and Tenreyro (2007).[32]

[32]Two natural extensions of the present framework remain. Recent evidence suggests that price responses to shocks can vary systematically with firms' cost structures (Ahlander et al., 2024). Thus, even though our model captures substantial heterogeneity in price adjustment through stochastic thresholds and firm- and product-specific shocks, extending our model to allow for heterogeneous pass-through could shed additional light on the sources of heterogeneity in price-setting behavior despite identification challenges. On the more technical side, extending the estimation framework to accommodate unbalanced panels would increase empirical applicability and reduce potential selection problems induced by the balancing of the dataset.

5.2 Temporal aggregation

As already pointed out, there are not that many studies for which one has monthly price observations for the whole manufacturing industry, and at the same time is able to link these product prices to detailed cost measures at the plant level. With our data, we are therefore able to analyse the importance of time aggregation when analysing pass-through. As pointed out in Giulietti et al. (2020), price rigidity is underestimated if prices are aggregated over time. Therefore, decisions made based on such data will underestimate the persistence of macroeconomic shocks.

To shed light on this, we use the annual price averages together with the annual wage cost variable. The model for the annual prices is a simplified version of the model described in Section 2.2. Since we the observe the labour cost on the same frequency as the prices the equations (7) and (8) is not present and $\ln F_{jt}$ is substituted with $\ln W_{jt}$ in equation (6). In this way we get the same features of the data as for instance Carlsson and Skans (2012), Carlsson (2017), and Amiti et al. (2019). Before doing any estimations, it is of course interesting to note that the overall price change inaction has dropped from 77% on a monthly basis to 22% per year for our sample of frequent price changers. The estimation results of this time-aggregation exercise are presented in Table 4.[33]

Most importantly, we observe that the estimates of the thresholds, $\mu_S^{\text{JanDec+}}$ and $\mu_S^{\text{JanDec-}}$, have significantly decreased compared to those obtained from monthly price data, Table 2. The corresponding conditional mean and standard deviations are 0.004 and 0.004, for price increases, and 0.045 and 0.014 for price decreases. Thus, the price rigidity, particularly for positive changes, will, for all practical purposes, be non-observable, when a model is fitted to the time-aggregated prices. This reduction in the means (and standard deviations) is not surprising, as the intermittency observed in monthly price data naturally diminishes

[33]Compared to the results shown in Table 1, the decrease in μ_α is approximately $\beta \cdot ln(12)$. This is because the regressor in the aggregate model is $ln(W)$, whereas in Table 1, it is $\ln(F)(\approx \ln(W/12))$. Furthermore, the estimate of β shows a deviation in comparison to the results in Table 1, which can be attributed to several factors. Firstly, certain elements present in the monthly model are absent in the yearly model: the dynamics of wage costs, including the drift, and the seasonal dummies in the thresholds (for obvious reasons). Secondly, variables undergo nonlinear transformations during aggregation (since $\ln(a) + \ln(b) \neq \ln(a + b)$). Finally, the estimated thresholds are smaller when using yearly data.

	With thresholds		No thresholds	
	Estimate	SE	Estimate	SE
Price equation				
μ_α	3.653	0.232	3.355	0.227
σ_α	2.697	0.120	2.695	0.120
β	0.310	0.026	0.359	0.025
σ_{pl}	0.131	0.000	0.137	0.003
σ_{id}	0.071	0.001	0.077	0.001
Price thresholds				
$\mu_S^{\text{JanDec+}}$	-0.036	0.002		
$\mu_S^{\text{JanDec-}}$	0.045	0.002		
σ_S	0.014	0.001		
plant-specific shocks	Yes		Yes	
asym. thresholds	Yes		–	

Table 4: Prices averaged over 12 months.
Notes: The number of plants, products, and observations is 105, 253, and 3,036, respectively.

considerably due to temporal aggregation. Therefore, time-aggregation plays a crucial role in analyzing both the intermittency and the pass-through of cost shocks to producer prices.

In the second column of Table 4 denoted "No thresholds", we show the results of what we refer to as a *two-way crossed random effects model*, i.e., our model where we ignore the thresholds completely. We see that both the order of magnitude and the significance of all the coefficients in the price equation are very similar to the one with thresholds (column 1), where we actually include the thresholds. Thus, if we were only interested in the pass-through parameter β we do quite well even if we ignore the threshold. Nevertheless, aggregating across time blurs the intermittency seen in the actual data.

6 Concluding remarks

The way in which micro-level prices respond to shocks is central to understanding macroeconomic dynamics. Accordingly, knowledge about firms' pricing decisions is crucial for constructing models used in macroeconomic analysis. Much of the existing empirical literature on price setting focuses on consumer goods, whereas the manufacturing sector—which underlies price rigidities in most macroeconomic models—has received comparatively less attention.

