UPPSALA
UNIVERSITET

Department of Economics

Working Paper 2014:6

Selection Effects in Producer-Price Setting

Mikael Carlsson

Department of Economics
Uppsala University
P.O. Box 513
SE-751 20 Uppsala
Sweden
Fax: +46 18 471 14 78

Working paper 2014:6
August 2014
ISSN 1653-6975

SELECTION EFFECTS IN PRODUCER-PRICE SETTING

MIKAEL CARLSSON

Selection Effects in Producer-Price Setting*

Mikael Carlsson[†]

August 19, 2014

Abstract

We use micro data on product prices linked to information on the firms that set them to test for selection effects (state dependence) in micro-level producer pricing. In contrast to using synthetic data from a canonical menu-cost model, we find very weak, if any, micro-level selection effects when running price change probability regressions on actual data. Moreover, when fitting a model that nests both time- and state-dependent elements (the CalvoPlus model of Nakamura and Steinsson, 2010) to the data, the resulting parameters mimic the standard Calvo (1983) model. Thus, upstream in the supply chain, price-setting is best characterized as time-dependent.

Keywords: Price-setting, Business Cycles, Micro Data.

JEL classifications: D4, E3, L16.

*I am grateful to Nils Gottfries, Oskar Nordström Skans, Andreas Westermarck and seminar participants at Uppsala University for useful comments and discussions. I am also grateful to Jonny Hall for helpful advice. The data used in this paper are confidential but the authors' access is not exclusive. Financial support from the Ragnar Söderberg Foundation is gratefully acknowledged. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank.

[†]Uppsala University, UCLS and Sveriges Riksbank. E-mail: mikael.carlsson@nek.uu.se

1 Introduction

In the canonical workhorse model of applied macroeconomics, the New Keynesian model, nominal frictions are the keystone for generating monetary non-neutrality and a role for monetary policy.¹ A key simplifying assumption in this model is that price-setting is time-dependent (TD). Thus, the pricing decision faced by the firm is only about the magnitude of the price change and not the timing of the change.² However, introducing state dependence (SD) in pricing, i.e. treating the timing (as well as the magnitude) of price changes as a regular profit-maximizing choice, can have a dramatic effect on the degree of monetary non-neutrality; see Caplin and Spulber (1987), Dotsey, King, and Wolman (1999), Golosov and Lucas (2007) and Midrigan (2011). The main driver behind this result is the self-selection mechanism in SD models that mitigates the real effects of money. That is, firms that change price in SD models are those that have the most to gain from it. This increases the effect on the price level from a monetary shock relative to a TD model and reduces the degree of monetary non-neutrality. Moreover, modeling pricing as TD or SD also affects other properties of the model, such as determinacy under a specific policy rule; see Dotsey and King (2005) for a discussion. Thus, whether self-selection by firms into the price-changing group is a key feature of observed firm behavior or not is an important question for macroeconomic analysis and the policy advice derived from it.

In this paper we address the empirical importance of the self-selection mechanism in pricing directly at the micro level. This paper is thus part of a very small, but growing literature that uses quantitative micro data linking prices to marginal cost. One strand of this literature focuses on data downstream in the supply chain that relates retail prices to costs (wholesale/spot prices or replacement cost for the vended product); see e.g. Levy, Dutta, and Bergen (2002), Davis and Hamilton (2004) and Eichenbaum, Jaimovich, and Rebelo (2011). In this paper, and as in Carlsson and Nordström Skans (2012), the focus is instead on price-setting behavior upstream in the supply chain and draws on very detailed annual Swedish data on product producer prices matched to a rich data set containing

¹See Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005)

²In the Taylor (1980) model the timing of price changes is a deterministic function of time, and in the Calvo (1983) model it is stochastic with a fixed probability of changing the price each period. The tractability gain from making the firm's pricing decision only about the magnitude of the price change comes from the reduced dimensionality needed when describing the evolution of the aggregate price level.

information on the activity of the firms that set these prices. Using the firm-level data, we construct a measure of marginal cost (i.e. unit labor cost) consistent with the vast majority of DSGE models in the literature. To our knowledge, this is the first data set where such detailed quantitative price data have been merged with detailed information on firm-level activity for a broad sample (702) of industrial firms.

Departing from the finding of sizeable nominal frictions reported in Carlsson and Nordström Skans (2012), this paper explores to what extent price-setting features sizable selection effects, if at all. Importantly, the focus here is directly on firm behavior and whether or not we observe self-selection on the micro level. This is a necessary condition if self selection will play a role in the degree of monetary non-neutrality. Note, however, that the overall importance of self-selection for monetary non-neutrality is also a question of the measure of marginal firms lying close to the adjustment threshold; see Midrigan (2011).

To impose discipline on the empirical exercise at hand, we first outline and calibrate a baseline SD model to match key moments in the data. The menu-cost model we rely on is along the lines of Nakamura and Steinsson (2008), but allows for fat-tailed idiosyncratic shocks to marginal cost (akin to Midrigan, 2011) in order to better match the micro-data.³ Moreover, the model is calibrated to a monthly frequency, which then allows us to gauge the effect of time aggregation in the annual data. Aggregating the simulated data in the same way as the actual data is aggregated, we find that time aggregation fills out the gap of very small price changes that is otherwise a hallmark of the price-change distribution in SD models. Actually, this type of data filtering takes the menu-cost model a long way in replicating the observed annual price change distribution. Thus, time aggregation is a complementary mechanism for generating small price changes in SD models to the economies of scope suggested by Lach and Tsiddon (2007), Midrigan (2011) and Alvarez and Lippi (2014) or stochastic menu costs as in Caballero and Engel (1999) and Dotsey, King, and Wolman (1999). Intuitively, large positive and negative monthly changes within a year nearly cancel one another out, which generates small overall price movements in the annual data. Also, the time-aggregation mechanism described here should be at work as soon as we leave ticker data and rely on data with intermittent price observations.

³The SD model of Nakamura and Steinsson (2008) builds in turn on work by Barro (1972), Sheshinski and Weiss (1977), Golosov and Lucas (2007) and others.

Next, we analyze the strength of the selection mechanism by running probability models along the routes of Cecchetti (1986), Buckle and Carlson (2000), Loupias and Sevestre (2013) and others. Specifically, we investigate if the absolute value of the change in the firm’s marginal cost affects the probability of a price change and compare the findings from observed data to those from synthetic time-aggregated data generated by the SD model. We find an order of magnitude smaller contemporaneous effect from the absolute value of the change in the firm’s marginal cost on the probability of a price change than expected if the SD model was generating the data. Neither do we find any effect from the lagged absolute value of the change in the firm’s marginal cost, which in a SD model would affect the price change probability through pent-up adjustment incentives. Moreover, when considering measurement issues pertaining to the classification of small price changes in the data, the small contemporaneous effect we find seems to be the result of upward bias.

To structurally quantify the regression results we also fit a price-setting model that nests both TD and SD elements to the data (i.e. a fat-tailed shocks version of the Calvo-Plus model outlined in Nakamura and Steinsson, 2010), which can generate an arbitrary degree of selection effects in the simulated micro data from the model. Importantly, the procedure to fit the model parameters can be constructed to be unaffected by the measurement issues that may bias the regression results. When choosing parameters so that the model matches empirical moments as closely as possible, the parameters are driven very close to a purely TD standard Calvo (1983) model. This thus again implies that the selection effects are not important feature of the data.

Thus, overall, timing adjustments of price changes to marginal-cost developments do not seem to be an important feature of observed price-setting behavior of goods-producing firms. As a result, a corollary to this finding is that a TD model seems to provide a reasonable description of the price-setting behavior in our data. Note that our data are drawn from firms upstream in the supply chain. Eichenbaum, Jaimovich, and Rebelo (2011) also links a measure of marginal cost, i.e. the replacement cost of the vended product, to the price set in data drawn from a large US food and drug retailer and documents a high degree of selection effects in pricing.⁴ This indicates that there seems

⁴Especially when considering reference prices (and costs) - i.e. when abstracting from high frequency variation such as sales commonly observed in consumer prices. As noted by Nakamura and Steinsson (2008), sales seem to be uncommon in producer price data.

to be considerable differences in pricing behavior along the supply chain. This is perhaps not surprising given differences in market conditions. In our data, a considerable share of the trades are likely to be business-to-business, and production processes might be need to be altered by both the buyer and the seller when a business relationship is formed. This, in turn, gives rise to hold-up problems and incentives to form long-term relationships. Moreover, in such situations, short-term sales are less likely to be an optimal strategy for influencing demand.

