

Do Menu Costs Make Prices Sticky?

Thomas A. Eife*

October 2008

This Version: November 2009

Abstract

This paper studies whether menu costs are large enough to explain why firms are so reluctant to change their prices. Without actually estimating menu costs, we can infer their relevance for firms' price setting decisions from observed pricing behavior around a currency changeover. At a currency changeover, firms have to reprint their price tags (menus) independently of whether or not they want to change prices. And if this is costly, firms' price setting behavior is altered in the months around the changeover. Using data from the Euro-changeover, the paper estimates that menu costs can explain a stickiness of around 30 days which is considerably less than the 7 to 24-month stickiness we observe in retailing and in the service sector. The reluctance of firms to adjust prices more frequently appears to be caused by factors other than menu costs.

Keywords: menu costs, price stickiness

JEL classification: E30

*University of Heidelberg (eife@uni-hd.de). I would like to thank Davide Cantoni, Christian Conrad, Johanna Kuehnel, and Peter S. Schmidt for helpful comments and suggestions.

1 Introduction

At a currency changeover, firms have to reprint their prices independently of whether or not they want to change prices (menus) and if changing prices is costly, firms will try to make the changeover coincide with a price change. This behavior will be reflected in the data. In the run-up to the changeover, firms will postpone price adjustments and, price changes originally planned for the months after the changeover will be anticipated. The higher the menu costs, the earlier firms will start postponing. Observing, for example, that an index is constant for six months before the changeover is a strong indication that menu costs can explain a stickiness of at least six months. We have to write “at least” because - as we will see below - firms might change prices more frequently in the run-up to the changeover than normal. The argument can be turned around: not observing firms postponing or anticipating is a strong sign that menu costs do not play a relevant factor in firm’s price setting.

Figure 1 illustrates this point. The figure shows restaurant prices in Germany in the years around the changeover. The vertical line denotes the changeover and, just as we would expect when firms anticipate and postpone, the index jumps at the changeover. Note, however, that menu costs can explain a jump only up to the extent that the jump is accompanied by periods of reduced inflation either before or after the changeover. The question is whether we observe a reduction in inflation. The dashed line indicates trend inflation in the period from 1996 until December 2000. In the 12 months before the changeover, inflation appears to be above trend and only in December 2001 do we observe firms postponing. Continuing this visual inspection and ignoring that after the changeover there appears to be no sign of firms anticipating, we could argue that menu costs cause a stickiness of one or two months in the restaurant sector. In fact, the restaurant sector is one of the few examples where we can observe at least some postponing. In general, menu costs do not appear to be a relevant factor in firms’s price setting decision. Using HICP data from the Euro-changeover, the paper estimates that menu costs can explain a stickiness of around 30 days which is considerably less than the 7 to 24-month stickiness we observe in retailing and in the services sector.

The years around the euro changeover in 2002 provide an interesting period for studying the menu cost hypothesis. For a number of years, inflation was steady and relatively low and since it took responsibility from the national central banks in 1999, the European Central Bank pursued a stable monetary policy with a strong emphasis on low inflation. The introduction of Euro coins and banknotes was not associated with a change in monetary policy and the discussion in the European public about whether the changeover lead to an increase in prices only started several months after the changeover (Del Giovane and Sabbatini [2005], Coombs and Eife [2007])

The next section gives an overview over the related literature and discusses this paper's contribution. In section 3, I set up and solve the firm's problem. Section 4 estimates the relevance of menu costs in firms' price setting decision and a short discussion in section 5 concludes the paper.

2 Related Literature

Menu costs, that is, the physical costs of changing prices and price tags and the costs of the decision process (sometimes called managerial costs) are a natural explanation for firms' reluctance to change prices more frequently. The literature that introduced menu costs as a microeconomic foundation for price stickiness started with Barro (1972), Akerlof and Yellen (1985), and Mankiw (1985).

In the empirical literature on menu costs, two strands are particularly relevant. The first strand, initiated by Levy, Bergen, Dutta and Venable (1997) estimates the size of menu costs. The important contribution of this literature was to show that menu costs are large enough to be regarded a non-trivial factor in firms' price setting decisions. An unfortunate drawback of this approach is that the size of menu costs does not provide information about how long firms hold back price adjustments because of these costs. It may well be that menu costs only impede weekly or daily changes. A firm's decision to change prices is driven by both the

cost and the benefit of doing so and estimates of the cost-side may not be enough to answer this question.

The second strand studies firms' price setting behaviour using time series on individual price notations. Even though the price data do not provide direct evidence on firms' motivation to change prices or not, there are certain features in the data that provide hints about the plausibility of the menu costs hypothesis. While the early studies by Carlton (1986), Cecchetti (1986) and Kashyap (1995) focus on specific products or markets, the more recent literature analyzes firms' price setting on the individual price data underlying official consumer (CPI) and producer (PPI) price indices. Access to these large data sets has been granted only recently by the statistical offices of a number of countries. For a study covering the U.S., see Nakamura and Steinsson (2008) and the earlier paper by Bils and Klenow (2004). Alvarez et al. (2006) and Dhyne et al (2007) summarize the findings for a number of country studies. The main findings are fairly robust across countries. In the retailing sector, prices are left unchanged for between 7 and 11 months on average. Restaurants (figure 1) keep prices constant for 12 to 24 months which is typical for most services. A feature in the data that is difficult to reconcile with menu costs is the large number of sales and also the fact that prices often return to the old regular price after a sale. Rotemberg (2004) and Nakamura and Steinsson (2005) discuss alternative explanations for price rigidity that appear to match quite well the observed features in the data.

The interesting aspect of the approach taken in the present paper is that it reveals directly whether menu costs hinder firms from changing prices more frequently. The benchmark reference for how a changeover affects firm's price setting decisions is Hobijn, Ravenna and Tambalotti (2006). Hobijn et al. solve the problem a firm faces during a changeover in a dynamic New-Keynesian framework. In this paper, I follow Hobijn et al. and adopt a partial equilibrium model but the firm's problem is modeled in a static setting where firms re-optimize their price path when the changeover is announced. This simpler approach has the merit that it draws attention to a common misperception in the literature: that observing

a “jump” in prices at the changeover is a sufficient condition for menu costs. The main predictions in the static model match those in Hobijn et al. This follows from the fact that finite menu costs only generate a transitory money non-neutrality. Higher than normal inflation at one point will be accompanied by unusually low inflation at some other point. Hobijn et al. (2006) argue that the jump in figure 1 can be explained by menu costs. The reason why I reach a different conclusion is due to the fact that I include the months around the changeover in the empirical analysis.

The present study focuses on the single issue of menu costs and does not answer what else might cause price stickiness; may it be implicit or explicit contracts, coordination failure (Cooper and John 1988), sticky information (Mankiw and Reis 2002), strategic complementarity in price setting (see for example the discussion in Woodford 2003) or the mechanism proposed by Rotemberg (2004) and Nakamura and Steinsson (2005). In addition, the data presented in this study raise a number of other questions. For instance, I do not discuss what might have caused the sudden increase in restaurant prices, why the increase is so persistent, and why we do not observe such a phenomenon in all euro-countries. Finding answers to these question is left for future work.