In this paper, we study price adjustment using a stochastic (S, s) panel data model, in which prices change only when the unobserved frictionless target price deviates sufficiently from the current price, giving rise to price rigidities. We combine high-quality monthly price survey data used to construct the Norwegian Producer Price Index with scrutinized plant-level register data for manufacturing producers. While prices are observed at a monthly frequency, production costs are available only annually, resulting in a mixed-frequency setting. We address this challenge by modeling latent monthly marginal costs that sum to the observed annual wage costs, within a pricing framework that allows us to estimate cost pass-through and price-adjustment thresholds.

Our analysis relies on estimation using Template Model Builder (TMB), which combines automatic differentiation with Laplace approximation to efficiently integrate out latent variables and random effects. This framework allows us to estimate a mixed-frequency stochastic (S, s) panel data model that combines a panel data structure — incorporating product-specific random effects $\boldsymbol{\alpha}$ and plant-specific shocks $\boldsymbol{\varepsilon^{pl}}$ — with a state-space model that includes a latent time-series process for production costs, $\boldsymbol{F}$. The generality of the Laplace approximation within TMB makes it possible to treat latent variables in a unified way, enabling the combination of panel data and state-space components within a single likelihood-based framework. From a practical perspective, this flexibility is valuable in empirical settings involving mixed frequencies and latent variables, as it allows rich forms of heterogeneity to be accommodated without relying on analytical approximations or numerical differentiation. More broadly, the approach provides a tractable likelihood-based framework for estimating high-dimensional latent-variable models in empirical settings with mixed observational frequencies.

The estimated thresholds imply that prices are adjusted only when deviations between the current price and the frictionless target price reach approximately 13–14% for price increases and 24–25% for price decreases. Thresholds are smaller at the beginning of the year, particularly for price increases. The immediate pass-through of marginal costs to prices is estimated to be approximately 0.40, consistent with monopolistic competition and standard assumptions about production technology. Substantial heterogeneity across sectors, plants, and products is present in the data and is captured by the stochastic elements of the model. At the same time, this heterogeneity can be accommodated within a one-size-fits-all specification, which can be estimated with high precision using a single model applied across industries and products. Taken together, these estimates provide empirically grounded inputs for state-dependent pricing models used to study inflation dynamics and monetary transmission, complementing existing approaches that largely rely on price-change moments.

References

Ahlander, E., M. Klein, and E. Pappa (2024, August). What flattens the supply curve? Working paper, Sveriges Riksbank.

Alvarez, F., H. Bihan, and F. Lippi (2016). The real effects of monetary shocks in sticky price models: A sufficient statistic approach. American Economic Review 106(10), 2817–2851.

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

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

Amiti, M., O. Itskhoki, and J. Konings (2019). International shocks, variable markups, and domestic prices. Review of Economic Studies 86, 2356–2402.

Aruoba, S. B., F. X. Diebold, and C. Scotti (2009). Real-time measurement of business conditions. Journal of Business & Economic Statistics 27(4), 417–427.

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

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

Blanco, A., C. Boar, C. J. Jones, and V. Midrigan (2024). Nonlinear dynamics in menu cost economies? evidence from us data. Technical report, National Bureau of Economic Research.

Buckle, R. A. and J. A. Carlson (2000). Menu costs, firm size and price rigidity. Economics Letters 66(1), 59–63.

Caballero, R. and E. Engel (2007). Price stickiness in ss models: New interpretations of old results. Journal of Monetary Economics 54(Supplement 1), 100–121.

Carlsson, M. (2017). Microdata evidence on the empirical importance of selection effects in menu-cost models. Journal of Money, Credit and Banking 49, 1803–1830.

Carlsson, M. and O. Skans (2012). Evaluating microfoundations for aggregate price rigidities: Evidence from matched firm-level data on product prices and unit labor cost. American Economic Review 102(4), 1571–1595.

Clements, M. P. and A. B. Galvão (2008). Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the united states. Journal of Business & Economic Statistics 26(4), 546–554.

Cox, D. R. and N. Wermuth (1991). A simple approximation for bivariate and trivariate normal integrals. International Statistical Review/Revue Internationale de Statistique, 263–269.

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. The Quarterly Journal of Economics 135(2), 561–644.

Dempster, A., N. Laird, and D. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society: Series B (Methodological 39(1), 1–22.

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

Dias, D., C. Marques, and F. Martins (2013). Choosing between time and state dependence: Micro evidence on firms' price reviewing strategies. The Scandinavian Journal of Economics 115(3), 756–780.

Dixon, H. D. and C. Grimme (2022). State-dependent or time-dependent pricing? new evidence from a monthly firm-level survey: 1980–2017. European Economic Review 150, 104319.

Dotsey, M., R. King, and A. Wolman (1999). State-dependent pricing and the general equilibrium dynamics of money and output. Quarterly Journal of Economics 114(2), 655–690.

Foroni, C., P. Guérin, and M. Massimiliano (2018). Using low frequency information for predicting high frequency variables. International Journal of Forecasting 34(4), 774–787.

