# What We can Learn About the Behaviour of Firms from the Average Monthly Frequency of Price-Changes: An Application to the UK CPI Data*

Huw David Dixon† and Kun Tian‡,§

†*Cardiff Business School, Colum Drive, Cardiff CF10 3EU, UK*
*(e-mail: dixonh@cardiff.ac.uk)*
‡*University of Xiangtan, Xiangtan, China*
§*Centre for Price & Inflation, Cardiff University, Cardiff, UK (e-mail: ktian@hotmail.co.uk)*

## Abstract

The monthly frequency of price-changes is a prominent feature of many studies of the CPI micro-data. In this paper, we see what the frequency implies for the behaviour of price-setters in terms of the cross-sectional distribution average of price-spell durations across firms. We derive a lower bound for the mean duration of price-spells averaged across firms. We use the UK CPI data at the aggregate and sectoral level and find that the actual mean is about twice the theoretical minimum consistent with the observed frequency. We construct hypothetical Bernoulli–Calvo distributions from the frequency data which we find can result in similar impulse responses to the estimated hazards when used in the Smets–Wouters (2003) model.

## I. Introduction

In recent years, there have been many studies using comprehensive micro-data on pricing. For the US we have Bils and Klenow (2004), Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008). In the Euro area, there has been the inflation persistence network (IPN) consisting of national studies of the CPI and PPI microdata,[1] which are summarized in Alvarez *et al.* (2006). There are other studies: for example Bunn and Ellis (2012) for the UK and Baharad and Eden (2004) for Israel. One common focus of these studies has been

[1] See Baudry *et al.* (2007) and Alvarez and Hernando (2006) for France and Spain *inter alia*.

the *frequency* statistic of the proportion of prices changing per month (this can either be an average over several months, or a monthly statistic). This statistic can be presented in several ways, depending on the level of disaggregation and the treatment of temporary sales and so on.[2] In this paper, we seek to analyse what this frequency statistic implies for the behaviour of 'firms' (or more accurately *price-setters*) in the economy in terms of the mean duration of prices set by firms. Each period firms set prices: they may either choose to leave the price unchanged or to change it. The proportion of firms resetting price corresponds to the proportion of prices changing (for simplicity we take a 1-1 correspondence between firms and prices[3]). The prices of some product types change frequently (e.g. gasoline, tomatoes) and some very infrequently (coin-operated laundromats).

We can think of the CPI data set as a panel of observations, each cross-section corresponding to the prices set by the price-setters at that date. The *cross-sectional* distribution and corresponding mean completed price spells can be seen as capturing the behaviour of the firms, which represents the 'structure' of the economy in this respect: what proportion of firms set prices of a particular duration, what is the average behaviour of the firms in the economy. The purpose of this paper was to explore what we can learn about the *cross-sectional* distribution from the frequency data. The cross-sectional distribution is also required if we are to calibrate price-setting on the economy as a Generalized Taylor economy (GT) where we consider the economy as made up of price-setters who set prices for different durations which are known *ex ante* (Taylor, 2016, Dixon and Le Bihan, 2012). What is the cost of losing information by summarizing the distribution of price changes based on the frequency data? In fact, we show that the loss can be quite small: it turns out that the behaviour of the economy using the hypothetical distributions can look quite similar to those based on the estimated distribution. However, a wide variety of hypothetical distributions are consistent with a given frequency and some can give rise to behaviour that are quite different. We need to know which hypothetical frequency-based distributions work best. In particular, Dixon and Kara (2010, 2011) assume that the frequencies correspond to the Bernoulli–Calvo distribution – there is a fixed *Bernoulli* probability each period that a firm's price will change.

Our approach to interpreting the frequency data is twofold: first, we ask a purely theoretical question – *what is the range of possible cross-sectional means consistent with a given frequency*; secondly, we compare the cross-sectional distribution estimated from the micro data with two hypothetical distributions – (a) the 'minimum distribution' generating the lowest mean durations and (b) the Bernoulli–Calvo distribution which is implied by the popular Calvo model of pricing. We use the UK CPI data for the period 1996–2007 and consider frequency data at three different levels of disaggregation: the 11 sector COICOP,[4] the 67 sector COICOP and the highest possible level of disaggregation at 570 items, to see how the actual data on price-spell durations compares to the hypothetical distributions consistent with the frequency data. We find that the cross-sectional mean duration

---

[2] See Kehoe and Midrigan (2015) for a discussion of the US data.

[3] See Alvarez and Lippi (2014) for the case of multi-product firms.

[4] COICOP stands for 'Classification of Individual Consumption According to Purpose' and is an international standard used for constructing consumer price indexes (see e.g. UN Statistics division http://unstats.un.org/unsd/cr/registry/regcst.asp?Cl=5).

estimated directly from the CPI data is 10.9 months and the median is 7.8 months. We use the full data set and include all price-spells including sales and other short-lived prices.

Turning first to the theoretical minima generated from the frequency data. In *Proposition 1* we find that the theoretical minimum consistent with a given frequency is only attained if all price-spells have the same or almost the same duration, whilst in the data the distribution contains a lot of heterogeneity in durations which implies a cross-sectional mean that will far exceed this theoretical minimum. Hence it is no surprise to find that the estimated mean is 62–90% higher than the theoretical minimum means generated at different levels of disaggregation.[5] The theoretical minimum is not just of theoretical interest: Carvalho, Dam and Lee (2015) take the frequency data from Bils and Klenow (2004) and Nakamura and Steinsson (2008) and use the minimum mean distribution as a prior in their Bayesian estimation of the cross sectional distribution.

Secondly, we interpret the frequency data under the hypothesis that within the sectors the frequency is generated by a Bernoulli–Calvo distribution as has been assumed in applied work by Dixon and Kara (2010, 2011). We look at this in two ways. We first take implied distribution of durations in each sector (for a given level of disaggregation) and then for each duration we aggregate over sectors using the CPI weights to find the aggregate distribution. We can then compare this implied distribution derived from the frequencies to the distribution estimated directly from the data survival function.

- The aggregate distribution derived under the Bernoulli–Calvo hypothesis at the sectoral level has a similar mean and median to the estimated distribution, with the mean increasing with the level of disaggregation. For example, the 570-item model yields a mean of 10.8 and median of 7.8 months, whilst the estimated values are 10.9 and 7.8 respectively.
- However, the implied Bernoulli–Calvo distributions differ to the estimated distribution in significant ways: (i) there is a 12-month spike in the estimated distribution absent from the Bernoulli–Calvo distributions and (ii) the proportion of short price-spells (1–3 months) is less in the Bernoulli–Calvo than in the estimated distribution.[6]

Whilst the Bernoulli–Calvo distributional hypothesis might not provide a good statistical fit, does this matter in terms of how the economy behaves? Since the Bernoulli–Calvo hypothesis yields a mean and median close to the estimated distribution, perhaps the differences will not result in different behaviour of the economy in terms of impulse response functions (IRFs). We explore this in the context of two macro models: a simple Quantity Theory (QT) model and the Smets and Wouters (2003) Euro area model. We also calibrate the models for quarterly and monthly versions. The pricing models used are the Generalized Taylor (GT) and Generalized Calvo (GC) as in Dixon and Le Bihan (2012), which are both consistent with any micro distributions of durations and can be calibrated both to the

---

[5]The theoretical *maximum* is only bounded by an upper bound on durations: a frequency of 20% could be generated by 20% of firms changing price each period, with 80% never changing price, yielding an infinite mean.

[6]We also examined whether the Bernoulli–Calvo is a good fit in each of the 11 COICOP sectors in an earlier version of the paper (Cesifo working paper 4226). Since the data set is large, even small deviations of the actual distribution from the hypothetical Bernoulli–Calvo distribution cause the Bernoulli–Calvo null to be rejected under the Kolmogorov–Smirnov test, which is indeed the case. However, whilst in some sectors the hypothesized Bernoulli–Calvo distribution looks completely different (e.g. in Health which has a very large 12-month spike), in others the Bernoulli–Calvo looks more similar (Transportation).

estimated distribution and the hypothetical distributions under both the Bernoulli–Calvo hypothesis and the theoretical minimum duration corresponding to Proposition 1.