The model in this paper is a model of state-dependent pricing based on the classical model of inventory management. The goal of the model is to understand what *would* happen if menu costs played a relevant factor in firms’ price setting. Whitin (1953) introduced the inventory management model to the economics profession. In monetary economics, the inventory management model has been used in two different fields. (1) Baumol (1952) and Tobin (1956) apply the model to study money demand by households and Miller and Orr (1966) study the demand for money by firms. (2) Slightly modifying and re-interpreting the variables, the model has also been used to study price setting by firms. Important contributions are Sheshinski and Weiss (1983), Danziger (1983), and Caplin and Spulber (1987).¹

¹See Alvarez and Lippi (2009) for a recent application of the Baumol-Tobin model in a general equilibrium

3 The Model

The central assumption of the model is that prices are sticky and that the stickiness is caused by menu costs. The next subsection presents the basic model. Subsection 3.2 shows how a firm alters its price setting during a changeover and a discussion about the model's predictions and a simulation exercise in subsection 3.3 concludes the theoretical part.

3.1 The Firm's Problem (general)

Assume that firms are price setters and let P denote the firm's price and P^* the optimal (profit maximizing) price. Without menu costs, firms would set $P = P^*$ in every period but when changing prices is costly, firms keep prices constant for a while before making a larger adjustment. In order to keep the model simple, I assume that P^* is exogenous and grows at a constant rate π^* .² Deviations of the actual from the optimal price are denoted by $D = P - P^*$. The menu costs are given by k and the costs of deviating from the optimal price is denoted by r . Since both k and r are strictly positive, firms face a trade-off between paying menu costs and deviating from the optimal price. A cycle is defined as the time span between two price changes. The optimal length of a cycle (s) will depend on the relative price k/r . Figure 2 shows how the firm's price fluctuates around the optimal price. The deviations from the optimal price are shaded. The firm's problem can be viewed as minimizing the shaded areas under the constraint that changing prices is costly.

The firm minimizes costs. Figure 2 makes clear that the average (absolute) deviation from the optimal price equals $\frac{\Delta P}{4}$. The number of price changes is given by $\frac{\pi^*}{\Delta P}$. The firm's problem is thus given by

$$\min_{\Delta P} cost = k \frac{\pi^*}{\Delta P} + r \frac{\Delta P}{4}. \quad (1)$$

Rearranging the first-order condition, we can express the optimally chosen length of a

set-up.

²I discuss a version where π^* is stochastic in section 3.3.

cycle as

$$s \equiv \frac{\Delta P}{\pi^*} = 2\sqrt{\frac{k}{r\pi^*}}. \quad (2)$$

The length of a cycle increases as the menu costs (k) increase, as the costs of deviating (r) decrease and as inflation (π^*) decreases.

3.2 The Firm's Problem during a Changeover

Now suppose that there is a currency changeover and assume that the firm has to reprint its price tag (menu) when the new currency is introduced. In order to save costs, the firm will try to make the changeover coincide with a price change and it will re-optimize its price-path as soon as the changeover is announced. This behavior makes the problem of the firm in the run-up to the changeover that of a finite horizon.

Figure 3 illustrates the firm's problem and defines several new variables. Note that in contrast to the previous figure, only deviations (D) from the optimal price are shown. The upper panel of figure 3 shows the original price-path before the firm learns about the changeover and re-optimizes its path. Let the *initial cycle* be the cycle in which the changeover is announced. The exact moment of the announcement in the initial cycle is indicated by the parameter $\gamma \in [0, 2)$. When $\gamma = 0$, the changeover is announced at the very end of the initial cycle. The changeover takes place in the *final cycle* and the exact moment of the changeover is indicated by $\sigma \in (0, 2]$. When $\sigma = 2$, the changeover takes place at the very end of the final cycle. Note the reversal of direction with respect to γ . It is convenient to describe the firm's problem in terms of triangles, rather than cycles. Each cycle consists of two triangles, one below and one above the horizontal line. Let $b = \frac{\varepsilon}{2}$ denote the length of a triangle.

The firm has two options to make the changeover coincide with a price change: postponing (lengthening the original cycles) or anticipating (shortening the original cycles). Which of the two possibilities the firm chooses depends on the costs. In case the changeover already coincides with a planned price change, the firm will not adjust its path. Let $b_P > b$ denote the triangle length in case of postponing and $b_A < b$ the triangle length in case of anticipating.

The following observation simplifies the presentation of the firm’s problem. The decision whether to postpone or anticipate is (almost) independent of when exactly the changeover is announced (independent of γ). For nearly all admissible values of γ and σ we can, therefore, assume that the announcement occurs in the center of the initial cycle (where $\gamma = 1$). The reason for the qualification is given in the next paragraph.

Suppose that a firm finds it optimal to postpone. Since $b_P > b$, the end of the new cycle lies outside the initial cycle. The size of b_P will, therefore, be independent of γ and without loss of generality we can assume that $\gamma = 1$. A similar assumption can be made when the firm anticipates but only as long as the announcement does not occur towards the very end of the initial cycle. Since $b_A < b$, there is the possibility that the announcement comes “too late” to choose b_A and some other $b'_A \leq b_A$ has to be chosen. In this case, it is optimal to change prices immediately after the announcement. Most of the time, the firm’s decision to postpone or anticipate is independent of γ so I will refer to the case where this is not the case as the “irregular part” of the initial cycle.³

We need three more variables to study the firm’s decision. Let T denote the time span between announcement and changeover and \bar{T} denote the time span between the center of the initial cycle and the changeover. The variable m denotes the number of triangles between the center of the initial and the final cycle. When the firm postpones, there are m triangles between the center of the initial cycle and the changeover. In case the firm anticipates, there are $m + 2$ triangles (see figure 3).

The relationship between b_A , b_P , and b is given by

$$b_A = b \frac{m + \sigma}{m + 2} \tag{3}$$

$$b_P = b \frac{m + \sigma}{m}, \tag{4}$$

where I used the fact that $\bar{T} = (m + 2) b_A = m b_P = m b + \sigma b$. Intuitively, the firm divides the

³Assuming, for example, that the changeover takes place five years after its announcement and that firms keep prices constant for around 10 months, less than 3 percent of the firms will anticipate immediately (i.e. choose a b'_A).

period between the center of the initial period and the changeover ($m + \sigma$) equally among the triangles. In case the firm anticipates, there are $m + 2$ triangles, and when the firm postpones, there are m triangles. The relationship between b'_A and b is given by

$$b'_A = b \frac{m - 1 + \gamma + \sigma}{m + 1}. \quad (5)$$

Unlike b_A or b_P , the length of b'_A depends on the exact moment of the announcement (on γ). In this case, the number of triangles whose lengths the firm can optimize is $m + 1$. More information about how equation (5) is derived is given in the appendix.

We can now turn to the decision problem of the firm.⁴ When deciding whether to postpone or anticipate, the firm needs to take into account (a) that there are less price changes when postponing (less menu costs to be paid) and (b) that on average, it will deviate more from the optimal price when postponing. The costs are given by

$$\begin{aligned} C_A &= r (m + 2) F_A + \frac{m + 1}{2} k \\ C_P &= r m F_P + \frac{m - 1}{2} k, \end{aligned}$$

where C_A (C_P) are the costs when anticipating (postponing). The costs include the costs of deviating from the optimal price (first term) plus the menu costs (second term). The costs of deviating from the optimal price is given by r times the number of triangles (m or $m + 2$) times the area of the triangles (F_P or F_A). F_A (F_P) has length b_A (b_P) and height given by the inflation rate. In general, when σ is small, the firm will postpone and when σ is large it will anticipate. This seems intuitive. The optimal length of a cycle would be b , but when b does not allow the changeover to coincide with a price change, some other triangle length has to be chosen. For a small σ , the new triangle length b_P is still close to b and will be preferred. For a large σ , b_A and b are similar and when $\sigma = 2$, b_A and b are equal.