Another important point, when thinking about the results found here, is that in the canonical New Keynesian model the TD price-setting frictions are usually added high up in the supply chain (intermediate goods sector), whereas down stream sectors (retail sector) are, for convenience, modeled as frictionless; see e.g. Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005). Thus, this class of models does not need price-setting frictions on all levels of the supply chain in order to generate significant monetary non-neutrality. This then implies that frictions found in the down stream sectors can only add to monetary non-neutrality and given the results presented here, they are not instrumental for the existence of sizable monetary non-neutrality.

This paper is organized as follows: Section 2 presents the data, section 3 outlines the SD model used as a benchmark, section 4 presents our results and, finally, section 5 concludes the paper.

2 Data

In this paper we rely on the same data set as in Carlsson and Nordström Skans (2012), i.e. quantitative price data on the product level that have been merged with information on the producing firm’s production level, inputs and costs for a broad sample of manufacturing firms. This data set combines information on detailed product-prices drawn from the Swedish IVP (“Industrins Varuproduktion”) survey with information on plant-level activity from the IS (“Industristatistiken”) survey.

The IVP micro data provides annual information on prices and quantities of products for all Swedish industrial plants with at least 10 (20) employees for the years 1990 – 1996 (1997 – 2002) and a sample of smaller plants. The product classification is at the 8/9-digit level of the Harmonized System (HS) for the years 1990 – 1995 and the Combined

Nomenclature (CN) for the years 1996 – 2002. The data allow us to follow the same product (or at least a very closely defined group of products) over time. The codes are fairly exact; an example of a product code is 84181010 for “*A combined freezer and cooler with separate exterior doors with a volume exceeding 340 liters intended for use in civilian aircrafts*”. The (unit) price for each product code is calculated by dividing the firms’ yearly reported value for the product code with the accompanying volume (in terms of the relevant measure, e.g. the number of products, cubic meters, metric tons, etc.). The data are thus based on actual transaction prices and not list prices.

A key novelty is that the price data can be matched to data on activity for the individual plant. The IS survey contains annual information on inputs and output for all Swedish industrial plants with 10 employees or more and a sample of smaller plants. We only use plants that are also a firm since pricing essentially is a firm-level and not a plant-level decision and since there is some scope for transactions between plants within a firm for tax reasons. In addition, we limit the analysis to firms that are in operation throughout the sample period since we want to identify normal behavior.

Following Rotemberg and Woodford (1999), Carlsson and Nordström Skans (2012) and others, we rely on unit labor cost as a measure of marginal cost.⁵ To construct unit labor cost we use the IS survey data on the firms’ wage bill divided by real output, where the latter variable is obtained by deflating nominal output from the IS survey (the value of total sales) using a firm-specific producer price index.⁶

Since the raw price data involve a few very large swings we apply a cleaning procedure in which we split the individual price series and give them a new unique plant-price identifier whenever a large change in the growth rate appears in the raw data. The cut-off levels are given by the 1.5 and 98.5 centiles of the full raw data distribution. We also remove firms that are subject to large swings in the observed marginal cost. As with prices, we use the full distribution of log changes in unit labor cost across all firms for which this variable can be computed and remove firms with growth rates outside the [1.5, 98.5] centiles in any one year of the sample period.

⁵As discussed in Carlsson and Nordström Skans (2012) this is a good measure of marginal cost under the assumption that firms are cost minimizers, wage takers and face a production technology that is approximately a Cobb-Douglas.

⁶The price index is constructed as a chained index with Paasche links combining the plant-specific unit prices described above and the most detailed product/producer-price indices available. The product/producer-price indices are used if the 8/9-digit unit value data are not available due to missing data, changes in the firm’s product portfolio, or when there are large swings (over the 1.5/98.5 centiles).

Figure 1: Histograms of data. The left-hand panel describes the distribution of log price changes across 13,772 observations (for 1,610 different products across 702 firms). The right-hand panel describes the distribution of log unit labor cost changes across 8,424 observations (for 702 firms). Bin size 0.01.

When merging data sets, we are left with 17,282 price observations (with a minimum spell length of two periods) across 1,610 unique product codes, 3,510 unique product/firm identities and 702 firms (as in Carlsson and Nordström Skans, 2012). These industrial firms are mainly medium to small firms with an average of 65 employees. See also Appendix A for more details on the data construction. There we also present evidence on the robustness of the results to more generous cut-off levels.

In Figure 1, we plot the final data distribution of log price changes (for the 8/9-digit unit price data). All in all, this comprises 13,772 price-change observations. Each bin represents a log difference of 0.01. Note that since these prices are calculated from reported values and volumes of sold products, there might be small rounding errors in the data. As can be seen in Figure 1, however, there is a substantial spike for the bin centered around zero. In fact, 13.6 percent of the price-change observations are confined within the ± 0.5 percent interval.

The observation that a substantial fraction of price spells remain fixed across years is well in line with existing survey evidence. When surveying 626 Swedish firms in 2002, Apel, Friberg, and Hallsten (2005) found that about 70 percent of the firms adjust their price once a year or less. Moreover, for the approximately 15,000 European firms surveyed in the Eurosystem Wage Dynamics Network, Druant *et. al.* (2012) reports that about half of the firms on average change their price once a year or less. In a wider perspective it is interesting to note that both studies report that manufacturing (upstream) firms seem to change prices less frequently than the economy-wide average.

In the right-hand panel of Figure 1, we plot the distribution of log changes in unit labor cost for the 702 firms (all in all 8,424 observations). As can be seen in the figure, there is no corresponding spike at the zero unit labor cost change bin.⁷ The shapes of the two distributions is thus indicative of nominal price rigidities in the sense that the spike in the price change distribution is not matched with a spike in the marginal-cost change distribution. Relying on the same data set and measurement, as employed here, Carlsson and Nordström Skans (2012) established that the marginal cost measure (unit labor cost) is an important driver of the magnitude of price changes and report empirical evidence in support of a nominal frictions interpretation of the data. Focusing on idiosyncratic variation for identification (i.e. including time fixed effects in all specifications), Carlsson and Nordström Skans (2012) first reports an instantaneous (within-year) pass-through of marginal cost to the price of about one-third, which speaks against a frictionless interpretation of the data. Secondly, when conditioning on price changers only, they found evidence that firms consider both current and future expected marginal cost when setting today's price (with the sum of coefficients not significantly different from unity). This is important since future marginal cost developments only matter for today's pricing decision in the presence of impediments to continuous and costless price adjustments as in SD or TD models. However, since SD or menu-cost models rely on a fixed cost to generate a mass point of zero adjustment, they also generate a region of inaction around the zero adjustment point. Thus, from the shape of the price-change distribution it may seem like a standard SD model could be taken out of the picture already at this point, but as we will see this is not the case when we explicitly consider the underlying time aggregation

⁷In fact, there are only three observations with exactly zero growth in marginal cost, whereas the corresponding number for price changes is 529.

of the annual data.⁸ A final important result from Carlsson and Nordström Skans (2012) is that the OLS and IV estimate of the pass-through of price to marginal cost is very similar.⁹ Thus, there does not seem to be any important endogenous variation in marginal cost, suggesting an approximately flat firm-level marginal cost curve. Also, classical measurement errors in the marginal-cost measure seem to be of minor importance since this would also drive a wedge between the OLS and the IV results.

3 A Baseline Menu-Cost Model

To obtain a benchmark for what micro-level selection effects to expect in the empirical work if the data were generated from a SD model, we rely on a partial equilibrium menu-cost model along the lines of Nakamura and Steinsson (2008).¹⁰ Moreover, we explicitly consider the effects of the time aggregation of our data by calibrating and simulating an underlying monthly menu-cost model from which we generate synthetic annual data by time aggregating the synthetic monthly data in the same way as our annual data are constructed.