Our findings for the IRFs are mixed. Whilst the best hypothetical distribution yields IRFs close to those generated by the estimated distribution, sometimes they can be far away. The minimum duration distributions all perform badly relative to the Bernoulli–Calvo distributions. Amongst the Bernoulli–Calvo distributions, in general we find that the most disaggregated gives the worst fit in terms of IRFs for the GT pricing model. This suggests that the 'best fit' in terms of the cross-sectional distribution does not imply the best calibration for IRFs. In almost all cases considered, we find that the *intermediate* level of disaggregation (67 COICOP) gives the IRFs closest to the one from the estimated distribution. We also find that the differences in calibration matter less in the Smets–Wouters (SW) model than in the simple QT: this is not surprising, since there are more sources of dynamics in the SW framework than nominal rigidity. Since most macroeconomists will be interested in using micro-data to calibrate in models of a similar degree of complexity to the SW model, our results suggest that using the sectoral frequency data under the Bernoulli–Calvo distributional hypothesis might be a useful shortcut and alternative to estimating the distribution using the hazard function. Indeed, where the hazard function is not available or not estimated reliably, we can be confident that the use of sectoral frequencies with the Bernoulli–Calvo distributional hypothesis can be a good working approximation. However, using the most disaggregated data might not lead to the best calibration. Since the Bils–Klenow data set used by Dixon and Kara (2010, 2011) is highly disaggregated with 350 frequencies, our results from the UK CPI data suggest that it would have been better for them to have used the frequency data in a more aggregated form.

The structure of the paper is as follows. In section II, we give a theoretical description of the steady state distributions of durations. In section III, we derive the propositions in which the average duration across firms consistent with the mean frequency of price changes. We show an application to the UK CPI data in section IV. In section V, we compare the different microeconomic distributions by simulating a simple QT model and a SW model, both with GT and GC price-setting respectively. We conclude in section VI.

## II.    Steady state distributions of durations

The statistical framework for understanding the CPI microdata is outlined in detail by Dixon (2012), so in this paper we just summarize in a less technical manner the key properties needed for this paper. There is a continuum of price-setting firms $f \in [0, 1]$, time is discrete[7] and infinite $t \in Z_+ = \{0, 1, 2 \ldots \infty\}$. The price set by firm $f$ at time $t$ is $p_{ft}$. A price spell is a duration, a sequence of consecutive periods that have the same price $p_{ft}$. For every $\{t, f\} \subseteq [0, 1] \times Z_+$ we can assign an integer $d(t, f)$ which is the duration of the price-spell to which $p_{ft}$ belongs.[8] The distribution of price-spell durations is simply the proportions of all durations having length $i = 1 \ldots F : \boldsymbol{\alpha}^d = \{\alpha_i^d\}_{i=1}^{F} \in \Delta^{F-1}$. We assume a steady-state, so

---

[7]Typically, CPI data are collected on a monthly basis, the price observations being obtained in the first two weeks of the calender month.

[8]Note that in assigning an integer to a duration, we start with 1 by convention: it would be equally valid to start with 0. With our convention, a new price-spell is 1 month old, rather than becoming 1 on completion of the first month.

that the distribution of durations of new price-spells is the same for each new cohort of price-spells. This means that the distribution of all price-spells is exactly the same as the distribution of new price spells at any period.

Whilst the distribution of durations $\boldsymbol{\alpha}^d \in \Delta^{F-1}$ is one way of looking at the microdata, it ignores the *panel* structure of the data, since it puts all of the price-spells together and ignores their start date. However, each row of the panel is a trajectory of prices corresponding to a particular firm (or more accurately product sold at an outlet). Each column is a cross-section of all of the prices set by firms at a point in time. The cross-sectional distribution of *completed* price-spell durations is $\boldsymbol{\alpha} \in \Delta^{F-1}$. This takes a representative $t$, and for each firm we see the completed price-spell duration at that time $d(f,t)$. $\alpha_i$ is then the proportion of firms that set prices for $i$ periods. Equivalently, it is the proportion of price-spells at any time $t$ which last for $i$ periods. We will call this cross-sectional distribution the *distribution across firms*, or DAF for short.

To see why we are interested in the cross-sectional DAF, we can consider the following example. Think of an economy with two sectors of equal size. In one, prices are perfectly flexible and change every month. In the other fix-price sector, prices only change once every 5 years. Suppose we have one observation taken from each sector each month over 5 years coinciding with one fix-price spell. We will have a panel with two rows and 60 columns: a total of 120 price observations. If we look at the price-spells, we have 61: one 5-year spell and 60 one month spells. The average spell duration is approximately 2 months (120/61). The frequency of price change is 50%. You might conclude that with price-spells being on average 2 months, prices were quite flexible. However, this is not the case: for half the economy prices are very inflexible and last 5 years. This is better captured by the cross-sectional mean of 30.5 months.

The proportion of firms re-setting price each month is denoted as $\bar{h}$: in the UK this is equal to 21%. We define the mean duration[9] of price-spells across the panel as a whole as

$$\bar{d}\left(\boldsymbol{\alpha}^d\right) = \sum_{i=1}^{F} i \alpha_i^d$$

and cross-sectional mean (across firms or price-setters) as

$$\bar{T}(\boldsymbol{\alpha}) = \sum_{i=1}^{F} i \alpha_i \tag{1}$$

Note that the cross sectional mean in general be larger than the mean duration $\bar{T} \geqslant \bar{d}$: this is because in cross-section you have length-biased sampling, since the probability of a price-spell being observed in cross-section is proportional to its duration. Indeed, the two can be equal ($\bar{T} = \bar{d}$) *if and only if* $\alpha_F = \alpha_F^d = 1$, so that all price-spells are $F$ months long and there is no heterogeneity to generate a length bias.

---

[9] Again, this way of defining the mean is consistent with our convention of assigning the integer 1 to the first time period. Had we instead assigned a 0 to this value, then we would have the expression $\sum_{i=1}^{F}(i-1)\alpha_i$ (as we do with human ages). An equally acceptable measure is to take the midpoint and have $\sum_{i=1}^{F}(i-0.5)\alpha_i$. We can move between these definitions simply by adding or subtracting a constant.

From Dixon (2012), we know that:[10]

$$\bar{h} = \bar{d}^{-1} \tag{2}$$

$$= \sum_{i=1}^{F} \frac{\alpha_i}{i} \tag{3}$$

That is, the proportion of firms resetting price is equal to the reciprocal of the mean duration $\bar{d}$. Furthermore, the proportion of firms resetting price is related to the cross-sectional distribution by equation (3). In steady-state, a proportion $i^{-1}$ of the $\alpha_i$ $i-$duration firms reset their price. The aggregate proportion is simply the sum over the durations $i = 1...F$.