⁴Here I only sketch the decision when the announcement happens during the regular part of the initial cycle. Details can be found in the appendix. The calculations for the irregular part are similar and presented in the appendix as well.

Interestingly, the firm is not indifferent between anticipating and postponing when $\sigma = 1$, but somewhat to the left of the center. The indifference point (σ_i) can be found by setting $C_A = C_P$ and is given by

$$\sigma_i = \sqrt{m(m+2)} - m. \quad (6)$$

Note that $\sigma_i \leq 1$ and that $\lim_{m \rightarrow \infty} \sigma_i = 1$. An intuition why $\sigma_i \leq 1$ is that the costs of deviating from the optimal price (the area given by F_A and F_P) are convex in b_A and b_P . Increasing b_P by one unit increases the costs by more than what is reduced by decreasing b_A by one unit. In this sense, postponing is “punished” more than anticipating. Note that figure 3 was drawn such that the firm is indifferent between postponing and anticipating. With $m = 3$ as in the figure, $\sigma_i \simeq 0.9$.

In the irregular part of the initial cycle, the firm is indifferent between postponing and anticipating when

$$\sigma_i^{irr}(\gamma) = \sqrt{2m(m+1)((1-\gamma)^2+1)} - m(2-\gamma). \quad (7)$$

Here, the point of indifference is a function of both m and γ . Note that $\sigma_i^{irr}\left(\frac{b}{b_A}\right) = \sigma_i$, and that for all γ , $\sigma_i \leq \sigma_i^{irr} \leq 1$.

The following decision tree summarizes the findings.

1. Before the changeover, the firm chooses s (or b) optimally, given the inflation rate (π^*) and the two types of costs k and r .
2. When the changeover is announced, the firm learns γ and T .
3. Knowing this, the firm can calculate σ , m and \bar{T} .
4. The decision whether to postpone or to anticipate is then given by equations (6) and (7) and the size of the new cycles can be calculated using (3), (4), and (5).

3.3 Discussion

In this section I discuss the model's predictions and present a simulation exercise. Using the steps outlined in the previous section, we can simulate the price path of an individual firm. An aggregate of many firms can be generated by averaging over a number of individual price paths.

Figure 4 shows three examples of such paths. For the individual firm, I assumed that its original cycle length is 12 months ($s = 12$). Re-optimizing after the announcement, the firm in this example chooses to anticipate so that the cycles between announcement and changeover are somewhat shorter than 12 months. As intended, the firm starts a new cycle at the changeover whose length is again 12 months.

The second path (aggregate) shows the average price level of 365 firms. The firms are identical ($s = 12$), but their cycles are shifted. The first firm's cycle starts on the first day of the year, the second firm's cycle on the second and so on. These shifts have two effects. First, they make the aggregate increase smoothly during the first years and second, all firms will be at a different stage of their cycle when the changeover is announced. Each firm will, therefore, react differently; some will anticipate, others postpone, all with the objective of making the changeover coincide with the beginning of a cycle. Because of this re-optimization, we observe the characteristic pattern discussed before, a discrete jump and constant prices before and after the changeover. By re-optimizing, the firms synchronized their price setting after the changeover so that the aggregate increases stepwise like the individual path.

This synchronization is a strong prediction and we only observe it because all firms are identical. The third path (aggregate, heterogeneous), is an aggregate of firms with different s . For the first twelfth, I set $s = 12$, for the second $s = 13$, and so on until $s = 24$. Within each group, the cycles are shifted as in the previous example. Again, we observe constant prices before and after the changeover and the characteristic jump. This time, however, the synchronization fades out soon after the changeover.

Another interesting prediction of the model is that menu costs have a relatively short

effect on the level of prices. Even though the level jumps significantly at the changeover, it returns to its pre-changeover trend after only about half a cycle ($s/2$). After a full cycle, the level is predicted to be below its pre-changeover trend.

Figure 4 nicely illustrates the hypothesis of this paper, that firms' price setting behavior around the changeover reveals information about the "relevance" of menu costs. The higher the menu costs, the longer firms will postpone and anticipate. Turning the argument around, not observing firms postponing or anticipating is a strong sign that menu costs are only of minor concern to firms.

Recall that observing that an index is constant for 12 months before the changeover implies a stickiness of more than 12 months. This is because there are firms that anticipate, that is, their cycles are shorter than they would normally be. The model allows us to calculate the bias introduced by this change in behavior. Since $b'_A < b_A < b < b_P$, the shortest cycles are the cycles of firms that anticipate immediately after the announcement. The relationship between the observed s'_A and s , the variable we are interested in, is given by equation (5). The number of triangles (m) can be calculated from the time span between announcement and changeover and the values of γ and σ that maximize the wedge between s'_A and s are $\gamma = 0$ and $\sigma = \sigma_i$ (for smaller values of σ , the firm would postpone). Returning to the example above, observing that $s'_A = 12$ implies that s is nearly 15 months. For the period after the changeover, no such adjustment needs to be made.

In the model above, the optimal price follows a deterministic trend. Changes in wages, costs, or policy changes that affect the optimal price are incorporated by firms as long as these changes are expected. An important question is, however, how firms react to unexpected parameter changes. Unfortunately, the model though fairly elegant in its basic version becomes rather complicated when extended to allow for a stochastic inflation so I will only provide some intuition.

Consider the cycle immediately before the changeover and suppose that the optimal price unexpectedly rises. The firm then has the option to continue its originally planned cycle (and

bear the additional, unexpected costs) or to react and re-optimize its path. By reacting, the firm can reduce the unexpected deviations from the optimal price but this comes at the expense of paying menu costs an additional time; something the firm would do only for sufficiently large shocks. It can be shown that the firm's sensitivity is not constant but decreases towards the end of the cycle (towards the changeover). The closer the unexpected shock occurs to the changeover, the less likely the firm is to react.

4 Estimating the Relevance of Menu Costs

There are a number of points that I need to address before discussing how I test for the model's predictions. There are, for example, certain characteristics of the data that might introduce a bias in the estimates and I have to be careful that the data used are appropriate for the analysis. The section starts with a discussion about how price setting was regulated in Germany during the changeover and why using data from other countries might not be appropriate. I will then discuss the consequences of using aggregate data and other important characteristics. Sections 4.1 and 4.2 test the menu cost hypothesis, and a summary in section 4.3 concludes.

Price Setting Regulations An important assumption of the model is that firms change price-tags (menus) at the changeover. This is, however, not necessarily the case in practice because of the possibility to “dual price”. Dual pricing means that a firm denotes the price of an item in both the new and the old currency. Dual pricing allows firms to switch from one currency regime to another at a moment other than the changeover, violating the assumption. Using data from Germany allows us to get around this problem. In Germany, prices had to be denoted in the old currency until the changeover and dual pricing was optional. From the changeover on, prices had to be denoted in the new currency and dual pricing was again optional but only until the end of February 2002. From then on, dual pricing was not permitted. This means that in Germany all price tags had to be replaced within the eight

weeks from January 1st until February 28th 2002. I will argue, however, that most price tags were replaced already in the days after the changeover. There is no direct evidence for this behavior but the impact of the changeover on relative prices provides some hints. As figure 1 illustrates, the changeover affected relative prices in some sectors (both retailing and services) and the impact appears to have occurred almost entirely between December and January. Neither in February nor March 2002 there does appear to be a noticeable impact.⁵ Also in the estimations below, it is January 2002 where we notice an impact from the changeover and neither February nor March are conspicuous; providing additional support for the claim.