3.1 The Menu-Cost Model

Let firm j 's product demand at time t , Y_{jt} , be given by

$$Y_{jt} = Cp_{jt}^{-\theta}, \quad (1)$$

where C is (constant) aggregate demand determining the size of the market, $p_{jt} = P_{jt}/P_t$ is the relative price of firm j and $\theta(> 1)$ is the (negative) of the price elasticity of demand. To change the nominal price, P_{jt} , κ units of labor is needed. Following Nakamura and

⁸Other routes to generate small price changes in SD models are economies of scope as suggested by Lach and Tsiddon (2007), Midrigan (2011) and Alvarez and Lippi (2014) or stochastic menu costs as in Caballero and Engel (1999) and Dotsey, King, and Wolman (1999).

⁹Beside internal instruments (i.e. lags), Carlsson and Nordström Skans (2012) also exploits that they have access to detailed information on all employees within each firm in the private sector. Relying on this information, they construct an instrument based on the local-market valuation of the (lagged) skill composition of the firm normalized by the lagged production level.

¹⁰Which in turn builds on work by Barro (1972), Sheshinski and Weiss (1977), Golosov and Lucas (2007).

Steinsson (2008) we assume that the (constant) real aggregate wage is given by¹¹

$$W_t/P_t = \frac{\theta - 1}{\theta}. \quad (2)$$

Assuming a constant returns to scale technology, the firm's real profit can be written as

$$\Pi_{jt} = Cp_{jt}^{-\theta} (p_{jt} - mc_{jt}) - \kappa \left(\frac{\theta - 1}{\theta} \right) I_{jt}, \quad (3)$$

where mc_{jt} is the real marginal cost of firm j , and I_{jt} is an indicator that takes the value one if the nominal price is changed, i.e. $P_{jt} \neq P_{jt-1}$, and zero otherwise. The constant returns assumption is consistent with the finding of an essentially flat firm-level marginal-cost schedule presented by Carlsson and Nordström Skans (2012). Assuming that firm-level marginal-cost is independent from any decisions taken by the firm that affects the scale of production also motivates modeling marginal cost as an exogenous process. Here, the log of real marginal cost follows an AR(1) process

$$\log mc_{jt} = \lambda + \rho \log mc_{jt-1} + \epsilon_{jt}, \quad (4)$$

where $\lambda = (1 - \rho) \log((\theta - 1)/\theta)$ so that the expectation of long-run real marginal cost converges to the real wage. Moreover, $\epsilon_{jt} \sim \text{Laplace}(0, \sigma_\epsilon/\sqrt{2})$, implying a standard deviation of ϵ_{jt} equal to σ_ϵ . The assumption of a Laplace distribution is motivated by the non-normal shape of the observed annual marginal cost change distribution (when controlling for time dummies the kurtosis (skewness) coefficient equals 3.95 (0.01) and a standard test (D'Agostino, Belanger and D'Agostino, 1990) rejects the null of normality on the one-percent level due to the relatively high kurtosis). This assumption is also in line with the fat-tails assumption of Midrigan (2011). The log of the price level drifts with the rate μ ¹²

$$\log P_t = \mu + \log P_{t-1}. \quad (5)$$

¹¹Following Nakamura and Steinsson (2008) we make a flex-price approximation and normalize aggregate productivity. In the linear (in labor) technology framework of Nakamura and Steinsson (2008) this would amount to setting aggregate productivity to unity.

¹²Nakamura and Steinsson (2008) models the log of the price level to follow a random walk with drift. Here, the approach in the partial equilibrium setting is to keep all aggregate variables at the steady state and use time dummies when motivated to compute empirical moments to match to the model. Note also that, as documented by Carlsson and Nordström Skans (2012), idiosyncratic variation strongly dominates any common variation in the data we use and adding an i.i.d. normally distributed shock to (5) calibrated to match the monthly PPI series does not change the results to any noticeable degree.

Assuming that the firm discounts profit streams at a constant rate β and denoting the relative price the firm enters the period with as $p_{jt}^- = P_{jt-1}/P_t$, the value function of firm j can be written as

$$V(p_{jt}^-, mc_{jt}) = \max_{P_{jt}} [\Pi_{jt} + \beta E_t V(p_{jt+1}^-, mc_{jt+1})], \quad (6)$$

where E_t is the expectations operator. Following Nakamura and Steinsson (2008) we solve this problem by value function iterations on a grid and using the method of Tauchen (1986) to approximate the mc_{jt} process.¹³

3.2 Monthly Calibration

To calibrate the model, we first estimate the drift parameter of the inflation process to (μ) to 0.00138 using monthly data on the Swedish industrial producer-price index for the period 1990:1 to 2002:12. This implies an annualized average inflation rate of 1.7 percent, which is very close to the annual mean price change in the data (1.8 percent). We set $\beta = 0.96^{1/12}$ to generate an annualized real interest rate of about 4 percent. We set $\theta = 3$ which is in line with the firm-level estimate for the Swedish manufacturing sector reported in Carlsson, Messina, and Nordström Skans (2013) when using the instrumental variable approach outlined in Foster, Haltiwanger, and Syverson (2008).

Table 1: Menu-Cost Model Calibration

	Parameter	Value
μ	Inflation Drift	0.00138
β	Discounting	$0.96^{1/12}$
θ	Price elasticity of demand	3
C	Market size	1
ρ	Real marginal cost persistence	0.921
σ_ϵ	S.D. real marginal cost shock	0.0676
$\frac{\kappa(\theta-1)}{\theta}$	Menu Cost	0.0791

To calibrate the remaining parameters, we first normalize C to unity and then set ρ , σ_ϵ and κ so as to match the annual data in terms of (i) the persistence of log real marginal cost estimated in Carlsson and Nordström Skans (2012) (0.542), (ii) the standard

¹³Since the model presented here is just a slightly rewritten version of the model in Nakamura and Steinsson (2008) we rely heavily on their MATLAB code available at <http://www.columbia.edu/~js3204/papers/MenuCostModelCode.zip>.

deviation of the log real marginal cost change distribution (0.145) and (iii) the size of the zero bin in the log price change distribution (0.136). The statistics for real marginal cost variables derived from the unit labor cost data controls for time fixed effects.¹⁴ This procedure removes any aggregate, or common, factors (including deflating the nominal data).

As noted above, the prices are calculated from reported values and volumes of sold products. Since, e.g., survey respondents are asked to state the value of sold products in thousands of SEK, there will be rounding errors in calculated prices and thus small erroneous price changes in the data.^{15,16} In contrast, there are no measurement errors in the synthetic data from the model. This difference motivates calibrating the model to match the zero bin rather than to the share of observation that are exactly zero in the data. That is, as long as any measurement error is small enough to be confined within the zero bin, misclassification should not matter for the moment-matching exercise. Also, judging from the continuous shape of the log price change distribution on both sides surrounding the zero bin, there is no reason to believe that a wider band than the zero bin should be warranted.

Finally, to match annual statistics, we time-aggregate the monthly data using monthly output weights consistently with the annual data we observe. The annual unit price of firm j is constructed as

$$\begin{aligned} P_{jt} &= \frac{\text{Annual Sales}_{jt}}{\text{Annual Volume}_{jt}} = \frac{\sum_m P_{jt}^m Y_{jt}^m}{\sum_m Y_{jt}^m} = \\ &= P_{jt}^1 \frac{Y_{jt}^1}{\sum_m Y_{jt}^m} + \dots + P_{jt}^{12} \frac{Y_{jt}^{12}}{\sum_m Y_{jt}^m}, \end{aligned} \tag{7}$$

¹⁴The estimate of the annual persistence of log real marginal cost in Carlsson and Nordström Skans (2012) actually controls for time interacted by two-digit sector code (NACE). Using this procedure for the standard deviation of the log real marginal cost change distribution yields a very similar estimate to what is used here (0.142 vs. 0.145).

¹⁵Note that the median value of sold products across product codes for the firms in our sample is SEK 6.1 million.