**The average duration across firms consistent with $\bar{h}$: some theoretical results**

Now, for a given frequency $\bar{h}$ there are many possible distributions across firms (DAF) $\boldsymbol{\alpha} \in \Delta^{F-1}$ consistent with identity (3) and each distribution results in a corresponding mean across firms $\bar{T}(\boldsymbol{\alpha})$. We can define the mapping $H(\bar{h}) : [0, 1] \to \Delta^{F-1}$

$$H(\bar{h}) = \left\{ \boldsymbol{\alpha} \in \Delta^{F-1} : \sum_{i=1}^{F} \frac{\alpha_i}{i} = \bar{h} \right\}$$

$H(\bar{h})$ is the set of all *DAFs* which are consistent with a given mean duration of price-spells $\bar{d}$ expressed in terms of the corresponding proportion of firms resetting prices $\bar{h}$. Clearly, since the maximum duration is $F$, we have $\bar{h} \geqslant F^{-1}$ so that $H$ is non-empty. Since $H(\bar{h})$ is defined by a linear restriction on the sector shares $\boldsymbol{\alpha}, H(\bar{h}) \subset \Delta^{F-2}$ and is closed and bounded. Since from (1) $\bar{T}(\boldsymbol{\alpha})$ continuous, both a maximum and a minimum value for $T$ will exist consistent with $\boldsymbol{\alpha} \in H(\bar{h})$. The *minimum* $\bar{T}$ consistent with a given $\bar{h}$ is given by:

$$\bar{T}^{\min} = \min \bar{T}(\boldsymbol{\alpha}) \text{ s.t.} \boldsymbol{\alpha} \in H(\bar{h}) \tag{4}$$

The *maximum* $\bar{T}$ consistent with a given $\bar{h}$ is given by:

$$\bar{T}^{\max} = \max \bar{T}(\boldsymbol{\alpha}) \text{ s.t.} \boldsymbol{\alpha} \in H(\bar{h}) \tag{5}$$

*Proposition 1.* For a given frequency $\bar{h}$: the minimum mean from equation (4) is $\bar{T}^{\min} = \bar{h}^{-1} = \bar{d}$. The maximum average contract mean from equation (5) is $\bar{T}^{\max} = F(1-\bar{h}) + 1$.

Proofs are presented in the Appendix. To understand Proposition 1, we just need to think of what is generating the mean duration of spells $\bar{d}$ and the proportion of firms changing price each period $\bar{h}$. There is the unit interval of firms, divided into proportions $\alpha_i$ with different price-spell durations $i = 1...F$. Firms with price-spell lengths $i$ will set prices once every $i^{-1}$ periods: the longer the price-spell, the more infrequently the firm will reset price. Hence, we can have the same proportion of firms re-setting price (and hence same mean duration of spells) and increase the mean duration across firms by having more very short (1 period) and more longer price-spells. The maximum $T^{\max}$ is reached when we have as

[10] In continuous time, we have $\bar{d} = -\frac{1}{\log(1-\bar{h})}$, which allows for the price to change more than once per period. Again, we are using a discrete time setting in which durations are integer valued.

many $F$ period contracts as possible, consistent with $\bar{h}$, which means we have a mix of 1 period and $F$ period price-spells. The existence of a maximum relies on us assuming an upper bound $F$: clearly, as $F \to \infty$, $T^{\max} \to \infty$. The minimum occurs when all firms have price-spells of (almost) the same length: if $\bar{d}$ happens to be an integer, then all price-spells have exactly the same length and the distribution of price spells durations and the cross section are the same: $\boldsymbol{\alpha}^d = \boldsymbol{\alpha}$ (there is no length bias in cross-section). Note that $\bar{T}^{\max} = \bar{T}^{\min}$ if and only if $F = \bar{h} = 1$, with all prices changing every period.

Thirdly, we look at a particular distribution: the discrete time Bernoulli process where the probability of a price-change ending the price-spell in any particular period (the Bernoulli measure) is $\bar{h}$. This is important in the macroeconomic context, because the Calvo model of price-setting uses the Bernoulli measure as a key parameter (often referred to in macroeconomics as the *Calvo probability* in the context of wage or price-setting). In terms of the distribution of durations $\boldsymbol{\alpha}^d$ this gives rise to the familiar Geometric distribution:

$$\alpha_i^d = (1 - \bar{h})^{i-1}\bar{h}$$

As shown in Dixon and Kara (2006), this corresponds to the cross-sectional distribution $\boldsymbol{\alpha} = \{\alpha_i\}_{i=1}^{\infty}$ where:

$$\alpha_i = i.\bar{h}^2(1 - \bar{h})^{i-1} \tag{6}$$

We will denote equation (6) as the Bernoulli–Calvo (BC) specification of the cross-sectional DAF. The corresponding cross-sectional mean is:

$$\bar{T}^{BC} = 2.\bar{h}^{-1} - 1. \tag{7}$$

The Bernoulli–Calvo mean is at most almost twice the theoretical minimum from Proposition 1, since:

$$\bar{T}^{BC} = 2.T^{\min} - 1$$

Note that $\bar{T}^{BC} = \bar{T}^{\min} = d$ iff $\bar{h} = 1$. That is the Bernoulli–Calvo mean equals the minimum if and only if all are perfectly flexible and change every period. If some price remain unchanged each period ($\bar{h} < 1$), then $\bar{T}^{BC} > \bar{T}^{\min}$. For example, if the Bernoulli measure is 0.25, then $\bar{T}^{BC} = 7$ and $\bar{T}^{\min} = 4$; a Bernoulli measure of 0.10 implies $\bar{T}^{BC} = 19$ and $\bar{T}^{\min} = 10$.

## III.   An application to the UK CPI data

In this section, we take the frequency data from the UK and construct[11] the three hypothetical distributions corresponding to the upper and lower bounds for the cross-sectional mean duration and the Bernoulli–Calvo distribution. These three distributions are then compared to the estimated distribution. The data set we use for is the locally collected CPI price microdata covering the period January 1996 to December 2007 which is described

---

[11]This work contains statistical data from ONS which is Crown copyright and reproduced with the permission of the controller of HMSO and Queen's Printer for Scotland. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research data sets which may not exactly reproduce National Statistics aggregates

in detail in the Appendix. The period covered corresponds to the Great Moderation period when the frequencies were stationary.

First, we consider the *estimated distribution* $\hat{\boldsymbol{\alpha}} = \{\hat{\alpha}_i\}_{i=1}^{44}$ which is estimated using the whole disaggregated microdata set. We have truncated the distribution at $F = 44$ months. We estimate the Survivor function $S(i)$ using the non-parametric Kaplan–Meier method as was used in Dixon and Le Bihan (2012) for French data and Dixon (2012) for the same UK data. This gives us the sequence of estimates of survival probabilities $\{\hat{S}(i)\}_{i=1}^{44}$, where $\hat{S}(i)$ is the estimated probability of a price-spell surviving at least $i$ periods. The corresponding *hazard* $\hat{h}(i)$, the probability that the spell will end in the $i$th period conditional on having survived $i$ periods, can be derived from the survivor function.[12] From this we are able to estimate the aggregate cross-sectional distribution using the steady-state identity:

$$\hat{\alpha}_i = \bar{h}.i.(\hat{S}(i) - \hat{S}(i+1))$$

Note the estimated distribution is based on all of the price spells across all COICOP sectors appropriately weighted: $\bar{h}$ is the aggregate Bernoulli–Calvo measure ($\bar{h} = 0.214$). We believe that the use of the steady-state identity is justified given that our time period is part of the great moderation in the UK.

We now turn to the *hypothetical distributions* derived from the same data set but at differing levels of aggregation, to see how the level of disaggregation affects the results. For the UK, we have the following levels of disaggregation available from the ONS, based on COICOP categories described below:

- 11 COICOP categories.
- 67 disaggregated COICOP categories.
- 570 items.

Each of these disaggregations represents exactly the same data. To get an idea of the level of aggregation, we can depict the broad 11 COICOP categories (excluding education which is not included in the ONS data set) in Table 1. For example, there is the category 'food and non-alcoholic beverages' which represents 17.6% of the CPI weight in the subsample available in the data set. The second level of disaggregation subdivides these into a total of 67 COICOP subcategories. For example, within 'food and non-alcoholic beverages' there are 11 subcategories: two for drinks (tea, coffee and cocoa; mineral water, soft drinks and juices) and nine for food (such as meat, fish, fruit). The lowest level of disaggregation is the item level. An item is a particular product or service on which the price observation is made. For example: canned sweet corn (198–340 g); coffee – take-away; fresh lettuce (iceberg). The 570 items we include are all of the items which were included throughout the sample period – it excludes old items which were either discontinued or new items introduced within this period. These items represent over 66.4% of the total CPI.