In all other Euro-countries, dual pricing was allowed for a significantly longer horizon - in some cases six months in others apparently indefinitely - making it more difficult to justify the assumption. The advantage of a shorter period of dual pricing is that consumers are encouraged to learn the new prices quickly. It is still discussed which approach is preferable but at the time, the German authorities reckoned a short transition period better and assured that the regulations were observed.

Aggregate Data For the data analysis, I am using the individual series of Eurostat's HICP basket. Unfortunately, micro data from the consumer price basket are not publicly available but I will argue that aggregate data are adequate for the argument of this paper. Note that the model's predictions have been tested using micro data. Hoffmann and Kurz-Kim (2006) using (classified) micro data from the German CPI basket, do not find evidence of firms postponing or anticipating in the six months before and after the changeover. The same conclusion will be drawn below.

Two issues arise when using aggregate data, the first might be called "causality" and the second concerns the heterogeneity of firms. Studying aggregate data and deducing characteristics of the underlying individual series might appear problematic and in fact, causality clearly goes from the individual firms to the aggregate. Consider the following two events:

⁵For a detailed description how the changeover affected relative prices see Buchwald et al. (2002). Bundesbank (2004) and Hoffmann and Kurz-Kim (2006).

A: The firms in an industry keep their prices constant.

B: The aggregate price index of the industry is constant.

Clearly, A implies B but the converse does not hold as there might be firms that increased and others that decreased leaving the average unchanged. What we need for the analysis is, however, something like the contrapositive. Not observing that the aggregate is constant implies that there are firms that adjusted prices and this observation is what is required for the exercise of this paper.

A different problem arises when the firms aggregated in one index are heterogeneous. It would skew the results if, for example, half the restaurants in figure 1 had menu costs so high, that it forced them to keep prices constant for 24 months and the other half menu costs so low that they can adjust every week. In this case, we would not observe a constant index as presumed above but we would observe a reduced inflation. This is, however, a testable implication as well and I will return to this point in section 4.2.

Quality Adjusted Data and Data Precision Eurostat’s HICP contains monthly observations of 83 indices. The data start in January 1996 and end in December 2007.⁶ Only quality-adjusted data are available for Germany. This is a potential drawback as the quality-adjusted data might not reflect underlying price movements and possibly introduce a bias in the estimates. In principle, the bias introduced can go both in favor and against the point I want to make. Consider again the example of restaurant prices in figure 1. In the month before the changeover, the index is constant and I argued that this could be explained by menu costs. The constancy of the index might, however, only be an artifact of the quality-adjusted data. It could be that the “true” data increased or decreased and only the judgment of the statistical office about quality adjustments made the index constant. The opposite might occur as well; that the “true” data are constant and only after adjusting for quality, movements in prices are added.

⁶The Euro changeover was announced in October 1997 and took place in January 2002.

Fortunately, there is more information available about the bias and I will argue that the bias is small and more importantly, that the bias appears to work against the point I want to make. Hoffmann and Kurz-Kim (2006) estimate that the average HICP inflation of 1.2 percent over the period 1997 until 2003 would have been 1.5 percent had the statistical office not adjusted for quality. On average, the quality adjusted data exhibit a lower inflation than the “true” data which introduces a bias in favor of the menu costs hypothesis. This bias seems small, however. A bias of 0.3 percentage points per year amounts to a bias of less than 0.03 percentage points per month. For industrial goods, the bias is larger. In the industrial goods sector, Hoffmann and Kurz-Kim (2006) estimate a bias of 0.9 percent per year which amounts to 0.075 percentage points per month.

This takes me to a second characteristic of the data. Unlike in other countries, the German data are published with only one decimal place. The average absolute size of the indices at the time of the changeover is about 100 points and for many series an increase or decrease of 0.1 percent per month is fairly large. This means that there can be considerable movement in the underlying “true” data which is not reflected in the index. This characteristic is a potential source of bias in favor of the menu cost hypothesis and will be important when I study whether the indices are constant around the changeover in the next subsection.

Pricing Points An interesting feature of the HICP basket is that most consumer goods are sold at “pricing points”, or threshold prices such as 1.99 or 24.90. Depending on how one defines pricing points, the estimates range from 72 to 95 percent of the data (Holdershaw et al. 1997, Fengler and Winter 2001, Bergen et al. 2003). This and the fact that the exchange rate from Deutschmark to Euro was 1.95583 $\left[\frac{DM}{Euro}\right]$ means that firms not only needed to reprint new price tags but also needed to decide whether to round to a new pricing point. In the literature, the costs of the decision making process are often referred to as “managerial costs”. What makes this feature interesting for the exercise in this paper is that the menu costs we estimate not only include the costs of printing new price tags but are likely to include the managerial costs as well.

The Model's Predictions The model's predictions are clear. In the months before and after the changeover, firms should keep prices constant when menu costs play a relevant role in their price setting. In addition, if the index has a non-zero trend, we should also observe a jump of the index at the changeover. When taking these predictions to the data, I will make three simplifications, all three due to practical reasons, though I should mention that the simplifications are "conservative" in the sense that they tend to favor the menu cost hypothesis.

First, I will ignore whether an index has a trend. An index without trend might be constant before and after the changeover but no information is revealed about menu costs. Ignoring the trend means that certain patterns in the data are (incorrectly) attributed to menu costs. The reason for this simplification is that it is often difficult to decide whether a series has a trend; a difficulty that arises mainly because the answer to this question depends on the period under consideration and a priori it is not clear, for example, whether the period under consideration should include the whole sample period or only the years around the changeover.

Second, I will ignore the jump and focus only on the periods before and after the changeover. In principle, the size of the jump could give information about menu costs as well, but several difficulties arise in practice. As mentioned above, it is often not clear whether a series has a trend or even whether the trend is positive or negative which makes it difficult to decide whether one should expect an index to jump up or down at the changeover. There are a number of apparently downward trending indices that jumped up when the new currency was introduced. Another difficulty is caused by firms' quest for attractive prices. The tendency to price at pricing points forces firms to round up or down at the changeover so that the size of the jump we observe does not necessarily reflect menu cost considerations.

Third, according to the model, we should expect an index to be constant both before *and* after the changeover. In the test below, I will consider it as a sign of the menu cost if an index is constant before *or* after the changeover.

4.1 Descriptive Statistics

We can get a first impression about the relevance of menu costs by studying how many constant indices there are before and after the changeover.⁷

Table 1 shows that 34 (28) of the 83 series are constant for one month before (after) the changeover. There are 18 (14) that are constant in the two consecutive months before (after) the changeover. Compared to the size of the basket, these numbers are not very high. Given the low inflation in Germany during our sample period, it is not unusual to observe constant indices so it would be interesting to see how likely it is to observe, say, 34 indices to be constant in a particular month.

Figure 5 shows that the answer to this question depends on whether we take the years before or the years after the changeover as reference. From the beginning of the sample in 1996 until about 2000 it appears quite usual that more than a third of the indices remains constant from one month to the other. Towards the end of the sample the picture changes. I do not have an explanation for the downward trend in the figure though my impression is that it is unrelated to the changeover.