Fougère, D., E. Gautier, and H. Le Bihan (2010). Restaurant prices and the minimum wage. Journal of Money, Credit and Banking 42(7), 1199–1234.

Gautier, E. and H. Le Bihan (2011). Time-varying (s, s) band models: Properties and interpretation. Journal of Economic Dynamics and Control 35(3), 394–412.

Genakos, C. and M. Pagliero (2022). Competition and pass-through: Evidence from isolated markets. American Economic Journal: Applied Economics 14(4), 35–57.

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

Ghysels, E., P. Santa-Clara, and R. Valkanov (2004). The midas touch: Mixed data sampling regression models. Mimeo.

Giulietti, M., J. Otero, and M. Waterson (2020). Rigidities and adjustments of daily prices to costs: Evidence from supermarket data. Journal of Economic Dynamics and Control 116, 103927.

Griewank, A. and A. Walther (2008). Evaluating derivatives: principles and techniques of algorithmic differentiation. SIAM.

Guérin, P. and M. Marcellino (2013). Markov-switching midas models. Journal of Business & Economic Statistics 31(1), 45–56.

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

Karadi, P., R. Schoenle, and J. Wursten (2024). Price selection in the microdata. Journal of Political Economy Macroeconomics 2(2), 228–271.

Kimball, M. S. (1995). The quantitative analytics of the basic neomonetarist model. Journal of Money, Credit and Banking 27(4), 1241–1277.

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

Klenow, P. and B. Malin (2010). Microeconomic evidence on price-setting. Handbook of Monetary Economics 3, 231–284.

Klette, T. J. (1999). Market power, scale economies and productivity: estimates from a panel of establishment data. The Journal of Industrial Economics 47(4), 451–476.

Kristensen, K. (2023). RTMB - R bindings for TMB.

Kristensen, K., A. Nielsen, C. Berg, H. Skaug, and B. Bell (2016). TMB: Automatic differentiation and laplace approximation. Journal of Statistical Software 70, 1–21.

Letterie, W. and Ø. A. Nilsen (2022). Pricing behaviour and menu costs in multiproduct firms. Economica 89, 746–769.

Mariano, R. and Y. Murasawa (2010). A coincident index, common factors, and monthly real gdp. Oxford Bulletin of Economics and Statistics 72(1), 27–46.

Midrigan, V. (2010). Is firm pricing state or time dependent? evidence from U.S. manufacturing. Review of Economics and Statistics 92(3), 643–656.

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

Motegi, K., J. W. Dennis, and S. Y. Hong (2025). Regular and reverse midastar models: Threshold autoregression with mixed frequency data.

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

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

NOU (2019:6). Norges offentlige utredninger - nou 2019:6 grunnlaget for inntektsoppgjørene (eng.; official norwegian reports 2019:6 "the basis for wage settlements".

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

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

Rotemberg, J. and M. Woodford (1999). The cyclical behavior of prices and costs. In in (ed. J. Taylor and M. Woodford (Eds.), Handbook of Macroeconomics, Volume 1B, pp. 1051–1135. Amsterdam: Elsevier.

Schorfheide, F. and D. Song (2015). Real-time forecasting with a mixed-frequency var. Journal of Business & Economic Statistics 33(3), 366–380.

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

SSB (2015). Commodity Price Index for the Industrial Sector. Statistics Norway.

SSB (2020). Producer price index. Statistics Norway.

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

Vincent, D. (2019). Estimating (S, s) rule regression models. In London Stata Conference 2019. Stata Users Group.

Weyl, E. G. and M. Fabinger (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. Journal of Political Economy 121(3), 528–583.

Xu, Q., X. Zhuo, C. , Jiang, F. Sun, and X. Huang (2019). Reverse restricted midas model with application to us interest rate forecasts. Communications in Statistics - Simulations and Computation 50(2), 462–482.

Zadrozny, P. A. (1990). Estimating a multivariate ARMA model with mixed-frequency data: an application to forecasting US GNP at monthly intervals. Working Paper Series 90-06, Federal Reserve Bank of Atlanta.

Zbaracki, M., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). Managerial and customer costs of price adjustment: direct evidence from industrial markets. Review of Economics and Statistics 86, 514–533.

Zheng, N. and N. Cadigan (2021). Frequentist delta-variance approximations with mixed-effects models and tmb. Computational Statistics & Data Analysis 160, 107227.