First, in Table 1 we present the appropriate CPI weight and the frequency[13] for the 11 COICOP sectors. In the first column of the Table 1 is the COICOP sector, in the second is

---

[12] $\hat{h}(i) = (\hat{S}(i) - \hat{S}(i+1))/\hat{S}(i)$.

[13] In a given month, the percentage of current prices (for items at a location) that are different from the price set in the previous month for the same product at the same location. The figure excludes items for which there was no observation the month before (e.g. it is the first price observation of the item at the outlet). The monthly frequencies are averaged over the whole data period.

TABLE 1

*COICOP 11 sectoral frequencies*

| COICOP Category. | CPI adj. wt | freq | Min (in months) | Max (in months) |
|---|---|---|---|---|
| Transport | 10.4% | 36.0% | 2.8 | 29.2 |
| Alcoholic beverages and tobacco | 7.1% | 27.6% | 3.6 | 32.9 |
| Clothing and footwear | 9.3% | 27.2% | 3.7 | 33 |
| Food and non-alcoholic beverages | 17.6% | 26.0% | 3.8 | 33.6 |
| Furniture and home maintenance | 11.3% | 22.7% | 4.4 | 35 |
| Communications | 0.2% | 22.5% | 4.4 | 35.1 |
| Recreation and culture | 9.9% | 20.0% | 5 | 36.2 |
| Housing and utilities | 8.3% | 13.7% | 7.3 | 39 |
| Miscellaneous goods and services | 6.5% | 12.7% | 7.9 | 39.4 |
| Restaurants and hotels | 17.5% | 10.5% | 9.5 | 40.4 |
| Health | 1.9% | 10.4% | 9.6 | 40.4 |
| Mean | | 21.4% | 4.7 | 35.6 |

*Notes*: 'Freq' denotes the frequencies of prices changes, which are reported in percent per month. 'CPI adj' denotes the adjusted CPI expenditure weight of the CPI sectors after excluding the Education sector. 'Min' denotes the minimum average duration. 'Max' denotes the maximum average duration with $F = 44$.

the CPI weight for the sector, normalized so that they add up to 100 (since Education is excluded) and the third the frequency (to three d.p.). In the fourth we have the minimum average duration in that sector and in the fifth the maximum from Proposition 1 based on the assumption that the longest price-spell is $F = 44$ months.

We next generate the cross-sectional distribution in the whole economy corresponding to the minimum average duration consistent with the observed frequencies. First we have the COICOP 11 disaggregation. From Proposition 1, in each sector we will have one or two durations with a non-zero share. In Recreation & Culture, since 20% of prices change per month, there will only be 5-month price-spells. In Food & Non-Alcoholic Beverages there will be a mixture of 20% 3-month and 80% 4-month price-spells to yield the mean duration of 3.8 months. The shortest durations are 2 months (in Transport) and the longest 10 months (in Health). For each duration, we can then add up across the 11 sectors to get the weighted cross-sectional distribution. We also perform the same procedure for the COICOP 67 and the 570 item level. These are all depicted in Figure 1, along with the estimated distribution.

The minimum duration distributions share some common features when compared with the estimated distribution: they all put too little weight on month 1, month 12 and longer durations. They all put too much weight on months 3–5 and months 9 and 10. However, the minimum duration distributions generated from different levels of aggregation are also quite different from each other. The level of disaggregation clearly matters when constructing a possible cross-sectional distribution. We can see this from the mean and median durations in Table 2.

These are far too short, reflecting the fact that the minimum duration distributions put a large weight on the shorter distributions and do not have a long fat tail as in the data

Figure 1. Estimated DAF vs. minimum duration distribution. 'ED' denotes the cross-sectional distribution implied by the estimated survivor function, 'MD11', 'MD67' and 'MD570' denote the cross-sectional distributions derived from the 'minimum method' at different disaggregation level, corresponding to 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively

TABLE 2

*The minimum mean and median duration across firms comparison*

|  | *Mean DAF (in months)* | *Median DAF (in months)* |
|---|---|---|
| ED | 10.9 | 7.8 |
| MD11 | 5.5 | 4 |
| MD67 | 6.1 | 4.5 |
| MD570 | 6.7 | 5.8 |

*Notes*: 'ED' denotes the cross-sectional distribution implied by the estimated survivor function, 'MD11', 'MD67' and 'MD570' denote the cross-sectional distributions derived from the 'minimum method' at different disaggregation level, corresponding to 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively. 'Mean DAF' and 'Median DAF' denote the mean and median length of duration across firms in months.

except for the most-disaggregated 570 item. In fact, the minimum durations are just over half the estimated mean duration.

Carvalho *et al.* (2015) estimate the mean cross-sectional quarterly distribution from US data on the basis of a variation the 'minimum method'. Using the Bils and Klenow (2004) sticky prices table of frequencies, they use the continuous time estimate of the mean duration, which is different from the discrete time Bernoulli formulation. The continuous time mean is less than the discrete time one, because it allows for durations in between the integer values. However, as in discrete time, the cross-sectional mean can only equal this in continuous time if all spells have the same length. Carvalho *et al.* (2015) use this as the basis for a prior and then go on to estimate the cross-section using Bayesian methods.

Next we look at the Bernoulli–Calvo distribution hypothetical distribution. Within each sector $k = 1 \ldots N$, we observe a sectoral frequency of $\bar{h}_k$ and the implied cross-sectional distribution for that sector $\boldsymbol{\alpha}_k^{\mathrm{BC}} = \{\alpha_{ik}^{\mathrm{BC}}\}_{i=1}^{\infty}$ given by equation (6). Each sector has a CPI weight $c_k$. We can then aggregate across the $N$ sectors using the CPI weights to get the share of each duration across all sectors $\boldsymbol{\alpha} = \{\alpha_i\}_{i=1}^{\infty}$ where:

$$\alpha_i^{\mathrm{BC}} = \sum_{k=1}^{N} c_k \alpha_{ki}^{\mathrm{BC}} \tag{8}$$

The mean duration of the Bernoulli–Calvo distribution at the sectoral level is equation (7), so that the mean of the aggregate distribution is then:

$$\bar{T}^{\mathrm{BC}} = \sum_{k=1}^{N} c_k \frac{2 - \bar{h}_k}{\bar{h}_k}$$

This is the method used in Dixon and Kara (2010, 2011) for generating the Bils–Klenow distribution based on the Bils and Klenow (2004) Appendix data set of 350 sectoral frequencies for the US. The Bernoulli process does not have a finite upper bound on possible durations: in empirical applications we will have to truncate it at some $F$, bundling together all of the theoretical durations of at least $F$ periods into duration $F$. In this paper we assume $F = 44$ (months) to be consistent with our estimated distribution.[14]

It is important to note that by assuming a Bernoulli–Calvo *distribution*, we are not assuming a Calvo *pricing model* within each sector. We are simply describing the distribution of price-spell durations in each sector generated by a constant Bernoulli measure (Calvo probability) that is equal to the sectoral frequency. This is purely descriptive of the distribution. It is perfectly compatible with a Taylor model, where within each sector the length of the price-spells is known *ex ante*. What we are doing is constructing $\boldsymbol{\alpha} = \{\alpha_i\}_{i=1}^{\infty}$ using equation (8): that means we take out all of the $i$ duration spells from each sector $k$ and put them together into a 'duration sector' $\alpha_i$, which includes all of the price-spells of length $i$ in the economy across all the COICOP sectors. The key difference between the Calvo and Taylor pricing frameworks is that under Taylor the firms know the length of the price-spell when they set the price, whereas in the Calvo pricing model they do not.

In Figure 2, we represent the Bernoulli–Calvo distributions at different levels of aggregation: the one sector 'aggregate BC' (ABC) distribution based on the mean UK frequency of 0.2140; the 11 sector COICOP, the 67 sector COICOP and the 570 item level. We have truncated the theoretical Bernoulli–Calvo distributions at 44 months.