For a more formal analysis regarding the likelihood to observe an index to be constant, I estimate the probability of such an event by bootstrapping methods. Table 1 shows the results.⁸ As figure 5 already suggested, compared to the years before the changeover (first part of the sample) it is likely to observe 34 constant indices in a particular month. Looking at the whole sample, the estimated probability is lower. Regarding the month after the changeover, table 1 shows that observing 28 constant indices as we do in the month after the changeover is not unusual in any of the two subsamples. Table 1 also shows that a similar picture arises when we look at the number of series that are constant for two consecutive

⁷Recall that the German data are published with relatively low precision so that there can be considerable movements in the underlying data that is not reflected in the aggregate indices.

⁸The probabilities in table 1 are calculated by randomly resampling the observations and calculating for each draw a 90% confidence interval. Repeating this 1000 times and counting how often the number of interest falls within the two bounds gives the probability. Since I rely on the empirical and not on a theoretical distribution, extreme probabilities such as 0 or 1 are not unusual. These should not be taken literally.

months.

To sum up, if we only consider the months after the changeover, no signs of menu costs appear. Considering the months before, there are some signs. However, the effect of menu costs on firms' price setting does not appear compellingly strong. In more than half of the sectors, firms do not appear to be very keen to postpone and only few firms postpone for more than two months.

The approach in this section is very basic, but it allows us to calculate an average “stickiness that can be explained by menu costs”. The procedure is best explained by an example. Consider again the restaurant prices in figure 1 that are constant for one month before the changeover. Two factors have to be taken into account. First, the data are collected at around mid-month so that we have to add 15 days to the 31 days observed.⁹ Second, the model in the previous section showed that there are firms whose cycle in the run-up to the changeover is shorter than otherwise. Using the method outlined in section 3.3 to compensate for this behavior, we find that menu costs can explain a stickiness in restaurant prices of around 50 days. Doing this for all sectors and weighing each sector by its weight in the basket gives a stickiness of 34.1 days that can be explained by menu costs. Looking at the period after the changeover, the stickiness is somewhat lower with around 33.5 days.

This is, admittedly, a rather crude measure but it nicely illustrates the main findings so far. There are some signs of menu costs in the data but these seem to be quite small. In the next section, where I pursue a more elaborate estimation procedure, the evidence in favor of the menu costs hypothesis is even weaker.

4.2 Regression Analysis

The idea of the approach taken in this section is, using a regression analysis, to predict inflation in the months around the changeover and to use this prediction to test whether firms

⁹Here I assume that the 31 days were caused by menu costs, ignoring that it is not unusual to observe a constant price level in this sector in other months.

postponed or anticipated. Compared to the approach taken in the previous section, a regression analysis has two advantages. First, it allows us to test whether inflation in the months around the changeover was reduced. As discussed above, when the firms aggregated in one index are heterogeneous, we might not observe a constant index as presumed before but we would observe a reduced inflation. A regression analysis allows us to test for this implication. The second advantage is that a regression analysis allows us to control for seasonalities and other regularities overlooked in the previous subsection. Consider the following model

$$\pi_t = X_t\beta + D\delta + \varepsilon_t, \tag{8}$$

where π_t is monthly inflation of a price index and X_t is a matrix of variables to be specified below. The variable D is a period-specific dummy that takes on the value one for a specific period and zero for all other observations. The coefficient estimate for this dummy (δ) is the forecast error for that period, and the estimated variance of the coefficient estimate is the estimate of the variance of the forecast error. Let a denote the actual value of inflation in the period under consideration and $p = a - \delta$ its predicted value. Observing that

$$|a| < |p|$$

can be interpreted as a sign that firms postponed or anticipated. We have to take absolute values because some indices are downward trending.¹⁰

I use the same model for all 83 price indices. This comes at the expense that the model might not be optimal for all sectors but the approach is more transparent and has the advantage that the results are easily replicated. Concerning the specification, I started with a general model and reduced the number of lags using the standard information criteria and scrutinizing the residuals. The goal was to find a specification that provides a reasonable fit for a number of different indices and extend it to the whole basket. In the baseline model,

¹⁰The case where $\text{sign}(a) \neq \text{sign}(p)$ barely occurs in practice.

I assume that inflation follows an autoregressive process with lags 1, 2, 3, 6, and 12. The matrix X_t also includes a constant and three dummies; one for the changeover and two for the two VAT changes in April 1998 and January 2007. Table 4 in the appendix shows regression results for a model with seasonal dummies and without the twelfth lag. The results are similar but slightly favor the point I want to make. Two indices had to be dropped from the model with seasonal dummies because of multicollinearity.

I am interested in how many of the 83 series in the basket overestimate and how many underestimate actual inflation in the months around the changeover. On average we should expect the model to overestimate the actual value in about half of the months. Deviations from this is a sign that some factor external to the model (e.g. menu costs) affected firms' price setting.

Table 2 reports the regression results. First consider December 2001 in the upper right column of table 2. In 50 out of the 83 indices, the model overestimates actual inflation ($|a| \leq |p|$). This can be read as a sign that firms have postponed price adjustments until the changeover. These 50 indices include the 34 from the previous subsection that kept prices constant (see table 1). Interestingly, nearly a quarter of the indices in the basket (19 out of 83) show a significantly higher inflation than the model predicts.

In November and October 2001, the reduction we observed in December seems absent. Interestingly, looking at the year before the changeover (ignoring December), the model suggests that in nearly three quarters of the series (61 out of 83), inflation is unexpectedly high. With the exception of December 2001, one month before the changeover, the postponing we would expect when menu costs are a relevant factor for firms seems to be absent. In the model with seasonal dummies (shown in table 4 in the appendix) there is even less evidence in favor of menu costs. With seasonal dummies, the indication that firms postponed in December disappears.

A similar picture arises in the months after the changeover where we should see firms anticipating which, again, causes the model to overestimate actual inflation ($|a| \leq |p|$). For

February 2002 one could argue that rather than an anticipation we observe the opposite. But the numbers are not that convincing so that, on the whole, the model neither over- nor underestimates actual inflation in the months after the changeover.

Before concluding this section, I present a different way to analyze whether firms altered their price-setting behavior in the months around the changeover. The test is as well based on a regression but rather than studying the signs of dummy variables, I analyze the residuals of these regressions. I first re-run the 83 regressions dropping D from equation (8) and collecting the residuals in a matrix Ω (with size 83×131). Each row of Ω contains the residuals of one of the 83 price indices. In order to make the residuals of the various regressions comparable, I standardize Ω by dividing each row by its standard deviations so that all rows have mean zero and standard deviation of 1. For illustrative purposes, I am also interested in the January residuals (January 2002 and others years). These are generated by dropping both D and the changeover dummy from the model (equation 8).

Figure 6 illustrates the test. The upper panel of figure 6 shows the kernel density of the residuals in January 2002 (dashed line) and the kernel density of a “typical” January, which consists of the combined residuals of all Januaries in the sample except 2002 and 2007 (VAT change). At the changeover, many prices increased which causes the prominent right-shift of the density.