¹⁶Changes in the composition of buyers who pay different prices are another reason for small measurement errors when computing prices by dividing value with volume. Although common in retail prices, see Eichenbaum, Jaimovich, Rebelo, and Smith (2014), some of the price-setting practices in that sector, like discount coupons, two for one offers, and so on, are less likely to be prevalent in producer price setting. Also, Nakamura and Steinsson (2008) notes that sales seem to be uncommon in producer price data.

where m denotes month. Similarly we can write

$$\begin{aligned}
ULC_{jt} &= \frac{\text{Annual Wage Bill}_{jt}}{\text{Annual Volume}_{jt}} = \frac{\sum_m W_{jt}^m L_{jt}^m}{\sum_m Y_{jt}^m} = \\
&= \frac{W_{jt}^1 L_{jt}^1}{Y_{jt}^1} \frac{Y_{jt}^1}{\sum_m Y_{jt}^m} + \dots + \frac{W_{jt}^{12} L_{jt}^{12}}{Y_{jt}^{12}} \frac{Y_{jt}^{12}}{\sum_m Y_{jt}^m} = \\
&= ULC_t^1 \frac{Y_t^1}{\sum_m Y_t^m} + \dots + ULC_t^{12} \frac{Y_t^{12}}{\sum_m Y_t^m},
\end{aligned} \tag{8}$$

which motivates the use of monthly output weights.

The full calibration is presented in Table 1 and implies that the model needs a sizable menu cost, about 23 percent of the average monthly real gross profits, in order to match annual moments.¹⁷

3.3 Simulation Results

In Figure 2 we plot the monthly log price/marginal cost change distributions for 100,000 simulated monthly observations. For clarity we have omitted the spike at zero which contains 92 percent of the observations. Here we see that the high menu cost generates the usual price change distribution with no mass in a region around zero price adjustment.

In Figure 3 we plot the observed and the simulated annual data from the model, focusing on the interval $[-0.5, 0.5]$ log points. A first observation is that the log marginal cost change distribution is well replicated from the simulation. In terms of the similarity of the dispersion of the distributions this is no big victory since the standard deviation of the log real marginal cost change distribution is a target moment when fitting the model combined with a constant inflation rate in the model. Importantly, however, the kurtosis of the actual data (3.82) is not far from that of the simulated distribution (3.24). Turning to the log price change distribution, a key observation is that we find no regions of inaction in the time aggregated synthetic data, although we do see some difference in the observed log price change data and the time-aggregated synthetic data in that there is a lack of mass around the spike at the zero bin. Moreover, the simulated distribution is not dispersed enough, the observed/simulated standard deviations are 0.19 vs. 0.13 and the kurtosis of the actual data (8.62) is much higher than that of the simulated distribution (3.39). However, time aggregation gives a lot of mileage in replicating the observed log

¹⁷That is the ratio of $\kappa(\theta - 1)/\theta$ and the average of $Cp_{jt}^{-\theta}(p_{jt} - mc_{jt})$ in the simulated monthly data.

Figure 2: Histograms of simulated monthly data from the menu-cost model. The log price change distribution (left panel) omits the zero bin.

Figure 3: Histograms of actual (top panel) and simulated data from the menu-cost model (bottom panel). Bin size 0.01.

price change distribution with a stylized menu-cost model and provides a complementary mechanism for generating small price changes in SD models to the economies of scope suggested by Midrigan (2011) or stochastic menu costs as in Dotsey, King, and Wolman (1999). Also, the time-aggregation mechanism described here should be at work as soon as we leave ticker data and rely on a time average of prices or in any setting where big positive and negative observations can almost cancel each other out as in data with intermittent price observations.

4 Results

In this section we compare the empirical strength of the selection effects in the micro data to what is expected from the Menu-Cost model, outlined above, using regression methods. We also discuss whether these results can be interpreted as true selection effects and evaluate potential bias. In a final step, we then try to structurally quantify the regression results in a model that can generate an arbitrary degree of selection effects in the simulated data (i.e. the CalvoPlus model of Nakamura and Steinsson, 2010).

4.1 Probability Regressions

To compare the relative strength of the selection mechanism in the Menu-Cost model vs. the data, we run probability regressions inspired by the work of Cecchetti (1986) and later contributions by e.g. Buckle and Carlson (2000) and Loupias and Sevestre (2013). We first define an indicator for price changes outside the zero bin as

$$I_{gt}^{OZB} = \begin{cases} 1 & \text{if } (|d \ln P_{g,t}| > 0.005) \\ 0 & \text{otherwise} \end{cases}, \quad (9)$$

where $P_{g,t}$ denote the price of good g (produced by firm j) at time t . Next, we regress the absolute value of the change in (log) marginal cost ($|d \ln MC_{j,t}|$) on this indicator, i.e.

$$I_{gt}^{OZB} = \gamma_0 + \gamma_1 |d \ln MC_{j,t}| + \eta_{gt}, \quad (10)$$

where γ_0 and γ_1 are coefficients to be estimated and η_{gt} is a goods-specific error term. That is we run a linear probability model to try to determine whether or not movements in the forcing variable (i.e. marginal costs) have an impact on the price change probability,

or in other words, the timing of the price change. To account for the fact that $|d \ln MC_{j,t}|$ varies on the firm level and not the goods level we correct the standard errors by clustering on the firm level, which handles any type of error-term dependence within the firm over time.

Looking at a small band around zero (instead of the zero point) in the price change distribution is very useful when relying on annual data since it increases the variation in the dependent variable and also renders potential misclassification of small price changes a non-issue for the results when comparing the model to the data. Note, however, that this estimate is likely to be an upward-biased estimate of the true selection effects, since absent any such effects we are still likely to obtain a positive estimate. This is because even in the purely TD model small price changes (within the band) are associated with small marginal cost changes.¹⁸ Here, the main focus is to evaluate the structural model with respect to fitting data moments and for this purpose this bias does not matter since it should also be captured by the model. Below, however, we will try to evaluate the size of this potential bias in the regression model.

In Table 2, we present summary statistics of the data used in the probability regressions. Note that the mean of I_{gt}^{OZB} (0.864) reflects that 13.6 percent of the observations

Table 2: Summary Statistics of Regression Data

Variable	Obs	Mean	Std. Dev.	Min	Max
I_{gt}^{OZB}	13,772	0.864	0.343	0	1
$ d \ln MC_{jt} $	13,772	0.105	0.091	0	0.521

Note: $|d \ln MC_{jt}|$ is weighted as in the regressions.

are contained in the zero bin in the log price change distribution of Figure 1. We also see that there is a sizable variation in $|d \ln MC_{jt}|$ (s.d. of 0.091) as also reflected in the log unit labor cost change distribution of Figure 1.

In the first column of the top panel of Table 3 we present the results from running the linear probability model as outlined above. The estimated marginal effect is 0.13 (s.e. 0.05) and statistically significant on the five-percent level. Thus, taking the estimate at face value and disregarding any biases, there do seem to be selection effects in the

¹⁸Or, in other words, if we erroneously redefine observations in the dependent (dummy) variable to zero that at the same time have values on the independent variable that are below its mean, the estimate of the slope parameter from the probability model will be upward-biased.

sense that the timing of the pricing decision is state-dependent. However, in an economic sense, the effect is very small; a standard deviation change in $|d \ln MC_{jt}|$ implies only a 1.2 percent higher probability of the firm changing price.

Table 3: Estimation and Simulation Results

	(1)	(2)	(3)	(4)	(5)
Estimator	OLS	OLS	OLS	Probit	Logit
Time Dummies:	No	No	Yes	No	No
<i>Data</i>					
$ d \ln MC_{jt} $	0.129* (0.053)	0.114* (0.053)	0.100* (0.051)	0.118* (0.057)	0.120* (0.058)
$ d \ln MC_{jt-1} $		-0.014 (0.072)	-0.032 (0.071)	-0.014 (0.072)	-0.0143 (0.072)
<i>Simulation - Menu-Cost Model</i>					
$ d \ln MC_{jt} $	1.076 [0.031]	1.067 [0.033]	1.067 [0.033]	1.364 [0.044]	1.221 [0.040]
$ d \ln MC_{jt-1} $		0.308 [0.035]	0.308 [0.035]	0.262 [0.033]	0.247 [0.029]

Notes: Dependent variable takes on a value of one if the price change is outside the zero bin and zero otherwise. Columns (4) and (5) present marginal effects evaluated at the mean. Data panel: Superscripts * and ** denote estimates significantly different from zero at the five/one-percent level. Robust standard error clustered on the firm level is inside the parenthesis. The number of observations is 13,772 (12,292 when also including a lag in the model). Simulation panel: The coefficient denotes the average across 200 panel simulations. The standard deviation of the point estimate across 200 panels is inside the square bracket.