The first observation is that the level of aggregation influences the shape of the aggregate distribution. The distributions all have a 'hump' shape, which peaks at 2 months (COICOP 67 and Item 570), 3 months (COICOP 11) and 4 months (ABC). They all have far too few 1-month shares and of course miss the 12-month spike, as pointed out by Alvarez and Burriel (2010). However, from month 8, COICOP 67 and Item 570 both track the estimated distribution fairly well (except for month 12). ABC and COICOP 11 both overestimate the share of durations between 3 and 16 months, and they underestimate the

---

[14] In Dixon and Kara (2010) we set $F = 60$, with the longest duration being 20 quarters.

Figure 2. Estimated DAF vs. Bernuli–Calvo distributions. 'ED' denotes the cross-sectional distribution implied by the estimated survivor function, 'ABC', 'BC11', 'BC67' and 'BC570' denote the cross-sectional distributions derived from the 'Bernoulli–Calvo distribution' at different aggregate level, corresponding to one aggregate sector, 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively

share of durations longer than 16 months. However, COICOP 11 is relatively closer to the estimated distribution than ABC. If we look at the Item 570 Bernoulli–Calvo distribution, the most disaggregated one we have, this peaks at month 2 and is the only Bernoulli–Calvo distribution to be roughly close to (a little less than) the estimated proportion of 1 months. Furthermore, the Bernoulli–Calvo distribution generated by Item 570 is quite similar from the months 2 onwards, only missing the 12 months spike. The UK distribution has a fatter tail, but the Bernoulli–Calvo tails are certainly quite substantial. Essentially, the Bernoulli–Calvo distributions put too much weight on the shorter months (2–7) and hence puts less weight on the remaining durations. The level of disaggregation matters for the mean, since from equation (7) the mean of the DAF is a convex function of the frequency $\bar{h}$: hence with sectoral heterogeneity the mean will increase at more disaggregated levels due to Jensen's inequality.[15]

The means and medians of the different Bernoulli–Calvo distributions are listed in Table 3. Here we can see that as the category becomes more disaggregated, the mean and median of Bernoulli–Calvo distributions become closer to the estimated distribution. Indeed, the Bernoulli–Calvo distribution generated by Item 570 almost has the same mean and median value as what we get from the estimated distribution. If we compare the means of the Bernoulli–Calvo distributions, these are linked to the minimum distributions in Table 1, since the mean of the minimum duration distribution is

$$\bar{T}^{\min} = \sum_{k=1}^{N} c_k \bar{h}_k^{-1}$$

---

[15] This is of course also true for the minimum distributions in Table 2, since the mean DAF is given by Proposition 1 is also convex in $\bar{h}$.

TABLE 3

*Mean and median durations of Calvo distributions*

|  | Mean DAF (in months) | Median DAF (in months) |
|---|---|---|
| ED | 10.9 | 7.8 |
| BC11 | 9.1 | 6.8 |
| BC67 | 10.5 | 7.2 |
| BC570 | 10.8 | 7.8 |
| ABC | 8.3 | 6.5 |

*Notes*: 'ED' denotes the estimated cross-sectional distribution, 'ABC', 'BC11', 'BC67' and 'BC570' denote the cross-sectional distributions derived from the 'Bernoulli–Calvo distribution' at different aggregate level, corresponding to one aggregate sector, 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively. 'Mean DAF' and 'Median DAF' denote the mean and median length of duration across firm in months.

which yields the theoretical relation $\bar{T}^{\mathrm{BC}} = 2.\bar{T}^{\min} - 1$. Hence the Bernoulli–Calvo means are similar to the estimated mean, whilst the theoretical minimum is just over half. No such exact relation holds for the medians. If we look at Tables 2 and 3, we can see that the Bernoulli–Calvo means are less than $2.\bar{T}^{\min} - 1$ and equation (7). This is because we have *truncated* the Bernoulli–Calvo distributions at 44 months: if we extend this then the mean will approach its theoretical value. Truncation reduces the mean quite significantly, since the long tail of the Bernoulli–Calvo distribution will be allocated to the 44th month: whilst the shares of these longer durations are small, they are also long and so affect the mean.

Whilst the Bernoulli–Calvo assumption gives us a mean that is about right, it still differs considerably from the estimated distribution. There are not enough one-period price-spells: in the data, there are a lot of products that have perfectly flexible prices that change almost every month (petrol, vegetables, etc.). Hence the one period share $\alpha_1^{\mathrm{BC}}$ needs to be higher than in the standard Bernoulli–Calvo case. Second, there are too many 3- to 8-month spells. Then of course there is a 12-month peak in the estimated distribution which is absent from the smooth Bernoulli–Calvo distribution.

In this section, we have looked at the aggregate economy. We can also look at the distributions within each of the 11 COICOP sectors.[16] There is considerable heterogeneity in the sectoral distributions. Since we have such a large data set, the formal Kolmogorov–Smirnov test[17] rejects the null hypothesis that the estimated distribution and the Bernoulli–Calvo distributions are the same in all of the 11 sectors. In conclusion, we can say that the Bernoulli–Calvo distribution is not a good description of the data either at the aggregate level or the COICOP 11 level.[18]

[16] For details see an earlier version of this paper, CESifo working paper 4226, especially Appendix 3.

[17] Best and Rayner (2003) examine tests for the geometric distribution, which is the distribution of $\{\alpha_i^{\mathrm{d}}\}$ with a constant hazard. These tests all rely on the Kolmogorov–Smirnov statistics. Whilst the tests differ in small sample sizes, since we have a large sample size these they are all equivalent and give the same results as the Kolmogorov–Smirnov test we used. We would like to thank a referee for pointing this out to us.

[18] Fougere, Le Bihan and Sevestre (2007) adopt a Wald test on the hazard function estimates on sectors that are as disaggregated as possible (using French price quote data). They find that in 35% of sectors by CPI weight (Table 2) the null of a constant hazard rate cannot be rejected. Our test is always at the aggregate distribution using

## IV.    The simulation of different pricing models

We have found that we can use the sectoral frequencies to generate the corresponding hypothetical Bernoulli–Calvo distributions. For the UK data at least, we find that at high levels of disaggregation, the resultant hypothetical aggregate distribution matches the estimated distribution quite well in terms of both the mean and the median. There are significant differences, most notably the hypothetical distribution has no 12-month spike and too few flexible prices. Since the mean and median are close, do the differences matter at the aggregate level? If we simulate a Dynamics Stochastic General Equilibrium (DSGE) macro model using the hypothetical Bernoulli–Calvo distribution, will it yield a good approximation to the simulations using the estimated distribution found in the UK data? If the answer is 'yes', then it implies that the absence of the 12-month spike and too few flexible prices does not matter from the perspective of the macroeconomic properties of the DSGE model. This would validate the approach taken in Dixon and Kara (2010, 2011) and Kara (2010) which used the hypothetical Bernoulli–Calvo distribution derived from the Bils–Klenow table of sectoral frequencies in order to calibrate their US pricing models.

We will perform our simulations using two DSGE models: a simple QT model and the Smets and Wouters (2003) (SW) model. We will look at two pricing models in both of these cases: the GT and GC model as in Dixon and Le Bihan (2012). We will also consider both models using quarterly and monthly calibrations.

### Price setting

There are two general time-dependent models which are capable of reflecting the underlying distribution found in the micro-data: the GT[19] and GC models.[20] The key difference between the models is that in the GT the firms know how long the price spell will last when they set the price, and so each duration of price-spell will have a different reset price. In the GC, in contrast, the firms do not know how long the price spell will last and have a distribution over possible price-spells durations. All firms have the same distribution and hence there is only one reset price every period as in the simple Calvo model. In the *Generalized Taylor Economy* (*GT*) there are $N$ sectors, $i = 1,...,N$. In sector $i$ there are $i-$period contracts: each period a cohort of $i^{-1}$ of the firms in the sector sets a new price (or wage). If we think of the economy as a continuum of firms, we can describe the *GT* as a vector of sector shares: $\alpha_i$ is the proportion of firms that have price-spells of length $i$. If the longest observed price-spell is $F$, then we have $\sum_{i=1}^{F} \alpha_i = 1$ and $\alpha \in \Delta^{F-1}$ is the $F$-vector of shares $\boldsymbol{\alpha} = \{\alpha_i\}_{i=1}^{F}$. We can think of the 'sectors' as 'duration sectors': we can classify the economy by the length of price-spells. The essence of the Taylor model is that when they set the price, firms know exactly how long its price is going to last. The simple Taylor economy is a special case where there is only one length of price-spell (e.g. $\alpha_2 = 1$ is a simple Taylor '2 quarters' economy).