The idea of the test is the following. Most indices have an upward trend so that we should expect the residuals to move left before and after the changeover when firms postpone or anticipate. The residual densities should, therefore differ from what we “typically” observe in other years. Let $f(x)$ denote the probability density function of the residuals in the period of interest and $g(x)$ the “typical” residuals from the same period in other years. Then the null hypothesis we wish to test is

$$H_0 : f(x) = g(x).$$

A convenient distance measure between two distributions $f(x)$ and $g(x)$ is their integrated square difference given by

$$J = \int [f(x) - g(x)]^2 dx.$$

Note that $J = 0$ under H_0 , and $J > 0$ if H_0 is false. I follow Li (1996) who proposes a test statistic based on J where $f(x)$ and $g(x)$ are replaced by kernel estimators as shown in figure 6.

Table 3 reports the test results.¹¹ For January 2002, the null hypothesis is clearly rejected. For all other months around the changeover, we cannot reject that the two sets of residuals come from the same underlying distribution. The lower two panels of figure 6 show the kernel densities for the two quarters before and after the changeover. In the fourth quarter 2002, the residuals shift slightly to the right and in the first quarter the residuals shift slightly left. In both cases, the shift is not strong enough to allow for a rejection of the null. The kernel densities for the two months immediately before and after the changeover (December 2001 and February 2002) look similar. December shows a slight shift to the right and February a slight shift to the left. Again, in both cases the shift is not large enough to reject the null. According to this test, no significant evidence appears in the data that firms altered their price setting around the changeover.

4.3 Summary

Overall, it is difficult to find convincing evidence that firms postponed or anticipated. There are signs that in the month immediately before the changeover some firms postponed, but these findings are not robust to different testing methods and quantitatively the effect is small. Regarding the type of sectors where we observe firms postponing, no real pattern emerges. Services firms, such as restaurants, seem to be somewhat more prone to postponing than others but the effect is minimal.

As mentioned before, the findings are in line with Hoffmann and Kurz-Kim (2006) who,

¹¹The p -values are bootstrapped; more information is given in the appendix.

using micro data from the German CPI basket, do not find evidence of firms postponing or anticipating in the six months before and after the changeover. It seems unlikely, that the reason why we do not see more firms postponing is that they had to react to unexpected changes in input costs or wages. No such shocks have been reported in the literature and recall that the precision of the published data was low so that we automatically allowed for - in some cases substantial - variation in the underlying prices. Bundesbank (2004) analyzes for a number of selected items (including restaurants) the evolution of major cost components (producer prices, wages, rents, and other input prices) and reports that the development of these cost components was stable in the months around the changeover. The accelerated inflation we observe in restaurant prices in the twelve months before the changeover, for example, appears not to be driven by supply-side factors. The authors argue that two years after the changeover, restaurant prices are above their pre-changeover price trend; something already suggested by figure 1.

5 Conclusion

This paper measures the “relevance” of menu costs by studying firms’ price-setting behavior around a currency changeover. At a changeover, firms have to reprint their price tags (menus) and if this is costly, firms will try to make the changeover coincide with a price change. In the run-up to the changeover, firms will postpone price adjustments and price changes originally planned for the months after the changeover will be anticipated.

Using data from the Euro-changeover in January 2002, the paper estimates that menu costs can explain a stickiness of around 30 days which is considerably less than the 7 to 24-month stickiness we observe in retailing and in the services sector. As argued, this estimate is based on relatively basic (though conservative) assumptions but it nicely illustrates that firms do not seem to care much about paying menu costs. It is difficult to find evidence of firms postponing or anticipating in the months around the changeover.

Though not directly the subject of this paper, the analysis raises the question of what caused the pointed increase in restaurant prices. In quite a large number of sectors, prices increased when Euro coins and banknotes were introduced. In the case of restaurant and some other services prices, it appears that these have stabilized at a higher level. There is a fairly large number of studies that describe the price movements but only few attempts have been made to discuss this phenomenon from a theoretical side.

6 References

- Akerlof, George A. and Janet L. Yellen (1985). “A Near-Rational Model of the Business Cycle, With Wage and Price Inertia.” *The Quarterly Journal of Economics* 100: 823-838.
- Alvarez, Fernando and Francesco Lippi (2009). “Financial Innovation and the Transactions Demand for Cash.” *Econometrica* 77(2): 363-402.
- Álvarez, Luis J., Emmanuel Dhyne, Marco Hoeberichts, Claudia Kwapil, Hervé Le Bihan, Patrick Lünnemann, Fernando Martins, Roberto Sabbatini, Harald Stahl, Philip Vermeulen and Jouko Vilmunen (2006). “Sticky Prices in the Euro Area: A Summary of New Micro-Evidence.” *Journal of the European Economic Association* 4(2-3): 575-584. *Journal of the European Economic Association* 4(2-3): 575-584., 04-05.
- Barro, Robert J. (1972). “A Theory of Monopolistic Price Adjustment.” *The Review of Economic Studies* 39(1): 17-26.
- Baumol, William J. (1952). “The Transactions Demand for Cash: An Inventory Theoretic Approach.” *The Quarterly Journal of Economics* 66(4): 545-556.
- Bergen, Mark, H Chen and Daniel Levy (2003). *Making Sense of Ignoring Cents: Another Implication of Rational Inattention*. Central European University Workshop. Bucarest.

- Bils, Mark and Peter J. Klenow (2004). “Some Evidence on the Importance of Sticky Prices.” *The Journal of Political Economy* 112(5): 947-985.
- Buchwald, Wolfgang, Jürgen Chlumsky and Nadine Engelhardt (2002). “Zum Einfluss des Euro auf die Preisentwicklung in Deutschland.” *Wirtschaft und Statistik* 3: 182 - 188.
- Bundesbank (2002). “Introduction of Euro Coins and Banknotes - Price Developments in January 2002.” *Monthly Report*, March: 8 - 11.
- Bundesbank (2004). “The Euro and Prices Two Years On.” *Monthly Report*, January: 15 - 28.
- Caplin, Andrew S. and Daniel F. Spulber (1987). “Menu Costs and the Neutrality of Money.” *The Quarterly Journal of Economics* 102(4): 703-726.
- Carlton, Dennis W. (1986). “The Rigidity of Prices.” *The American Economic Review* 76(4): 637-658.
- Coombs, W. Timothy and Thomas Eife (2007). *Coping with People’s Inflation Perceptions During a Currency Changeover*. Working Paper, University of Heidelberg.
- Cooper, Russell and Andrew John (1988). *Coordinating Coordination Failures in Keynesian Models*, The MIT Press. 103: 441-463.
- Danziger, Leif (1983). “Price Adjustments with Stochastic Inflation.” *International Economic Review* 24(3): 699-707.
- Del Giovane, Paolo and Roberto Sabbatini (2005). “La Divergenza tra Inflazione Rilievata e percepita in Italia”. In “L’Euro e l’Inflazione”. Edited by Paolo Del Giovane, Francesco Lippi and Roberto Sabbatini, Il Mulino.