A The log-likelihood function

Let $d_{ijt} = \alpha_{ij} + \beta \ln F_{jt} - \ln P_{i,j,t-1}$ and remember that $\mu^s_{ijt} = a_s + b_s D(\text{JanFeb}_t) + g_s D(\text{neg}_{ijt})$ is the expected value of the distribution of the thresholds S_{ijt}. Furthermore, let $\boldsymbol{\alpha}$, $\boldsymbol{F}$ and $\boldsymbol{\varepsilon^{pl}}$ be vectors containing the (unobserved) variables α_{ij}, F_{jt} and $\varepsilon^{\text{pl}}_{jt}$ for all i, j and t, respectively. Then, the log-likelihood, treating $\boldsymbol{\alpha}$, $\boldsymbol{F}$ and $\boldsymbol{\varepsilon^{pl}}$ as observed, can be written

$$
\begin{aligned}
\ln L(\boldsymbol{\theta}|\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon}) = & \sum_{j=1}^{J}\sum_{i=1}^{n_j} \ln \phi\left(\frac{\alpha_{ij} - \mu_\alpha}{\sigma_\alpha}\right) + \sum_{j=1}^{J}\sum_{t=1}^{T} \ln \phi\left(\frac{\varepsilon^{\text{pl}}_{jt}}{\sigma_{\text{pl}}}\right) \\
& + \sum_{j=1}^{J}\sum_{t=2}^{T} \ln \phi\left(\frac{\ln F_{jt} - \phi_0 - \ln F_{j,t-1}}{\sigma_\eta}\right) \\
& + \sum_{j=1}^{J}\sum_{i=1}^{n_j}\sum_{t=2}^{T} I(\Delta \ln P_{ijt} > 0) \ln\left(\Phi\left(\frac{\Delta \ln P_{ijt} - \mu^s_{ijt}}{\sigma_s}\right) - \Phi\left(\frac{-\Delta \ln P_{ijt} - \mu^s_{ijt}}{\sigma_s}\right)\right) \\
& + \sum_{j=1}^{J}\sum_{i=1}^{n_j}\sum_{t=2}^{T} I(\Delta \ln P_{ijt} > 0) \ln \phi\left(\frac{\ln P_{ijt} - \alpha_{ij} - \varepsilon^{\text{pl}}_{jt} - \beta \ln F_{jt}}{\sigma_{\text{id}}}\right) \\
& + \sum_{j=1}^{J}\sum_{i=1}^{n_j}\sum_{t=2}^{T} I(\Delta \ln P_{ijt} < 0) \ln\left(\Phi\left(\frac{-\Delta \ln P_{ijt} - \mu^s_{ijt}}{\sigma_s}\right) - \Phi\left(\frac{\Delta \ln P_{ijt} - \mu^s_{ijt}}{\sigma_s}\right)\right) \\
& + \sum_{j=1}^{J}\sum_{i=1}^{n_j}\sum_{t=2}^{T} I(\Delta \ln P_{ijt} < 0) \ln \phi\left(\frac{\ln P_{ijt} - \alpha_{ij} - \varepsilon^{\text{pl}}_{jt} - \beta \ln F_{jt}}{\sigma_{\text{id}}}\right) \\
& + \sum_{j=1}^{J}\sum_{i=1}^{n_j}\sum_{t=2}^{T} I(\Delta \ln P_{ijt} = 0) \ln\left(\frac{\Phi_2\left(\frac{\mu^s_{ijt}}{\sigma_s}, \frac{\mu^s_{ijt} - d_{ijt}}{\sqrt{\sigma^2_{\text{id}} + \sigma^2_s}}, \frac{\sigma_s}{\sqrt{\sigma^2_{\text{id}} + \sigma^2_s}}\right)}{\Phi\left(\frac{\mu^s_{ijt}}{\sigma_s}\right)} \right. \\
& \qquad \left. - \frac{\Phi_2\left(\frac{\mu^s_{ijt}}{\sigma_s}, \frac{-\mu^s_{ijt} - d_{ijt}}{\sqrt{\sigma^2_{\text{id}} + \sigma^2_s}}, -\frac{\sigma_s}{\sqrt{\sigma^2_{\text{id}} + \sigma^2_s}}\right)}{\Phi\left(\frac{\mu^s_{ijt}}{\sigma_s}\right)}\right)
\end{aligned} \tag{12}
$$

where Φ and ϕ are, respectively, the CDF and PDF of the standard normal distribution and $\Phi_2(x, y, \rho)$ is the bivariate normal distribution with zero means, unit variances and correlation ρ. Note that μ^s_{ijt} is a function of covariates $D(\text{JanFeb}_t)$ and $D(\text{neg}_{ijt})$ and that the corresponding parameters a_s, b_s and g_s are contained in $\boldsymbol{\theta} = (\mu_\alpha, \sigma_\alpha, \beta, \sigma_{\text{pl}}, \sigma_{\text{id}}, \phi_0, \sigma_\eta, a_s, b_s, g_s, \sigma_s)'$. It should be noted that the yearly wages, W, only enters the log-likelihood implicitly through (8).

Details on how to derive the likelihood above, (12), can apart from the involved latent variables, be found in Dhyne et al. (2011).[34] Here, we only give a brief explanation of the different terms. The first three terms (sums) are the likelihood contributions from α, ε^{pl} and F which are subsequently integrated out. The fourth and fifth terms are the contributions from the observations with a positive price change while the sixth and seventh are contributions from the ones with a negative price change. Finally, the eight term is the likelihood contributions from observations with zero price changes.