The log-linearized equation for the aggregate price $p_t$ is a weighted average of the sectoral prices $p_{it}$, where the weights are $\alpha_i$:

---

$$p_t = \sum_{i=1}^{F} \alpha_i p_{it} \tag{9}$$

In each sector $i$, a proportion $i^{-1}$ of the $\alpha_i$ firms reset their price at each period. Assuming imperfect competition and standard demand curve, the optimal reset price in sector $i, x_{it}$ is given by the first-order condition of an intertemporal profit-maximization program under the constraint implied by price rigidity. The log-linearized equation for the reset price, as in the standard Taylor set-up, is then given by:

$$x_{it} = \left( \frac{1}{\sum_{k=0}^{i-1} \beta^k} \right) \sum_{k=0}^{i-1} \beta^k E_t p_{t+k}^* \tag{10}$$

where $\beta$ is a discount factor, $E_t$ is the expectation operator conditional on information available at date $t$ and $p_{t+k}^*$ is the optimal flex price at time $t+k$. The reset price is thus an average over the optimal flex prices for the duration of the contract (or price-spell). The formula for the optimal flex price will depend on the model: clearly, it is a markup on marginal cost. We will specify the exact log-linearized equation for the optimal flex-price when we specify the exact macroeconomic model we use.

The sectoral price is simply the average over the $i$ cohorts in the sector:

$$p_{it} = \frac{1}{i} \sum_{k=0}^{i-1} x_{it-k} \tag{11}$$

In each period, a proportion $\bar{h}$ of firms reset their prices in this economy: proportion $\iota^{-1}$ of sector $i$ which is of size $\alpha_i$ resulting in equality (3).

In the GC, firms have a common set of duration-dependent reset probabilities: the probability of resetting price $i$ periods after you last reset the price is given by $h_i$. This is a time-dependent model, and the profile of reset probabilities is $h = \{h_i\}_{i=1}^{F}$. Clearly, if $F$ is the longest price-spell we have $h_F = 1$ and $h_i \in [0, 1)$ for $i = 1 \ldots F - 1$. Again, the duration data can be represented by the hazard function. The estimated hazard function can then be used to calibrate $h$. Since any distribution of durations can be represented by the appropriate hazard function, we can choose the GC to exactly fit micro-data.

In economic terms, the difference between the Calvo approach and the Taylor approach is that when the firm sets its price, it does not know how long its price is going to last. Rather, it has a *survivor function* $S(i)$ which gives the probability that its price will last at up to $i$ periods. The survivor function in discrete time is:[21]

$$S(1) = 1$$
$$S(i) = \prod_{j=1}^{i-1} (1 - h_j) \quad i = 2,\ldots,F \tag{12}$$

Thus, when they set the price in period $t$, the firms know that they will last one period with certainty, at least two periods with probability $S(2)$ and so on. The Calvo model is a

---

[21] Note that the discrete time survivor function effectively assumes that all 'failures' occur at the end of the period (or the start of the next period): this corresponds to the pricing models where the price is set for a whole period and can only change at the transition from one period to the next.

special case where the hazard is constant $h_i = \bar{h}$, $S(i) = (1 - \bar{h})^{i-1}$ and $F = \infty$. Of course, in any actual data set, $F$ is finite.

In the GC model the reset price is common across all firms that reset their price. The optimal reset price, in the same monopolistic competition set-up as mentioned above, is given in log-linearized form by:

$$x_t = \frac{1}{\sum_{i=1}^{F} S(i)\beta^{i-1}} \sum_{i=1}^{F} S(i)\beta^{i-1} E_t p_{t+i-1}^* \tag{13}$$

The evolution of the aggregate price-level is given by:

$$p_t = \sum_{i=1}^{F} S(i) x_{t-i+1} \tag{14}$$

That is, the current price level is constituted by the surviving reset prices of the present and previous $F - 1$ periods.

**A simple Quantity Theory model with price-setting**

We will first examine the GC and GT models of prices in a QT model with labour as the only input of production. This model has the great advantage of being very simple, so that almost all its dynamic properties are generated by the pricing models alone. DSGE models like the SW model in contrast are quite complicated with dynamic properties emerging from the interaction of pricing with many other features of the model. The model we present is in its log-linearized version (see Ascari, 2003, Dixon and Kara, 2010 for the derivation from microeconomic foundation). We present both a monthly and a quarterly calibrations. We look at how the model behaves using the estimated distribution, denoting resultant IRFs as ED. We can then compare ED to the IRFs derived from the frequency data using the Bernoulli–Calvo (these IRFs denoted as BC11, BC 67 and BC570) and 'minimum duration' distributions (denoted as MD11, MD67 and MD570).

To model the demand side, we use the constant-velocity Quantity Theory:

$$y_t = m_t - p_t$$

where $(p_t, y_t)$ are aggregate price and output and $m_t$ the money supply. We model the monetary growth process as an autoregressive process of order one $AR(1)$:

$$m_t = m_{t-1} + \varepsilon_t$$
$$\varepsilon_t = v\varepsilon_{t-1} + \xi_t$$

where $\xi_t$ is a white noise error term (effectively a monetary growth shock). Following Chari, Kehoe and McGrattan (2000), we set $v = 0.5$ for the quarterly model and the equivalent of 0.85 in the monthly model.[22]

---

[22]This calibration is also used in Dixon and Kara (2010) and Mankiw and Reis (2002).

The optimal flexible price $p_t^*$ at period $t$ in all sectors is given by:

$$p_t^* = p_t + \gamma y_t \qquad (15)$$

The key parameter $\gamma$ captures the sensitivity of the flexible price to output.[23] As discussed in Dixon and Kara (2010), there are a range of calibrated and estimated values for $\gamma$: for illustrative purposes, we use the 'moderate' case of $\gamma = 0.1$ as in Mankiw and Reis (2002). As shown by Ascari (2003) and Edge (2002), the value of $\gamma$ can be interpreted as resulting from either wage or price-setting. The value of $\gamma$ is common across the monthly and quarterly models.

In Figure 3a–d, we see the IRFs for the quarterly and monthly QT model responding to a monetary 1% monetary growth shock derived from the Bernoulli–Calvo and the minimum duration distributions. There are *seven* IRFs: the IRF ED generated from the estimated UK distribution (ED), and the IRFs generated from the three Bernoulli–Calvo (BC) and minimum duration (MD) distributions derived from the sectoral frequencies at different levels of aggregation.

Turning first to the Bernoulli–Calvo IRFs, we can summarize in a few points:

(a) The IRFs for output and inflation BC11, BC67 and BC570 are similar to ED for both pricing models (GC and GT) and calibrations (monthly and quarterly). In the monthly calibration they are very close.

(b) The GT has a hump-shaped reaction function for inflation, the GC does not (monthly and quarterly).

Turning to the minimum distribution IRFs, we can summarize:

(c) For the GT pricing model, the IRFs for output and inflation MD11, MD67 and MD570 are quite similar to each other, but significantly different to ED for both the quarterly and monthly calibration. The MD IRFs all posses a hump shape for output and inflation. The monthly calibration results in a jagged MD11 and MD67 for inflation, which reflects the irregular nature of the underlying minimum distributions at the more aggregated levels.