- Dhyne, Emmanuel, Luis J. Alvarez, Herve Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lunnemann, Fabio Rumler and Jouko Vilmunen (2006). “Price Changes in the Euro Area and the United States: Some Facts from Individual Consumer Price Data.” *Journal of Economic Perspectives* 20(2): 171-192.
- Dutta, Shantanu, Mark Bergen, Daniel Levy and Robert Venable (1999). “Menu Costs, Posted Prices, and Multiproduct Retailers.” *Journal of Money, Credit and Banking* 31(4): 683-703.
- Fengler, Matthias and Joachim Winter (2000). “Price Variability and Price Dispersion in a Stable Monetary Environment - Evidence from German Retail Market.” *Sonderforschungsbereich 373: Quantification and Simulation of Economic Processes*.
- Hobijn, Bart, Federico Ravenna and Andrea Tambalotti (2006). “Menu Costs at Work: Restaurant Prices and the Introduction of the Euro.” *Quarterly Journal of Economics* 121(3): 1103 - 1131.
- Hoffmann, Johannes and Jeon Kurz-Kim (2006). “Consumer Price Adjustment Under the Microscope, Germany in a Period of low Inflation.” *European Central Bank Working Paper* 652.
- Levy, Daniel, Mark Bergen, Shantanu Dutta and Robert Venable (1997). “The Magnitude of Menu Costs: Direct Evidence From Large U. S. Supermarket Chains.” *The Quarterly Journal of Economics* 112(3): 791-825.
- Holdershaw, Judith, Philip Gendall and Ron Garland (1997). “The Widespread Use of Odd Pricing in the Retail Sector.” *Marketing Bulletin* 8: 53-58.
- Kashyap, Anil K. (1995). “Sticky Prices: New Evidence from Retail Catalogs.” *The Quarterly Journal of Economics* 110(1): 245-274.

- Li, Qi (1996). “Nonparametric testing of closeness between two unknown distribution functions.” *Econometric Reviews, Taylor and Francis Journals* 15(3): 261-274.
- Li, Qi and Jeffrey Scott Racine (2006). *Nonparametric Econometrics: Theory and Practice*. Princeton, Princeton University Press
- Mankiw, N. Gregory (1985). “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly.” *The Quarterly Journal of Economics* 100(2): 529-537.
- Mankiw, N. Gregory and Ricardo Reis (2002). “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *The Quarterly Journal of Economics* 117(4): 1295-1328.
- Miller, Merton and Daniel Orr (1966). “A Model of the Demand for Money by Firms.” *The Quarterly Journal of Economics* 80(3): 413-435.
- Nakamura, Emi and Jon Steinsson (2005). “Price Setting in a Forward-Looking Customer Market.” *Macroeconomics EconWPA* 0509010.
- Nakamura, Emi and Jon Steinsson (2008). “Five Facts About Prices: A Reevaluation of Menu Cost Models.” *Quarterly Journal of Economics* 123(4): 1415 - 1464.
- Rotemberg, Julio J. (2004). “Fair Pricing.” NBER Working Paper 10915.
- Sheshinski, Eytan and Yoram Weiss (1983). “Optimum Pricing Policy under Stochastic Inflation.” *The Review of Economic Studies* 50(3): 513-529.
- Tobin, James (1956). “The Interest-Elasticity of Transactions Demand For Cash.” *The Review of Economics and Statistics* 38(3): 241-247.
- Whiting, Thomson M. (1953). *The Theory of Inventory Management*. Princeton, Princeton University Press.
- Woodford, Michael (2003). *Interest and Prices*. Princeton, Princeton University Press.

A Appendix

A.1 Details to Section 2

This appendix describes in detail the decision problem of the firm. For convenience, the variables of the model are repeated here.

- P : actual price, set by the firm
- P^* : optimal price (assumed to increase over time at a constant rate $\pi^* > 0$)
- k : menu costs (independent of size of price change)
- r : costs of deviating from optimal price
- s : length of a cycle
- b : length of a triangle ($2b = s$)
 - $b_A(b_P)$: length of a triangle when the firm anticipates (postpones)
- m : number of triangles
- $\gamma \in [0, 2)$ indicates exact position of the announcement in the initial cycle
- $\sigma \in (0, 2]$ indicates exact position of the changeover in the final cycle
- T : time span from announcement to changeover.
- \bar{T} : time span from center of initial cycle to changeover
- $F_A(F_P)$: area of a triangle when the firm anticipates (postpones)

Describing the firm's problem simplifies if we use trigonometric functions. Let α denote the angle described by the slope of P^* , that is, by the inflation rate π^* . When inflation is

zero, $\alpha = 0$. Using this, we can derive an equilibrium condition linking the menu costs to the optimally chosen triangle length (b). In section 3.1, we found that

$$s \equiv \frac{\Delta P}{\pi^*} = 2\sqrt{\frac{k}{r\pi^*}}.$$

Following the definition of a cycle, there is exactly one price change within a cycle which implies that within a given number of cycles we must have that $\Delta P = \pi^*$ so that

$$k = r\frac{\Delta P}{4}.$$

Graphically, the menu costs (k) are proportional to the two triangles of an (optimally chosen) cycle. Rewriting this using α we get that

$$k = r \tan \alpha \times b^2.$$

A.2 The Decision Problem of the Firm

A.2.1 Regular Part

Whether the firm postpones or anticipates depends on the costs.

- C_A : costs when anticipating when announcement occurs in the regular part of the initial period.
- C_P : costs when postponing when announcement occurs in the regular part of the initial period.

The firm is indifferent when $C_A = C_P$. Without loss of generality, C_A and C_P are calculated as if the announcement was at the center of the initial cycle. (The additional terms that arise in the more general case would cancel when we compare C_A and C_P). Also

note that the menu costs that need to be paid at the changeover are not counted. This is just a convention without any effect on the final result. The additional terms would cancel when we compare C_A and C_P .

$$\begin{aligned}
C_A &= \underbrace{r (m+2) \frac{\tan \alpha}{2} b_A^2}_{r \times \text{number of triangles} \times \text{area of triangles}} + \underbrace{\frac{m+1}{2} k}_{\text{menu costs}} \\
C_P &= \underbrace{r m \frac{\tan \alpha}{2} b_P^2}_{r \times \text{number of triangles} \times \text{area of triangles}} + \underbrace{\frac{m-1}{2} k}_{\text{menu costs}}
\end{aligned}$$

$$\begin{aligned}
C_A &\geq C_P \\
r (m+2) \frac{\tan \alpha}{2} b_A^2 + \frac{m+1}{2} k &\geq r m \frac{\tan \alpha}{2} b_P^2 + \frac{m-1}{2} k \\
(m+2) b_A^2 + 2b^2 &\geq m b_P^2
\end{aligned}$$

Substituting $b_A = b \frac{m+\sigma}{m+2}$ and $b_P = b \frac{m+\sigma}{m}$ we find that

$$(m+2) \left(\frac{m+\sigma}{m+2} \right)^2 + 2 \geq m \left(\frac{m+\sigma}{m} \right)^2.$$

Setting both sides equal, re-arranging and solving for σ gives equation (6) in the text.

$$\sigma_i = \sqrt{m(m+2)} - m$$

A.2.2 Irregular Part

- C'_A : costs when anticipating, when the announcement occurs at the end of the initial cycle (the irregular period).
- C'_P : costs when postponing, when the announcement occurs at the end of the initial

cycle (the irregular period).