B The Laplace approximation

Denote by $\boldsymbol{u} = (\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})$ the joint vector of latent variables as they appear in (11), and by $\hat{\boldsymbol{u}}(\boldsymbol{\theta})$ the value of $\boldsymbol{u}$ minimizing $f(\boldsymbol{u}, \boldsymbol{\theta}) = -\ln L(\boldsymbol{\theta}|\boldsymbol{u})p(\boldsymbol{u})$. Hence, $\hat{\boldsymbol{u}}(\boldsymbol{\theta})$ is maximizing the integrand of (11) with respect to the latent variables (for a fixed value of $\boldsymbol{\theta}$). The Laplace approximation of (11) is then, in the notation of Kristensen et al. (2016), given as

$$L^*(\theta) = \sqrt{2\pi}^{\dim u} \det \left\{H(\boldsymbol{\theta})\right\}^{-1/2} \exp\left\{-f\left(\hat{\boldsymbol{u}}(\boldsymbol{\theta}), \boldsymbol{\theta}\right)\right\}, \tag{13}$$

where $\dim u$ denotes the dimension of $\boldsymbol{u}$ and $\det(H)$ is the determinant of the Hessian matrix

$$H(\boldsymbol{\theta}) = \frac{\partial^2}{\partial \boldsymbol{u}^2} f\left(\boldsymbol{u}, \boldsymbol{\theta}\right)\Big|_{\boldsymbol{u}=\hat{\boldsymbol{u}}(\boldsymbol{\theta})}. \tag{14}$$

`TMB` evaluates $H(\boldsymbol{\theta})$ numerically using automatic differentiation (AD), which is both computationally efficient and precise. Note that AD is not the same as taking "finite differences", which is imprecise and computationally costly. Neither is it the same as obtaining an analytical expression for $H(\boldsymbol{\theta})$, followed by numerical evaluation, which would be impractical for a likelihood as complex as ours. Further, `TMB` uses AD for maximizing $\hat{\boldsymbol{u}}(\boldsymbol{\theta})$, and for evaluating the gradient $\partial/\partial\boldsymbol{\theta}\ L^*(\boldsymbol{\theta})$. The latter involves third order derivatives evaluated

[34] In the parlour of the EM-algorithm, see Dempster et al. (1977), (12) is the log-likelihood for the *full data*.

by AD, and allows the use of quasi Newton methods to maximize $L^*(\theta)$.

For the Laplace approximation to be accurate, the function $L(\boldsymbol{\theta}|\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}}) \cdot p(\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})$ must be unimodal as a function of $\boldsymbol{F}$, $\boldsymbol{\alpha}$ and $\boldsymbol{\varepsilon^{pl}}$, and the closer the functional form is to a multivariate Gaussian density, the better the approximation is. (This is because (13) is based on a second order Taylor approximation of $f(\cdot, \boldsymbol{\theta})$.) For this reason, the Laplace approximation is carried out with respect to $\ln \boldsymbol{F}$ (rather than $\boldsymbol{F}$), because according to (7), the prior distribution of $\ln \boldsymbol{F}$ is multivariate Gaussian.

As a byproduct of the Laplace approximation `TMB` returns the maximum a posteriori (MAP) estimates of $\boldsymbol{F}$, $\boldsymbol{\alpha}$ and $\boldsymbol{\varepsilon^{pl}}$. The MAP estimate of $\boldsymbol{u}$ is simply $\hat{\boldsymbol{u}}(\hat{\boldsymbol{\theta}})$, where $\hat{\boldsymbol{\theta}}$ is the maximizer of $L^*(\boldsymbol{\theta})$, i.e. the approximate maximum likelihood estimator. The reason that the Laplace approximation is expected to work well in our setting because the posterior distributions of $\boldsymbol{\alpha}$, $\ln \boldsymbol{F}$ and $\boldsymbol{\varepsilon^{pl}}$ are all close to Gaussian.[35] A framework in which parameters $\boldsymbol{\theta}$ are estimated by maximum likelihood, while a plug-in ($\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$) Bayes estimator is used for the latent variables, often referred to as "empirical Bayes". By inserting into (6) the MAP of α_{ij}, $\ln F_{jt}$, $\varepsilon^{\mathrm{pl}}_{jt}$, and zero for $\varepsilon^{\mathrm{id}}_{ijt}$, we obtain a MAP estimate of $\ln P^*$. However, because S is integrated out of the likelihood prior to applying the Laplace approximation, `TMB` does not return a MAP estimate for S.