(d) For the GC pricing model, the impact effect of all MD is closer to ED than all BC. However, after the first period BC are closer than MD. The MD all display a hump shape for inflation (monthly and quarterly) that is absent in the BC (and ED).

It is rare for a GC to display a hump shape for inflation. The reason we find it here is that the minimum distributions are *myopic* relative to the Bernoulli–Calvo and estimated distributions. Myopia means that the reset price puts a greater weight on the short-run and hence tends to adjust less on impact, potentially leading to a delayed hump-shape response. Because the expected duration of a price-spell is relatively low in the minimum distributions, it results in myopia despite the forward looking 'infinite horizon' feature of the GC pricing model. Sheedy (2010) showed that a hump-shaped IRF could result with GC pricing if the hazard function was upward sloping.

---

[23] This can be due to increasing marginal cost and/or an upward sloping supply curve for labour. See for example Walsh (2003) chapter 5) and Woodford (2003) chapter 3).
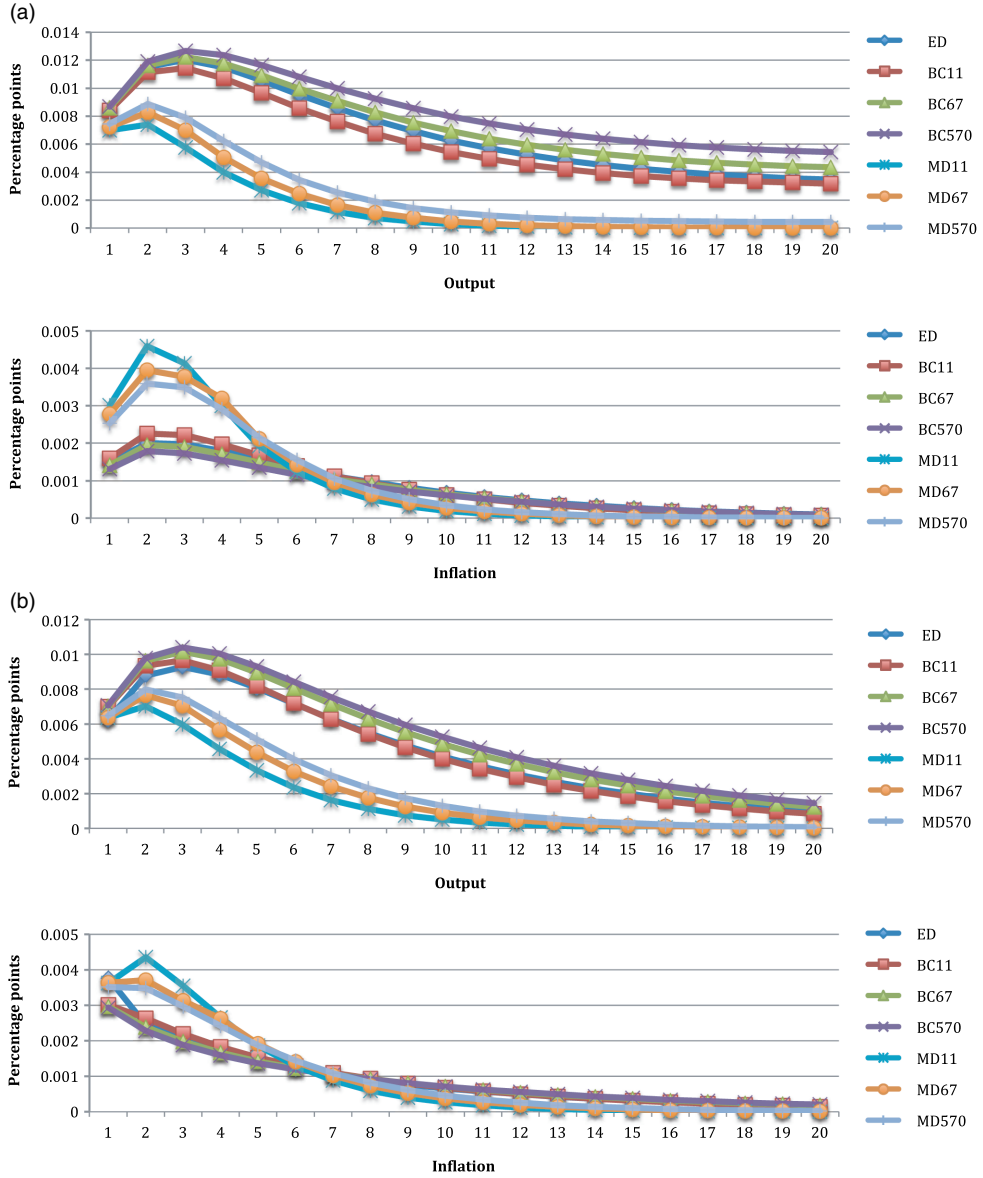
Figure 3. (a) Impulse response functions (IRFs) for monetary shock in a quarterly Quantity Theory (QT) model with Generalized Taylor (GT) pricing. (b) IRFs for monetary shock in a quarterly QT model with Generalized Calvo (GC) pricing. (c) IRFs for monetary shock in a monthly QT model with GT pricing. (d) IRFs for monetary shock in a monthly QT model with GC pricing

In order to more explicitly quantify differences between the IRFs, we define the point-by-point absolute difference between ED and the particular BC IRFs as a percentage of the mean ED: $\delta_i = \frac{|IRF_{ED} - IRF_{BC}|}{|mean(IRF_{ED})|} \times 100\%$. We then find the average of these differences over 5 years (20Q or 60 months) to get the *average relative difference* (ARD),[24] which is shown

---

[24] For robustness we also calculated Tables 4–7 using the Average absolute differences and Quadratic differences. The results were consistent with existing Tables 4–7 using ARD. The results are available from the authors on request.
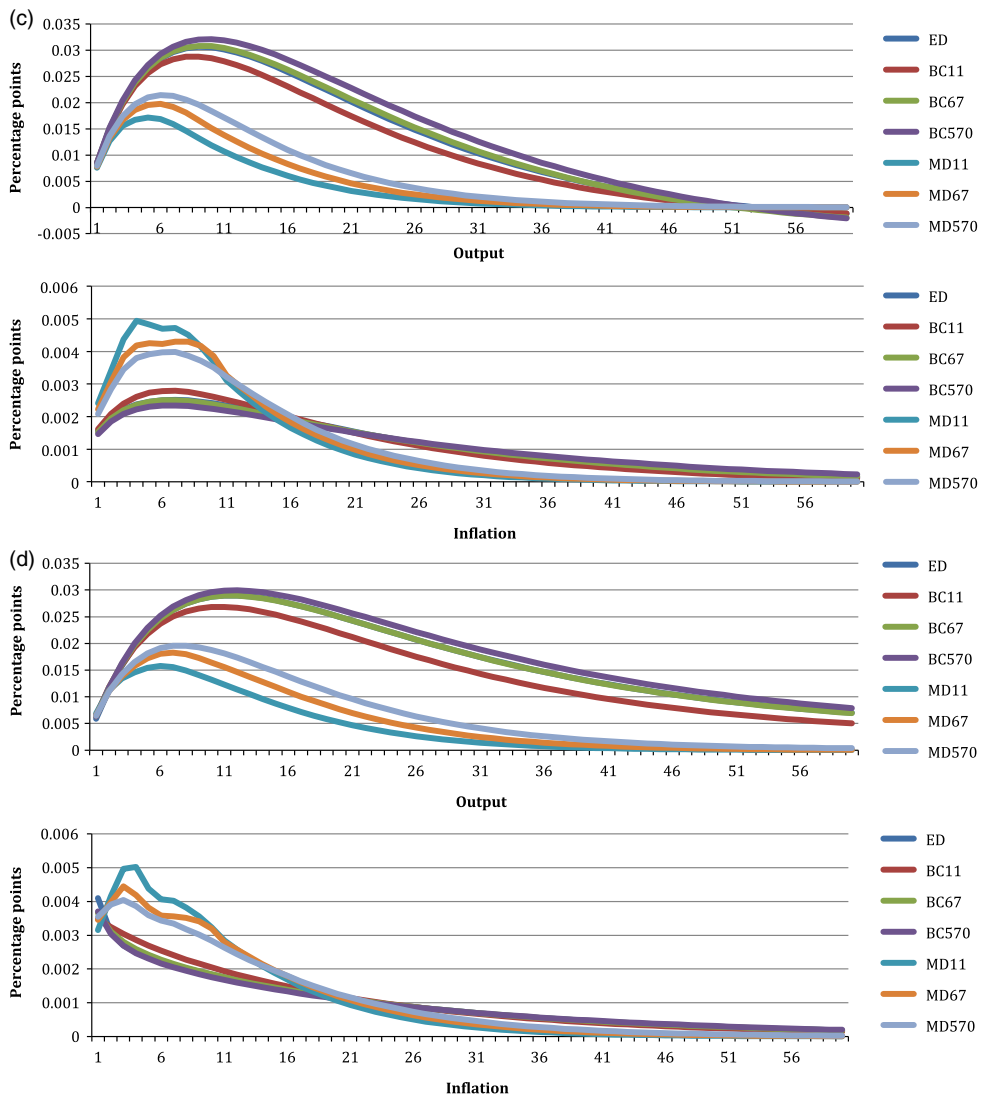
Figure 3.   (Continued)