This should be compared to the results from doing the same exercise on simulated and time-aggregated data from the Menu-Cost model. Here, we use the monthly Menu-Cost model to generate panels of simulated, time-aggregated annual data consisting of 3,510 price identities (as in the data) observed for five years (the average number of observations per price identity is 4.92 years in the data). The average estimate of the linear probability model across 200 simulated panels is presented in the first column in the bottom panel of Table 3 together with the standard deviation of the point estimate across all repetitions. As can be seen from the table the point estimate does not move much across simulations and the mean, 1.08, is about eight times larger than found in actual data, implying that a standard deviation increase in $|d \ln MC_{f,t}|$ should increase the probability of price adjustment by 9.8 percent.

In column 2 of Table 3 we also include lagged changes in marginal cost, i.e. $|d \ln MC_{jt-1}|$. In a SD model we would also expect lagged changes to matter due to pent-up adjustment incentives. As can be seen in the column 2 of the bottom panel of Table 3 this prediction is confirmed in the simulated and time-aggregated data with a mean point estimate of 0.31 (s.d. of 0.03) on $|d \ln MC_{jt-1}|$. However, we do not see this effect in the observed data. The point estimate is very close to zero -0.01 (s.e. 0.07) and naturally statistically and economically insignificant.

Columns 3-5 show that the results are robust to including time dummies or using Probit and Logit estimators instead of the linear probability model.¹⁹ Thus, across models, we get the same message that the timing adjustments of price changes in response to marginal-cost developments do not seem to be an important feature of observed price-setting behavior of producing firms.

4.2 Structural Evaluation - The CalvoPlus Model

As noted above, the Menu-Cost model generates selection effects that are much too strong. In order to structurally quantify the selection effects implied by the regression results above, we fit a price-setting model that nests TD and SD elements and thus can generate an arbitrary degree of selection effects. To this end we use the CalvoPlus model outlined in Nakamura and Steinsson (2010). As compared with the menu cost model outlined in section 3, the firms now get an opportunity with probability $(1 - \alpha)$ to change price at a low cost κ_L , and to a high cost κ_H otherwise. Thus, this model nests the standard Calvo (1983) model with $\kappa_L = 0$ and $\kappa_H \rightarrow \infty$, as well as the baseline menu cost model presented above with $\alpha = 1$ (or 0) or $\kappa_L = \kappa_H$.

The firm's real profit in the CalvoPlus economy can be written as

$$\Pi_{jt}^{CP} = C p_{jt}^{-\theta} (p_{jt} - mc_{jt}) - (\kappa^L (1 - I_{jt}^H) + \kappa^H I_{jt}^H) \left(\frac{\theta - 1}{\theta} \right) I_{jt}, \quad (11)$$

where I_{jt}^{High} is an indicator that takes on the value one if the the firm faces the high menu

¹⁹The results are also robust to including firm-level fixed effects. Thus, heterogeneity in average price-change probabilities across firms does not seem to be a big issue in our data.

cost and zero otherwise. The value function can be written as,

$$V^{CP}(p_{jt}^-, mc_{jt}, I_{jt}^H) = \max_{P_{jt}} [\Pi_{jt}^{CP} + \beta E_t V^{CP}(p_{jt+1}^-, mc_{jt+1}, I_{jt+1}^H)], \quad (12)$$

where

$$I_{jt+1}^H \sim \text{Bernoulli}(\alpha), \quad (13)$$

and subject to the processes (5) and (4) above.

To fit this model, we again set $\mu = 0.00138$, $\beta = 0.96^{1/12}$, $\theta = 3$ and normalize C to unity. To keep computations feasible we set ρ and σ_ϵ to the same values as for the menu-cost model. The remaining parameters, κ_H , κ_L and α are set so as to minimize the criterion function $\mathbf{M}'\mathbf{M}$ where

$$\mathbf{M} = \begin{bmatrix} (\bar{I}_{Model}^{IZB} - \bar{I}_{Data}^{IZB})/\sigma(\bar{I}_{Data}^{IZB}) \\ (\gamma_{1,Model} - \gamma_{1,Data})/\sigma(\gamma_{1,Data}) \\ (\gamma_{2,Model} - \gamma_{2,Data})/\sigma(\gamma_{2,Data}) \end{bmatrix}, \quad (14)$$

and $\bar{I}^{IZB}$ is the average of $1 - I_{gt}^{OZB}$ and $\gamma_{1,Data}$ and $\gamma_{2,Data}$ denote the coefficients on contemporaneous and lagged $|d \ln MC_{jt}|$, respectively, presented in column 2 of Table 3.²⁰ Finally, σ denotes the standard errors of the observed data moments.²¹ The resulting parameter values, as well as observed and synthetic data moments, for the CalvoPlus model are presented in Table 4. The data wants a menu-cost setup that is in line with the standard Calvo (1983) model with a very high menu cost in the high cost state (about 14 months of average monthly real gross profits) and a very low menu cost in the low cost state (about 22 minutes of average real gross profits for a continuously operating firm). Thus, this exercise speaks against any important selection effects in the data. Moreover, the data wants a Calvo parameter, $\alpha = 0.89$, that is not too far from estimates from macro-data studies. Adolfson, Laséen, Lindé, and Villani (2008) present a quarterly estimate of α of 0.84 using Swedish data, which translates into a monthly Calvo parameter of 0.94. Moreover, Carlsson and Nordström Skans (2012) presents estimates of 0.562 (s.e. of 0.165) on current marginal cost and 0.364 (s.e. of 0.154) on expected future marginal

²⁰Note that the Menu-Cost model could be calibrated to exactly match the data moments used for that model. Thus, any sensible weighting of the moments would return the same parameters.

²¹To find the minimum of the weighted squared deviations we use a combination of a global minimization method (the ga algorithm in MatLab), to rule out local minimums, and a simplex method (fminsearch in MatLab). To make computations feasible, the number of grid points for the state space as well as the number of simulated panels of firms is gradually increased in this process.

Table 4: CalvoPlus Model Calibration

<i>Monthly Calibration</i>			
Parameter		Value	
μ	Inflation drift	0.00138	
β	Discounting	$0.96^{1/12}$	
θ	Price elasticity of demand	3	
C	Market size	1	
ρ	Real marginal cost persistence	0.921	
σ_ϵ	S.D. real marginal cost shock	0.0676	
$\frac{\kappa_H(\theta-1)}{\theta}$	Menu cost (High State)	4.733	
$\frac{\kappa_L(\theta-1)}{\theta}$	Menu cost (Low State)	0.000153	
α	Calvo probability	0.892	
<i>Annual Moments Match</i>			
Moment		Model	Data (S.E.)
Persistence of log real marginal cost		0.544	0.542 (0.042)
S.D. log real marginal cost change distribution		0.143	0.145 (0.002)
Price spike $\bar{I}^{IZB}$		0.135	0.136 (0.008)
Parameter $ d \ln MC_{jt} $		0.172	0.114 (0.053)
Parameter $ d \ln MC_{jt-1} $		0.121	-0.014 (0.072)

Note: Robust standard error clustered on the firm-level within parenthesis in the moments-match panel.

cost when estimating the first-order condition for pricing in the standard Calvo (1983) model on the same data as used in this paper. Interestingly, solving for these coefficients using the first-order condition from the Calvo (1983) model and setting $\alpha = 0.89$ and $\beta = 0.96^{1/12}$ yields expected coefficients of 0.763 on current marginal cost and 0.181 on expected future marginal cost, which is well within the 95-percent confidence interval of the reduced form estimates.²²

In the bottom panel of Table 4 the model moments are compared to their targets in the annual observed data (with standard errors clustered on the firm level). Although the model is not able to exactly match the targets, it does a good job when considering the confidence bands for the observed moments and notably so when it comes to replicating the regression estimates as compared to the coefficients obtained from the canonical Menu-Cost model. Next, in Figure 4, we plot the implied annual log price/marginal change distributions and compare them to both the observed data and the simulated data

²²These coefficients are given by $(1 - \alpha\beta) \cdot \sum_{m=0}^{11} (\alpha\beta)^m$ and $(1 - \alpha\beta) \cdot \sum_{m=12}^{23} (\alpha\beta)^m$, respectively (see, e.g., equation (8) in Carlsson and Nordström Skans, 2012).