`TMB` also returns a large-sample approximation to the covariance matrix of the MLE of $\boldsymbol{\theta} = (\mu_\alpha, \sigma_\alpha, \beta, \sigma_{\mathrm{pl}}, \sigma_{\mathrm{id}}, \phi_0, \sigma_\eta, a_s, b_s, g_s, \sigma_s)'$, and optionally also the covariance matrices of $\boldsymbol{u} = (\boldsymbol{F}, \boldsymbol{\alpha}, \boldsymbol{\varepsilon^{pl}})$. The formula used by TMB to calculate the covariance matrix of $\hat{\boldsymbol{u}}$ is

$$\Sigma_{\hat{\boldsymbol{u}}} = H(\boldsymbol{\theta})^{-1} + \boldsymbol{v}\Sigma_{\hat{\boldsymbol{\theta}}}\boldsymbol{v}', \tag{15}$$

where $\Sigma_{\hat{\boldsymbol{\theta}}} = -\left\{\partial^2/\partial\boldsymbol{\theta}^2\, L^*(\boldsymbol{\theta})\right\}^{-1}$, and $\boldsymbol{v} = \partial/\partial\boldsymbol{\theta}\, \hat{\boldsymbol{u}}(\boldsymbol{\theta})$ is the Jacobian matrix of $\hat{\boldsymbol{u}}(\boldsymbol{\theta})$. Note that $\boldsymbol{v}$ has the interpretation of being the sensitivity of $\hat{\boldsymbol{u}}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$. The first term on the right hand side of (15) is the ordinary large-sample approximation to the covariance matrix of $\hat{\boldsymbol{u}}$, while the second term is the uncertainty in $\hat{\boldsymbol{u}}$ arising from

[35]Partly due to a central limit effect caused by the large amount of data, but also partly due to the Gaussian (prior) distribution placed on the latent variables, such as on η_{jt} in (7).

uncertainty in $\hat{\boldsymbol{\theta}}$. Zheng and Cadigan (2021) give a frequentist interpretation of (15) as the covariance matrix of $\hat{\boldsymbol{u}} - \boldsymbol{u}$ under repeated sampling.

C Monte Carlo study of the parameter estimator

The performance of the point estimates and their standard errors is examined using a Monte Carlo study. For the reader's convenience, the model formulation used for simulation is repeated.

$$\ln P^*_{ijt} = \alpha_{ij} + \beta \ln F_{jt} + (\varepsilon^{\mathrm{pl}}_{jt} + \varepsilon^{\mathrm{id}}_{ijt})$$

$$\ln F_{jt} = \phi_0 + \ln F_{j,t-1} + \eta_{jt}$$

$$W_{j,12s} = F_{j,12s-11} + F_{j,12s-10} + \ldots + F_{j,12s}$$

$$\ln P_{ijt} = \begin{cases} \ln P_{ij,t-1} \text{ if } -S^-_{ijt} \leq (\ln P^*_{ijt} - \ln P_{ij,t-1}) \leq S^+_{ijt} \\ \ln P^*_{ijt} \qquad \text{otherwise} \end{cases}$$

where the stochastic variables α_{ij}, $\varepsilon^{\mathrm{pl}}_{jt}$, $\varepsilon^{\mathrm{id}}_{ijt}$, η_{jt} and S_{ijt} are independent of each other; they are neither auto-correlated nor cross-correlated. S_{ijt} follows a truncated normal distribution and the other four normal distributions. The notation S^+_{ijt} and S^-_{ijt} shows that the thresholds can have different distributions for positive and negative changes. The distributional assumptions are as follows:

$$\alpha_{ij} \sim N(\mu_\alpha, \sigma^2_\alpha)$$

$$\varepsilon^{\mathrm{pl}}_{jt} \sim N(0, \sigma^2_{\mathrm{pl}})$$

$$\varepsilon^{\mathrm{id}}_{ijt} \sim N(0, \sigma^2_{\mathrm{id}})$$

$$\eta_{jt} \sim N(0, \sigma^2_\eta)$$

$$S^{+/-}_{ijt} \sim TN(\mu^s_{ijt}, \sigma^2_{\mathrm{S}}, 0, \infty)$$

μ^s_{ijt} is specified as

$$\mu^s_{ijt} = a_s + b_s D(\mathrm{JanFeb}_t) + g_s D(\mathrm{neg}_{ijt}),$$

where $D(\mathrm{JanFeb}_t)$ and $D(\mathrm{neg}_{ijt})$ are dummy variables for January-February and negative changes, respectively. It implies that the seasonal effect (b_s) is the same for positive and negative price changes, while g_s (> 0) introduces asymmetry such that we observe less frequent price decreases relative to price increases. We set the threshold parameters $a_s = 0.01$, $b_s = -0.02$, and $g_s = 0.20$, while $\ln \sigma_s = -1.80 \Leftrightarrow \sigma_s = 0.165$. The chosen threshold parameters generate the following conditional means and standard deviations for the truncated normal distribution:

	μ_S	σ_s	$E(S\|\cdot)$	$sd(S\|\cdot)$
JanFeb+	$a_s + b_s = -0.01$	0.165	0.13	0.10
JanFeb-	$a_s + b_s + g_s = 0.19$	0.165	0.23	0.14
MarDec+	$a_s = 0.01$	0.165	0.14	0.10
MarDec-	$a_s + g_s = 0.21$	0.165	0.24	0.14

Table 5: Conditional means and standard deviations for the truncated normal distribution. $E(S|\cdot)$ is a short form notation for $E(S|D(\text{JanFeb}), D(\text{neg}))$, and similarly for $sd(S|\cdot)$.