$$\begin{aligned}
C'_A &= \underbrace{r (m+1) \frac{\tan \alpha}{2} (b'_A)^2}_{r \times \text{number of triangles} \times \text{area of triangles}} + \underbrace{\frac{m+1}{2} k}_{\text{menu costs}} \\
C'_P &= C_P - \underbrace{r \frac{\tan \alpha}{2} [(1-\gamma)b]^2}_{r \times \text{area of triangle in initial period to the left of the announcement}}
\end{aligned}$$

$$\begin{aligned}
C'_P &\geq C'_A \\
r (m+1) \frac{\tan \alpha}{2} (b'_A)^2 + \frac{m+1}{2} k &\geq C_P - r \frac{\tan \alpha}{2} [(1-\gamma)b]^2 \\
(m+1) (b'_A)^2 + 2b^2 &\geq m (b_p)^2 - (1-\gamma)^2 b^2
\end{aligned}$$

using $b'_A = b^{\frac{m-1+\gamma+\sigma}{m+1}}$ and $b_p = b^{\frac{m+\sigma}{m}}$ we can rearrange terms to get

$$\frac{1}{m(m+1)} (m+\sigma)^2 + 2(m+\sigma) \frac{(1-\gamma)}{m+1} \geq \frac{(\gamma-1)^2}{m+1} + (1-\gamma)^2 + 2.$$

Setting both sides equal, re-arranging and solving for σ_i^{irr} gives

$$\sigma_i^{irr} \equiv \sigma(\gamma) = \sqrt{2m(m+1)((1-\gamma)^2 + 1) - m(2-\gamma)}.$$

The relationship between b'_A and b above was derived using the fact that

$$\begin{aligned}
T &= (m+1) b'_A \\
T &= (m-1)b + \gamma b + \sigma b.
\end{aligned}$$

A.2.3 Graphical Presentation

There is a nice graphical illustration of the decision problem of the firm. For a given m , the parameters σ and γ are sufficient to describe the problem. Figure 7 summarizes the findings. For small σ (changeover occurs shortly after an originally planned price change), it is always optimal to postpone. For large σ , it is always optimal to anticipate. The boundary where the firm is indifferent is constant and equal to σ_i for large values of γ . For small values of γ (smaller than $\frac{2-\sigma}{m+2}$), the boundary varies with γ . The area below $\gamma < \frac{2-\sigma}{m+2}$ (the dashed line in the figure), is the irregular part. The figure is drawn for $m = 1$ for illustrative purposes. For larger m , the irregular part is smaller.

A.2.4 Testing the Equality of two Distributions

The test is based on Li (1996), see also Li and Racine (2007). The test is asymptotically normal distributed,

$$T_n(J) \xrightarrow{d} N(0, 1)$$

but a small-sample bias has been reported, so I follow Mammen (1992) who suggests using bootstrapping methods to better approximate the null distribution of the test statistic. This is accomplished by randomly sampling with replacement from the pooled data. Let $\{X_i\}_{i=1}^{n_1}$ and $\{Y_i\}_{i=1}^{n_2}$ be the two sets of residuals and assume that X has a PDF $f(\cdot)$ and Y has a PDF $g(\cdot)$. Since under the null hypothesis both f and g are drawn from the same underlying distribution, we can pool them. Letting Z_i denote the i^{th} sample realization for the pooled data, I randomly draw n_1 observations from $\{Z_i\}_{i=1}^{n_1+n_2}$ with replacement, calling this sample $\{X_i^*\}_{i=1}^{n_1}$. Next I draw randomly n_2 observations with replacement from $\{Z_i\}_{i=1}^{n_1+n_2}$ and call this $\{Y_i^*\}_{i=1}^{n_2}$. The test statistic T_n^* is computed in the same way as $T_n(J)$ except with X_i and Y_i being replaced by X_i^* and Y_i^* , respectively. This procedure is repeated 1000 times. The reported p-value is the percentage of the 1000 bootstrapped T_n^* s above $T_n(J)$.




Figure 1: The figure shows German restaurant prices in the four years around the changeover (denoted by a vertical line). Upper left corner: same data for a longer horizon. Dashed line: trend inflation from January 1996 until December 2000.




Figure 2: The figure shows how the firm's price (P) fluctuates around the optimal price (P^*). The length of a cycle (s) is result of the firm's optimization problem.




Figure 3: Upper panel: original path before firm learns about changeover and adjusts its path. The changeover is announced in the initial cycle and takes place in the final cycle. Middle and lower panel: price path when firm adjusts cycles. The figure is drawn for $m = 3$, $\sigma = 0.9$, and $\gamma > b - b_A$.




Figure 4: The figure shows three simulated price paths. The announcement and the changeover are denoted by vertical lines. See text for more information.




Figure 5: The figure shows the number of indices in the basket (out of a total of 83) whose level remains constant in a particular month. The changeover and two VAT increases are denoted by vertical lines.




Figure 6: Kernel densities of the residuals. "Typical" refers to the residuals in the same period in other years. *First quarter is the first quarter after the changeover (February, March, April).




Figure 7: The figure shows the firm’s optimal choice for all possible combinations of σ and γ . For low (high) values of σ a firm will postpone (anticipate). The area below the dashed line is the “irregular” part. If a firm finds it optimal to anticipate in the irregular part it would do so immediately.

Constant Price Indices Around Changeover				
length of constant spell (in months)	before	(prob.)	(prob.)	after
	(full sample)		(full sample)	
1	34	(0.99)	(0.68)	28 (0.98) (1)
2	18	(0.97)	(0.67)	14 (0.93) (1)
3	6			9
4				
11	6			2
12	1			1

Table 1: The table displays the number of series in the basket (out of a total of 83) with no price change immediately before or after the changeover for the number of months given in the left column. Bootstrapped probabilities are shown in brackets. The probability is 1 if not otherwise indicated.

Predicted Versus Actual Inflation				
	Jan. - Nov. 2001	Oct. 2001	Nov. 2001	Dec. 2001
$ a > p $ (significant)	61 (15)	47 (36)	41 (31)	33 (19)
$ a \leq p $ (significant)	22 (2)	36 (17)	42 (22)	50 (22)
	Feb. - Dec. 2002	Feb. 2002	Mar. 2002	Apr. 2002
$ a > p $ (significant)	42 (9)	48 (24)	43 (19)	53 (23)
$ a \leq p $ (significant)	41 (7)	35 (22)	40 (26)	30 (25)

Table 2: The table reports the number of regressions in which the model's prediction (p) over- or underestimates the actual value (a) of inflation. Significance based on Newey-West standard errors.

Testing Equality of Distributions					
$H_0 : f(x) = g(x)$					
period	Nov. 2001	Dec. 2001	Jan. 2002	Feb. 2002	Mar.2002
p -value	0.14	0.18	0.00	0.92	0.46
period	Fourth Quarter 2001			First Quarter* 2002	
p -value	0.98			0.47	

Table 3: The test is whether the residuals in period x differ from the "typical" residuals in period x in other years. The p -value is the probability that accepting H_0 is wrong. *First Quarter is the first quarter after the changeover (February, March, April).

Predicted Versus Actual Inflation				
	Jan. - Nov. 2001	Oct. 2001	Nov. 2001	Dec. 2001
$ a > p $ (significant)	58 (14)	46 (28)	36 (24)	40 (20)
$ a \leq p $ (significant)	23 (2)	35 (13)	45 (15)	41 (14)
	Feb. - Dec. 2002	Feb. 2002	Mar. 2002	Apr. 2002
$ a > p $ (significant)	40 (5)	51 (15)	50 (9)	50 (18)
$ a \leq p $ (significant)	41 (4)	30 (17)	31 (14)	31 (15)

Table 4: The table reports the number of regressions in which the model's prediction (p) over- or underestimates the actual value (a) of inflation. Significance based on Newey West standard errors. Model: baseline without lag(12) and 11 seasonal dummies. 2 series had to be dropped due to multicollinearity.