Figure 4: Histograms of actual data (top panel), simulated data from the menu-cost model (middle panel) and simulated data from the CalvoPlus model. Bin size 0.01.

from the Menu-Cost model. As compared to the dispersion generated by the Menu-Cost model (s.d. of 0.13), the dispersion of the simulated log price-change distribution (s.d. of 0.08) is actually further away from the observed dispersion (s.d. of 0.19). However, what is clear from the figure is that the CalvoPlus model is better at capturing the high kurtosis observed in the data (8.62) and the overall shape of the log price change distribution. The kurtosis of the log price change distribution of the CalvoPlus model is 4.71 as compared to 3.39 from the menu cost model. Importantly, the results presented here support the view that the CalvoPlus model provides a sensible basis for a structural investigation of the data.

4.3 Selection Effects and Estimation Bias

As discussed above, the small positive contemporaneous selection effect we find in the regression exercise may be due to the way we define the zero band. Note that shrinking the I^{IZB} band in the analysis will have two consequences in that it reclassifies true

price changes as price changes in the data and potentially reclassifies true non-changing observations as price changes in the data. First, reclassifying small true price changes as price changes in the data would reduce the positive bias discussed above and drive down the point estimate in the probability regression. Second, to the extent there are small rounding errors in the price data, shrinking the I^{IZB} band creates misclassified price changes in the data. In a pure TD model this will not bias the point estimate in the probability model since the probability of being stuck with the old price and the measurement error in prices are independent of marginal cost. However, in a SD model, firms that do not change the price do so because they typically had small changes in marginal costs. Thus, reclassifying true non-changing observations as price changes will bias the point estimate downwards if the data is generated by a SD model. For this reason, comparing the baseline regression results with those obtained when shrinking the band towards only including exactly zero price changes yields an interval within which the true selection effect lies.

In the baseline formulation we include 1,349 small price changes in the I^{IZB} definition. Here we see what happens when we throw out a sizeable part of the very small price changes and shrink the I^{IZB} definition. To this end, we define I_{gt}^{IZB} to take on the value of one for $|d \ln P_{g,t}| \leq 0.0003$ and zero otherwise, implying that we throw out 75 percent of the included observed very small price changes from the $I_{gt}^{IZB} = 1$ definition. All in all, this leaves us with 866 $I_{gt}^{IZB} = 1$ observations, which constitutes 6.3 percent of the observations. Comparing column (1) and (2) in the top left panel of Table 5, we see that narrowing the band lowers the point estimate from 0.129 to 0.030 as expected, and that the coefficient is statistically insignificant in the latter case. It is also interesting to see that the standard error is actually about 20 percent smaller in column (2) as compared to column (1), thus speaking against that we have too little variation in the dependent variable when shrinking the I^{IZB} band. In a next step, we go all the way and only use the 529 observed exactly zero price change observations. This further diminishes the share of zero observations to 3.8 percent, but does not change the results qualitatively with an insignificant point estimate of 0.011. Interestingly, the standard error (0.037) falls by an additional fifteen percent.

One might be worried about using a linear probability model when the average probability of changing the price is 96.2 percent. However, first note that with a point estimate

Table 5: Estimation and Simulation Results - Band Size

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Band Size: I^{IZB}	Base 13.6%	Narrow 6.3%	Zero 3.8%	Base 13.6%	Zero 12.1%	Base 13.5%	Zero 7.9%
	Data			Simulations			
				Menu Cost		CalvoPlus	
$ d \ln MC_{jt} $	0.129* (0.053)	0.030 (0.043)	0.011 (0.037)	1.076 [0.031]	0.957 [0.029]	0.175 [0.031]	0.009 [0.025]
$ d \ln MC_{jt} $	0.114* (0.053)	0.017 (0.042)	-0.001 (0.035)	1.067 [0.033]	0.948 [0.031]	0.173 [0.033]	0.009 [0.026]
$ d \ln MC_{jt-1} $	-0.014 (0.072)	-0.052 (0.071)	-0.060 (0.071)	0.308 [0.035]	0.334 [0.031]	0.122 [0.036]	0.017 [0.030]

Notes: The dependent variable takes on a value of one if the price change is outside the zero band defined in the first row above and zero otherwise. Data panels: Superscripts * and ** denote estimates significantly different from zero at the five/one-percent level. The number of observations is 13,772 (12,292 when also including a lag in the model). Robust standard error clustered on the firm level insiden the parenthesis. Simulation panel: The coefficient denotes the average across 200 panel simulations. Standard deviation of the point estimate across 200 panels is inside the square bracket.

of 0.011 it requires more than a 38 standard deviation shock to marginal cost in order to predict a price change with unit probability. Also, all results in Table 5 are qualitatively unchanged from using a Probit or Logit estimator instead of the linear probability model.

In columns (4) and (5) of the top right panel of Table 5 we redo the experiment above on synthetic data from the Menu-Cost model. Here, we still expect a positive estimate when using only exactly zero price-change observations since in the SD world firms choose not to change price due to small changes in marginal cost and vice versa. Comparing the results in columns 4 and 5, we see that the point estimate falls slightly with 0.119 when going from using the zero bin to exactly zero price change observations in the probability regression (average point estimates across simulations are 1.076 vs. 0.957). Since there are no measurement errors and thus no associated cases of misclassified price changers in the synthetic data, this result gives a measure of the size of the positive bias from misclassifying small true price changes when relying on the baseline definition of I^{IZB} . Moreover, not much happens in terms of the share of observations when using the zero bin (13.6%) versus only exactly zero price changes (12.1%). Thus, in the Menu-Cost model the bulk of the observations in the baseline zero bin is exactly zero price-changes.

In columns (6) and (7) we do the same experiment in the calibrated CalvoPlus model. The point estimate drops from 0.175 to 0.009 when shifting the dependent variable from the baseline zero bin to only looking at exactly zero price changes. The intuition is that since the data want a calibration of the CalvoPlus model that is, for all relevant aspects, a standard Calvo model, there are no selection effects.²³ This exercise thus confirms that the time aggregation does not affect the basic intuition for the mechanisms at work. Moreover, the difference between the estimates, 0.166, gives a similarly sized estimate of the positive bias from including small positive price changes in the I^{IZB} definition as compared to the Menu-Cost model. It is also interesting to see that the CalvoPlus model wants about 7.9 percent of exactly zero observations. Thus, narrowing the band in the CalvoPlus case removes quite a share of true small price changes from the same.

In the bottom panels of Table 5 we redo the exercises outlined above when also including a lag in the model. As can be seen in the two bottom rows of Table 5, results are qualitatively unchanged from this extension. Also, comparing the results in columns (3) and (4) we see that the lagged effect in the Menu-Cost model is qualitatively unchanged from using the baseline zero bin or only the observations that are exactly zero.

The results suggest that the difference between estimated selection effects in the data when comparing the baseline with the results from relying on only the exactly zero observations is well in line with the bias estimates from the simulated data. In fact the point estimate of the drop (0.118) when shrinking the band is actually smaller than in the models, thus pointing away from the hypothesis that the estimate when only relying on exactly zero observation in the data is downward-biased due to misclassification of price changes in combination with state dependence in price-setting. Moreover, the results from fitting the Calvo Plus model, which indicate very little state dependence, suggest that the estimate in column (3) is more or less an unbiased estimate of the true selection effect. Thus, taken together, the results presented here lend support to the TD interpretation of the data and that the small contemporaneous effect reported in Table 3 is the result of upward bias from including small price changes in the zero bin.

²³In fact, setting $\kappa_L = 0$ and $\kappa_H = 150$ in the CalvoPlus model gives rise to nearly identical results to those presented in Table 5.