We see that the resulting conditional means $E(S|D(\text{JanFeb}), D(\text{neg}))$ range from 0.13 to 0.24, with corresponding conditional standard deviations between 0.10 and 0.14. Dhyne et al. (2011) - Table 2) — based on a model with normally distributed thresholds — report estimates of the means of the price inaction band (excluding Energy and Services) in the interval [0.203, 0.522] and the corresponding average standard deviations in the interval [0.135, 0.254]. Our simulated thresholds are therefore somewhat lower and less dispersed. This is consistent with the fact that Dhyne et al. (2011) analyze consumer prices (CPI), while our model focuses on producer prices (PPI). As shown by Vermeulen et al. (2012), producer prices adjust more frequently than consumer prices.

The two α parameters, $\mu_\alpha = 4.0$ and $\ln \sigma_\alpha = 1 \Leftrightarrow \sigma_\alpha = 2.718$, are set such that together with the chosen value of $\beta = 0.40$ (see our discussion of β in Section 2) and given the distribution of annual wages per employee, they describe the distribution of actual prices $\ln P_{ijt}$. Based on a two-way crossed random effects model (see (Baltagi et al., 2001)) of the $\ln P$ on $\ln(W/12)$ we obtain hints about the magnitudes of the spread of the two shock parameters ε_{jt}^{pl} and ε_{ijt}^{id}. Choosing them to be equally sized, we set $\ln \sigma_{\text{id}} = \ln \sigma_{\text{pl}} = -2 \Leftrightarrow \sigma_{\text{id}} = \sigma_{\text{pl}} = 0.135$. The constant term of the wage equation, ϕ_0, is set to 0.003, which corresponds to an annual increase of 3.7%, a reasonable value based on the trend of mean wages in the relevant period in Norway. The parameter $\ln \sigma_\nu = -4.00 \Leftrightarrow \sigma_\nu = 0.018$. Together, the chosen parameter values provide a transparent and empirically grounded basis for our simulation exercises, designed to evaluate the appropriateness of the model specification. With the parameter values described above, we simulate 105 plants producing 253 products, mirroring the sample in our empirical analyses.[36] We track these simulated plants over 12 toyears, making 1000 replications and averaging over these to report the moments.

	Parameter		St. error	
	True	MC mean	True	MC mean
Price equation				
μ_α	4.000	4.0612	0.1796	0.1738
$\ln \sigma_\alpha$	1.000	0.9936	0.0439	0.0445
β	0.400	0.3842	0.0124	0.0097
$\ln \sigma_{pl}$	-2.000	-1.9743	0.0111	0.0104
$\ln \sigma_{id}$	-2.000	-2.0080	0.0065	0.0062
Wage evolvement				
ϕ_0	0.003	0.0028	0.0005	0.0004
$\ln \sigma_\eta$	-4.000	-4.7633	0.1046	0.0047
Price thresholds				
a_s	0.010	0.0373	0.0066	0.0075
b_s	-0.020	-0.0061	0.0060	0.0060
g_s	0.200	0.2342	0.0058	0.0060
$\ln \sigma_s$	-1.800	-1.7223	0.0103	0.0131

Table 6: Monte Carlo (MC) evaluation of point estimates and standard errors.
True parameter: The values from which we generated the data. MC mean: The average across 1000 MC simulations of either the point estimate or the st. error as reported by TMB. True st. error: empirical standard deviation of point estimates across the 1000 simulations.

Ninety-eight percent of the replications converged, and the Monte Carlo study took about 9 hours and 20 minutes to complete on a desktop computer. As shown in Table 6, the averages of the point estimates are close to their population (true) counterparts. Moreover, the standard errors reported by TMB generally provide good estimates of the standard deviations of the sampling distributions. A few exceptions are observed for $\ln \sigma_\eta$, which shows discrepancies in both the parameter and standard deviation estimates, and

[36]The distribution of products per plant is the same during the simulations as in the estimation sample.

for a_s, b_s, g_s, and $\ln \sigma_s$, which exhibit small biases. A possible explanation for the bias in the price-threshold parameters is the approximation of the bivariate normal cumulative distribution function (BNCDF). A smaller simulation experiment with a reduced threshold variance, $\sigma_s = 0.049$, indicates that the bias is all but gone. Hence, an improvement might be achieved—albeit at the cost of computational speed—by using a more refined approximation of the BNCDF. Despite the small biases reported in Table 6, the simulation results demonstrate that the estimator used in the main analysis performs well in capturing the key properties of the model.