5 Concluding Discussion

We use very detailed Swedish micro data on product producer prices linked to a detailed data set containing information on the firms that set these prices to test the empirical relevance of selection effects in micro-level producer pricing. To impose discipline on the empirical exercise at hand, we first outline and calibrate a baseline SD model to match key moments in the data. The menu-cost model we rely on follows Nakamura and Steinsson (2008), but allows for fat-tailed idiosyncratic shocks to marginal cost (akin to Midrigan, 2011) in order to better match the micro data. Moreover, the model is calibrated to a monthly frequency, which then allows us to gauge the effect of time aggregation in the annual data we observe. Aggregating the data the same way as actual data is aggregated, we find that time aggregation gives a lot of mileage in replicating the observed price change distribution with a stylized menu-cost model. This is because the time aggregation filter fills out the gap of very small price changes otherwise expected in the price-change distribution from an SD model. Thus, time aggregation is a complementary mechanism for generating small price changes in SD models to the economies of scope suggested by Lach and Tsiddon (2007), Midrigan (2011) and Alvarez and Lippi (2014) or stochastic menu costs as in Caballero and Engel (1999) and Dotsey, King, and Wolman (1999). Intuitively, large positive and negative monthly changes within a year nearly cancel one another, which generates small overall price movements in the data. Also, the time-aggregation mechanism described here should be at work as soon as we leave ticker data and rely on data with intermittent price observations.

To analyze the strength of the selection mechanism we investigate if the absolute value of the change in the firm's marginal cost affects the probability of a price change and compare the findings from observed data to those from synthetic time-aggregated data generated by the SD model. We find a much smaller contemporaneous effect from the absolute value of the change in the firm's marginal cost on the probability of a price change than we would expect in the SD model. Moreover, we do not find any effect from the lagged absolute value of the change in the firm's marginal cost, which in an SD model would affect the price change probability through pent-up adjustment incentives. Moreover, when considering measurement issues pertaining to the classification of small price changes in the data, the small contemporaneous effect we find does seem to be the result of upward bias.

To structurally quantify the regression results we also fit a price-setting model that nests both TD and SD elements to the data (i.e. a fat-tailed shocks version of the Calvo-Plus model outlined in Nakamura and Steinsson, 2010), which can generate an arbitrary degree of selection effects in the simulated micro data from the model. Importantly, the procedure to fit the model parameters can be constructed to be unaffected by the measurement issues that may bias the regression results. When choosing parameters so that the model match empirical moments as closely as possible, the parameters are driven very close to a purely TD standard Calvo (1983) model. Thus, again pointing away from selection effects being an important feature of the data.

Thus, overall, timing adjustments of price changes in response to marginal-cost developments do not seem to be an important feature of observed price-setting behavior of goods-producing firms. Note that our data are drawn from firms upstream in the supply chain. Eichenbaum, Jaimovich, and Rebelo (2011) also link a measure of marginal cost, i.e. the replacement cost of the vended product, to the price set in data drawn from a large US food and drug retailer and documents a high degree of selection effects in pricing. This indicates considerable differences in pricing behavior along the supply chain. This is perhaps not surprising given differences in conditions between consumer and business-to-business markets.

Another important point, when thinking about the results found here, is that in the canonical New Keynesian model the TD price-setting frictions are usually added high up in the supply chain (intermediate-goods sector), whereas downstream sectors (retail sector) are, for convenience, modeled as frictionless; see e.g. Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005). Thus, this class of models does not need price-setting frictions throughout the whole supply chain in order to generate significant monetary non-neutrality. This then implies that frictions found in the downstream sectors can only add to monetary non-neutrality and given the results presented here, they are not instrumental for the existence of sizable monetary non-neutrality.

References

- ADOLFSON, M., S. LASÉEN, J. LINDÉ, AND M. VILLANI (2008): “Evaluating An Estimated New Keynesian Small Open Economy Model,” *Journal of Economic Dynamics and Control*, 32, 2690–2721.
- ALVAREZ, F., AND F. LIPPI (2014): “Price Setting With Menu Cost for Multi-Product Firms,” *Econometrica*, 82, 89–135.
- APEL, M., R. FRIBERG, AND K. HALLSTEN (2005): “Microfoundations of Macroeconomic Price Adjustment: Survey Evidence from Swedish Firms,” *Journal of Money, Credit and Banking*, 37, 313–338.
- BARRO, R. J. (1972): “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies*, 39, 17–26.
- BUCKLE, R., AND J. CARLSON (2000): “Inflation and Assymetric Price Adjustment,” *Review of Economics and Statistics*, 82, 157–160.
- CABALLERO, R., AND E. ENGEL (1999): “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach,” *Econometrica*, 67, 783–826.
- CALVO, G. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–398.
- CAPLIN, A., AND D. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *Quarterly Journal of Economics*, 102, 703–725.
- CARLSSON, M., J. MESSINA, AND O. NORDSTRÖM SKANS (2013): “Firm-Level Shocks and Labor Adjustments,” Mimeo, Uppsala University.
- CARLSSON, M., AND O. NORDSTRÖM 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, 1571–1595.
- CECCHETTI, S. (1986): “The Frequency of Price Adjustments: A Study of the Newsstand Prices of Magazines,” *Journal of Econometrics*, 31, 225–274.
- CHRISTIANO, L., M. EICHENBAUM, AND C. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.

- D'AGOSTINO, R. B., A. J. BELANGER, AND R. B. D'AGOSTINO JR. (1990): "A Suggestion for Using Powerful and Informative Tests of Normality," *American Statistician*, 44, 316–321.
- DAVIS, M., AND J. HAMILTON (2004): "Why are Prices Sticky, The Dynamics of Wholesale Gasoline Prices," *Journal of Money Credit and Banking*, 36, 17–37.
- DOTSEY, M., AND R. KING (2005): "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics*, 52, 213–242.
- 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, 655–690.
- DRUANT, M., S. FABIANI, G. KEZDI, A. LAMO, F. MARTINS, AND R. SABBATINI (2012): "How are Firm's Wages and Priced Linked: Survey Evidence in Europe," MNB Occasional Papers 102.
- EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): "Reference Prices, Costs and Nominal Rigidities," *American Economic Review*, 101, 234–262.
- EICHENBAUM, M., N. JAIMOVICH, S. REBELO, AND J. SMITH (2014): "How Frequent Are Small Price Changes?," *American Economic Journal: Macroeconomics*, 6, 137–155.
- FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 98, 394–425.
- GOLOSOV, M., AND R. E. LUCAS (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199.
- LACH, S., AND D. TSIDDON (2007): "Small Price Changes and Menu Costs," *Managerial and Decision Economics*, 28, 649–656.
- LEVY, D., S. DUTTA, AND M. BERGEN (2002): "Heterogeneity in Price Rigidity: Evidence from a Case Study Using Microlevel Data," *Journal of Money, Credit and Banking*, 34, 197–220.
- LOUPIAS, C., AND P. SEVESTRE (2013): "Cost, Demand and Producer Price Changes," *Review of Economics and Statistics*, 95, 315–327.

- MIDRIGAN, V. (2011): “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 79, 1139–1180.
- NAKAMURA, E., AND J. STEINSSON (2008): “Five Facts About Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123, 1415–1464.
- (2010): “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” *Quarterly Journal of Economics*, 125, 961–1013.
- ROTEMBERG, J., AND M. WOODFORD (1999): “The Cyclical Behavior of Prices and Costs,” in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford, vol. 1B, pp. 1051–1135. Elsevier.
- SHESHINSKI, E., AND Y. WEISS (1977): “Inflation and Costs of Price Adjustment,” *Review of Economic Studies*, 44, 287–303.
- SMETS, F., AND R. WOUTERS (2003): “An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area,” *Journal of the European Economic Association*, 1, 1123–1175.
- TAUCHEN, G. (1986): “Finite State Markov-Chain Approximation to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177–181.
- TAYLOR, J. (1980): “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy*, 88, 1–23.

Appendix

A Data

The data we use are drawn from the Industri Statistiken (IS) survey for plant-level data and the Industrins Varuproduktion (IVP) survey for the 8/9-digit price data, which can be linked to the producing plant.

The IVP survey provides plant-level information on prices and quantities for the years 1990 – 2002 at the finest (i.e. 8/9 digit) level of the Harmonized System (HS) for the years 1990 – 1995 and according to the Combined Nomenclature (CN) for the years 1996 – 2002. Although these two coding systems are identical only down to the 6-digit level, the change means that we have no overlap in the raw data at the most detailed level between 1995 and 1996. To avoid throwing away too much information, we need to merge spells across these two coding systems while minimizing the risk of creating spells of price observations for non-identical products. Thus, we take a very cautious approach by only merging price spells for products produced by firms that only produce a single product in 1995 and 1996 and whose product code is identical between 1995 and 1996 at the 6-digit level.

In the left-hand panel of figure 5, we plot the raw data distributions of log price changes (for 8/9-digit unit value data) for all price changes that we can match to the firms in the IS data (including the merged price spells in 1995/1996). All in all, this comprises 18,878 observations for 2,059 unique product codes and 4,385 unique product/firm identities across 934 firms. Each bin represents a log difference of 0.01. As can be seen in the figure, there is a substantial spike for the bin centered around zero. About 13.2 percent of the price-change observations are confined within the ± 0.5 percent interval (with 714 observations identically equal to zero, i.e. 3.8 percent).

Since the raw price data involve quite a few large swings (Max/Min. in the log price change distribution is 7.08/−7.65) we apply a cleaning procedure for the data used in the analysis. We are concerned with two types of errors in the price data. First, there may be measurement errors (of some magnitude) which show up as a zigzag pattern in the growth rate of the price and, second, there may be significant changes in, say, the quality of a product within a 8/9-digit product group, which will show up as a large

Figure 5: Histograms of raw data of log changes truncated at ± 1.1 . The left-hand panel describes the distribution of log price changes across 18,878 observations (for 2,463 different products across 943 firms). The right-hand panel describes the distribution of log unit labor cost changes across 17,760 observations (for 1,480 firms). Dashed lines indicate truncation limits. Bin size 0.01.

one-period increase in the difference. To remove the impact of this type of observations on the results, we split the individual price series and give them a new unique plant-price identifier whenever a large change in the growth rate appears in the data. We use the full distribution of log price change and determine the cut-off level as given by the 1.5 and 98.5 centiles of this distribution, depicted in the left-hand panel of figure 5. We also correct the firm-specific producer price index used to compute real output in unit labor cost by not using unit-value data in them for these observations. Moreover, price spells with holes in them are given separate unique plant-price identifiers for each separate continuous spell.

For the data from the IS database we start out with standard data quality checking, removing obviously erroneous observations like negative sales or a zero wage bill. Moreover, after constructing the firm-level variables needed, we remove firms which are subject to large swings in unit labor cost, since we aim at capturing normal behavior and not firms in extreme circumstances. In the right-hand panel of figure 5, we plot the log changes in firm-level unit labor cost for all firms (1,480) for which we can compute this measure in the IS data, in sum, 17,760 observations. The distribution is much less spread out as compared to the price change distribution with the Max/Min at 3.52/−3.79. Similarly, as with prices, we only keep firms that have unit labor cost changes that are inside the 1.5 and the 98.5 percentile of this distribution in all years (the limits are depicted by dashed lines in the right-hand panel of figure 5).

All in all, this then leaves us with 702 firms with at least one price spell that is longer than one period. The sample of industrial firms is dominated by small to medium sized firms with an average of 65 employees. The firms are distributed across 22 two-digit sectors (NACE). The four industries with most firms represented are industry 28 (Fabricated metal products, except machinery and equipment), industry 20 (Wood and products of wood and cork), industry 15 (Food products and beverages) and industry 29 (Machinery and equipment) with altogether 422 firms (out of the 702). The four smallest sectors, industry 14 (Other mining and quarrying products), industry 23 (Coke, refined petroleum products and nuclear fuels), industry 32 (Radio, television and communication equipment and apparatus) and industry 37 (Secondary raw materials), only have one firm.

When experimenting with more generous cut-off rules for prices and unit labor cost, we find the regression results presented in the top panel of Table 3 and the left panel of Table 5

in the main text to be qualitatively robust and, in fact, barely affected numerically. More specifically, we tried using the 1 and 99 centiles instead, leaving us with an estimation sample of 767 firms and 14,990 price-change observations in the final sample (751 firms and 13,368 price-change observations when also including a lag in the regression).

WORKING PAPERS*

Editor: Nils Gottfries

- 2012:16 Karolina Stadin, Vacancy Matching and Labor Market Conditions. 51 pp.
- 2012:17 Anne Boschini, Jan Pettersson, Jesper Roine, The Resource Curse and its Potential Reversal. 46 pp.
- 2012:18 Gunnar Du Rietz, Magnus Henrekson and Daniel Waldenström, The Swedish Inheritance and Gift Taxation, 1885–2004. 47pp.
- 2012:19 Helge Bennismarker, Erik Grönqvist and Björn Öckert, Effects of contracting out employment services: Evidence from a randomized experiment. 55 pp.
- 2012:20 Pedro Carneiro and Rita Ginja, Partial Insurance and Investments in Children. 32pp.
- 2013:1 Jan Pettersson and Johan Wikström, Peeing out of poverty? Human fertilizer and the productivity of farming households. 43 pp.
- 2013:2 Olof Åslund and Mattias Engdahl, The value of earning for learning: Performance bonuses in immigrant language training. 52 pp.
- 2013:3 Michihito Ando, Estimating the effects of nuclear power facilities on local income levels: A quasi-experimental approach. 44 pp.
- 2013:4 Matz Dahlberg, Karin Edmak and Heléne Lundqvist, Ethnic Diversity and Preferences for Redistribution: Reply. 23 pp.
- 2013:5 Ali Sina Önder and Marko Terviö, Is Economics a House Divided? Analysis of Citation Networks. 20 pp.
- 2013:6 Per Engström and Eskil Forsell, Demand effects of consumers' stated and revealed preferences. 27 pp.
- 2013:7 Che-Yuan Liang, Optimal Inequality behind the Veil of Ignorance. 26 pp.
- 2013:8 Pia Fromlet, Monetary Policy Under Discretion Or Commitment? -An Empirical Study. 57 pp.
- 2013:9 Olof Åslund and Mattias Engdahl, Open borders, transport links and local labor markets. 41 pp.
- 2013:10 Mohammad Sepahvand, Roujman Shahbazian and Ranjula Bali Swain, Time Investment by Parents in Cognitive and Non-cognitive Childcare Activities. 31 pp.
- 2013:11 Miia Bask and Mikael Bask, Social Influence and the Matthew Mechanism: The Case of an Artificial Cultural Market. 13 pp

* A list of papers in this series from earlier years will be sent on request by the department.

- 2013:12 Alex Solis, Credit access and college enrollment. 54 pp
- 2013:13 Alex Solis, Does Higher Education Cause Political Participation?: Evidence From a Regression Discontinuity Design. 48 pp.
- 2013:14 Jonas Kolsrud, Precaution and Risk Aversion: Decomposing the Effect of Unemployment Benefits on Saving. 37 pp.
- 2013:15 Helena Svaleryd, Self-employment and the local business cycle. 28 pp.
- 2013:16 Tobias Lindhe and Jan Södersten, Distortive Effects of Dividend Taxation. 22 pp.
- 2013:17 Jonas Poulsen, After Apartheid: The Effects of ANC Power. 42 pp.
- 2013:18 Magnus Gustavsson, Permanent versus Transitory Wage Differentials and the Inequality-Hours Hypothesis. 11 pp.
- 2013:19 Lovisa Persson, Consumption smoothing in a balanced budget regim. 35 pp.
- 2013:20 Linuz Aggeborn, Voter Turnout and the Size of Government. 50 pp.
- 2013:21 Niklas Bengtsson, Stefan Peterson and Fredrik Sävje, Revisiting the Educational Effects of Fetal Iodine Deficiency. 48 pp
- 2013:22 Michihito Ando, How Much Should We Trust Regression-Kink-Design Estimates? 69 pp.
- 2013:23 Bertil Holmlund, What do labor market institutions do? 25 pp.
- 2014:1 Oscar Erixson and Henry Ohlsson, Estate division: Equal sharing as choice, social norm, and legal requirement. 45 pp.
- 2014:2 Eva Mörk, Anna Sjögren and Helena Svaleryd, Parental unemployment and child health. 35 pp.
- 2014:3 Pedro Carneiro, Emanuela Galasso and Rita Ginja, Tackling Social Exclusion: Evidence from Chile. 87 pp.
- 2014:4 Mikael Elinder och Lovisa Persson, Property taxation, bounded rationality and housing prices. 38 pp.
- 2014:5 Daniel Waldenström, Swedish stock and bond returns, 1856–2012. 51 pp.
- 2014:6 Mikael Carlsson, Selection Effects in Producer-Price Setting. 34 pp.