The importance of latent F versus observed $W/12$

In a final exercise, we examine to what extent using a much simpler approach—based on the observed $W/12$ rather than the latent variable F as defined in (7) (and linked to observed annual wages W through (8)) affects the parameter estimates of our model. We do this by simulating data from the model using the "true" parameter values reported in Table 6, Column 1, thereby generating the series $\ln F_{jt}$ for all plants j and months t. We then aggregate F over twelve months to construct W, and, instead of using $\ln F_{jt}$, we estimate the model using $\ln(W/12)$). In this specification, the latent variable F and the parameters ϕ_0 and σ_η (and therefore the whole "wage evolvement" equation) are therefore removed from the model. The results are reported in Table 7.

Both the point estimates and the standard errors for the parameters in the price equations are very close to the 'True' values and the corresponding 'MC mean' reported in Table 6. Thus, if our primary interest is the pass-through parameter β, it appears that simply "splitting" annual costs evenly across months performs well. However, this approach does affect the threshold parameter a_s which is smaller than the 'true' value (-0.036 vs 0.010, estimated and 'true', respectively). This suggests that using $W/12$ instead of the latent variable F may introduce a small negative bias in the threshold estimates.[37]

[37] The estimated difference is 0.046 (= 0.010 − -0.036) with a standard error of 0.007, yielding an approximate 95% confidence interval of (0.032, 0.060). Thus, the standard error would have to be substantially

	Coef	SE
Price equation		
μ_α	4.068	0.177
$\ln \sigma_\alpha$	1.009	0.044
β	0.408	0.010
$\ln \sigma$	-1.974	0.010
$\ln \sigma$	-2.011	0.006
Price thresholds		
a_s	-0.036	0.007
b_s	-0.006	0.006
g_s	0.231	0.006
$\ln \sigma_s$	-1.715	0.013

Table 7: Monte Carlo results when replacing the latent variable F with observed $W/12$. Notes: Parameter estimates and st. errors (SE), using simulated data generated from the model with true parameters given in Table 6, Column 1.

underestimated for the confidence interval to include zero.

D Figures

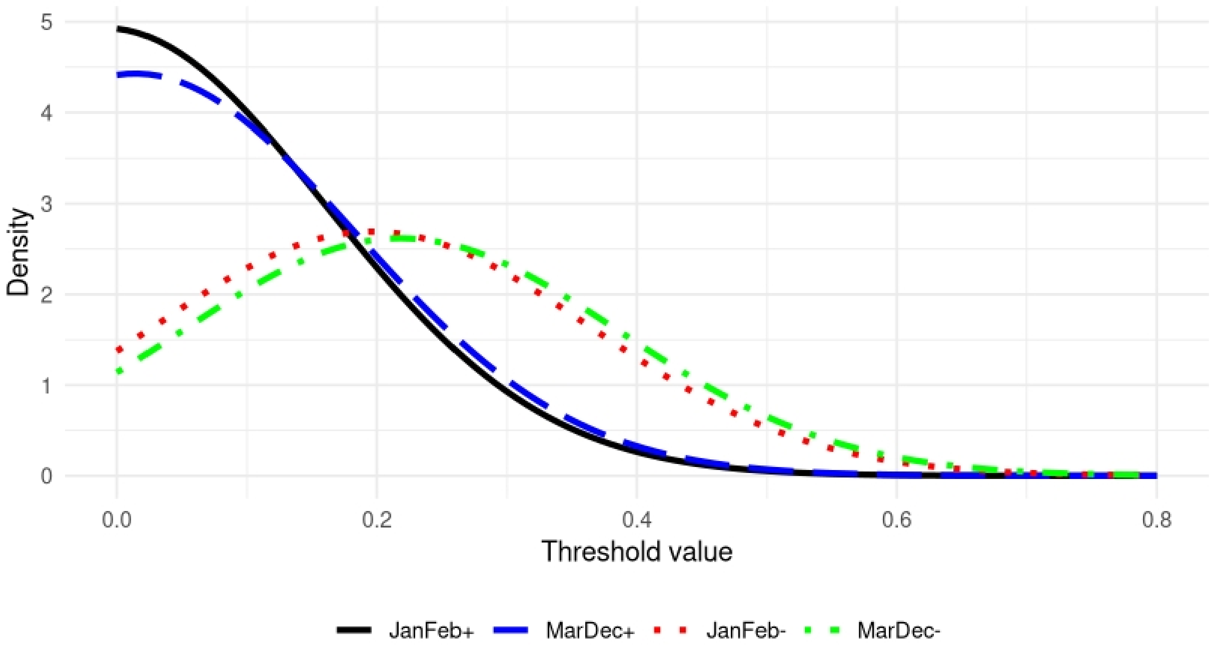

Figure 1: Estimated truncated normal distributions for the thresholds in Model 4.

NHH

NORGES HANDELSHØYSKOLE
Norwegian School of Economics

Helleveien 30
NO-5045 Bergen
Norway

T +47 55 95 90 00
E nhh.postmottak@nhh.no
W www.nhh.no