in Table 4.[25] The *ARD* appear as percentages. There are three columns for each pricing model (GT and GC), with two rows (output and inflation) for each calibration (quarterly and monthly).

Here we can see that with the GT model, BC67 is the one closest to ED for both calibrations. However, in the GC model, the results are more mixed. For the quarterly IRF

[25] For the quarterly model, we have

$$ARD = \frac{1}{20} \sum_{i=1}^{20} \delta_i$$

and analogously for the monthly calibration.

TABLE 4

*The average relative difference AD in IRF from QT model*

| QT | GT | | | GC | | |
|---|---|---|---|---|---|---|
| *ARD* | *BC11* | *BC67* | *BC570* | *BC11* | *BC67* | *BC570* |
| Quarterly output | 9.12% | 8.73% | 21.82% | 4.34% | 14.41% | 21.42% |
| Quarterly inflation | 9.93% | 5.36% | 11.87% | 9.01% | 7.99% | 10.49% |
| Monthly output | 11.23% | 2.54% | 11.26% | 14.01% | 0.32% | 6.50% |
| Monthly inflation | 8.48% | 2.43% | 8.98% | 9.50% | 1.69% | 3.35% |

*Notes*: 'BC11', 'BC67' and 'BC570' are the IRFs derived from the cross-sectional Bernoulli–Calvo distribution at different levels of aggregation, corresponding to 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively.

in output, the BC11 is the one having the smallest *ARD*. However, in both the quarterly and monthly calibrations, BC67 is the best.

If we compare the *ARD* in the quarterly and the monthly calibrations, we can see that the differences between ED and BC67 are smaller in the monthly case. For the GT, both quarterly and monthly, BC11 comes second and BC570 last. However, for GC matters are reversed with the monthly calibration: BC570 does better than BC67. For quarterly GC, BC11 outperforms BC570 and indeed has the lowest *ARD* for output.

Macroeconomic models are generally quarterly, which indicates that for inflation, for both GT and GC the intermediate level of disaggregation BC67 is closest to ED. For output, this also holds for GT, but in the case of GC the BC11 performs better. The most disaggregated level can perform quite badly: for both the GT and GC, the *ARD* with BC570 takes relatively high values over 20% for output. This is because the degree of persistence is far too great. Thus, although the 570 item Bernoulli–Calvo distribution has the closest mean and median to the estimated cross-sectional distribution, this does not imply that the resultant IRFs are closer to those generated by the estimated distribution in the simple QT model. For the GT, both BC11 and BC67 are within 10% *ARD* for both output and inflation. For the GC, only BC11 is within 10% for both output and inflation. To see the differences between monthly and quarterly calibrations, if we compare the 12 panels for monthly with the 12 panels for the quarterly calibration, in 9 out of 12 pairwise comparisons the *ARD* is smaller for the monthly calibration. For example, we can compare the panels for GT using BC67 for the output of the IRF: the monthly *ARD* is 2.54% as compared with the quarterly 8.73%.

Lastly, we perform the same exercise for the QT model using the minimum duration distributions derived in Proposition 1. The IRFs MD11, MD67 and MD570 are a long way from the ED, with very high *ARDs*, as illustrated in Table 5. We can see that the *ARD* is decreasing the more disaggregated we get and that MD 11, MD67 and MD570 work better for the GC than the GT. The poor performance of the minimum duration calibration shows that getting the microeconomic distribution wrong can have major implications for the way macroeconomic models behave. The minimum duration distributions lack a long-fat tail and give rise to far to little persistence in the IRFs, which results in the large *ARD*.

TABLE 5

*The average relative difference AD in IRF from QT model*

| QT | GT | | | GC | | |
|---|---|---|---|---|---|---|
| ARD | MD11 | MD67 | MD570 | MD11 | MD67 | MD570 |
| Quarterly output | 76.76% | 71.81% | 62.39% | 61.51% | 52.29% | 45.10% |
| Quarterly inflation | 74.19% | 64.84% | 55.28% | 47.09% | 37.55% | 32.04% |
| Monthly output | 70.25% | 62.34% | 53.82% | 74.93% | 67.57% | 59.65% |
| Monthly inflation | 61.51% | 54.05% | 45.71% | 58.17% | 48.64% | 40.16% |

*Notes*: 'MD11', 'MD67' and 'MD570' denote the IRFs corresponding to the cross-sectional distributions derived from the minimum duration distribution at different aggregate level, corresponding to 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively.

## A DSGE model: Smets and Wouters (2003)

In this section, we use the Smets and Wouters (2003) model of the euro area commonly employed for monetary policy analysis to compare the different calibrations. The SW model is much more complicated than the simple QT model we have just used: there are many sources of dynamics other than prices and wages, including capital adjustment, capital utilization, consumer dynamics with habit formation and a monetary policy reaction function. The behaviour of the model is the outcome of the interaction of all these processes together as it should be in a DSGE model. Hence the effect of pricing dynamics is not isolated as in the simple QT framework of the previous section. The details of the model and calibration are outlined in Dixon and Le Bihan (2012) using a notation consistent with this paper, and the Dynare program can be downloaded from huwdixon.org. The original SW model is quarterly: however, we calibrate an equivalent monthly version of the model.[26]

We depict the IRFs for an interest rate shock, which causes output and inflation to fall initially (as shown in Figure 4 for the quarterly model with GC and GT pricing).[27] Here we see that the differences between the BC11, 67 and 570 and ED are much smaller and less visible when compared with the QT model. This is probably because the structure on the dynamics is also determined by the rest of the model's complex dynamics, which leaves less room for the precise distribution of price-spell durations to matter. The *ARD* are shown in Table 6 and are quite small compared to the *ARD* in the QT model.

We can see that for both the GT and GC model, the intermediate level of disaggregation (COICOP 67) always yields the lowest *ARD*. The quarterly GT calibration has a lower *ARD* for all cases when compared with the monthly. For the GC, it is almost the opposite – the monthly calibration has *lower* values than the quarterly except for inflation with 570 item. If we compare the SW simulations with the QT, we can see that in general the differences are much smaller. The big exception is for inflation with GT pricing in the monthly calibration where the *ARD* is higher in the SW model than the QT.

Turning next to the IRFs MD11, 67 and 570, in Table 7 we can see that in the SW model the *ARD*s are much bigger than for any of the BC cases. In all cases, MD570 gives the lowest *ARD*: for both GT and GC, monthly and quarterly, output and inflation. The only

[26]The model and calibrations are summarized in Appendix 3.

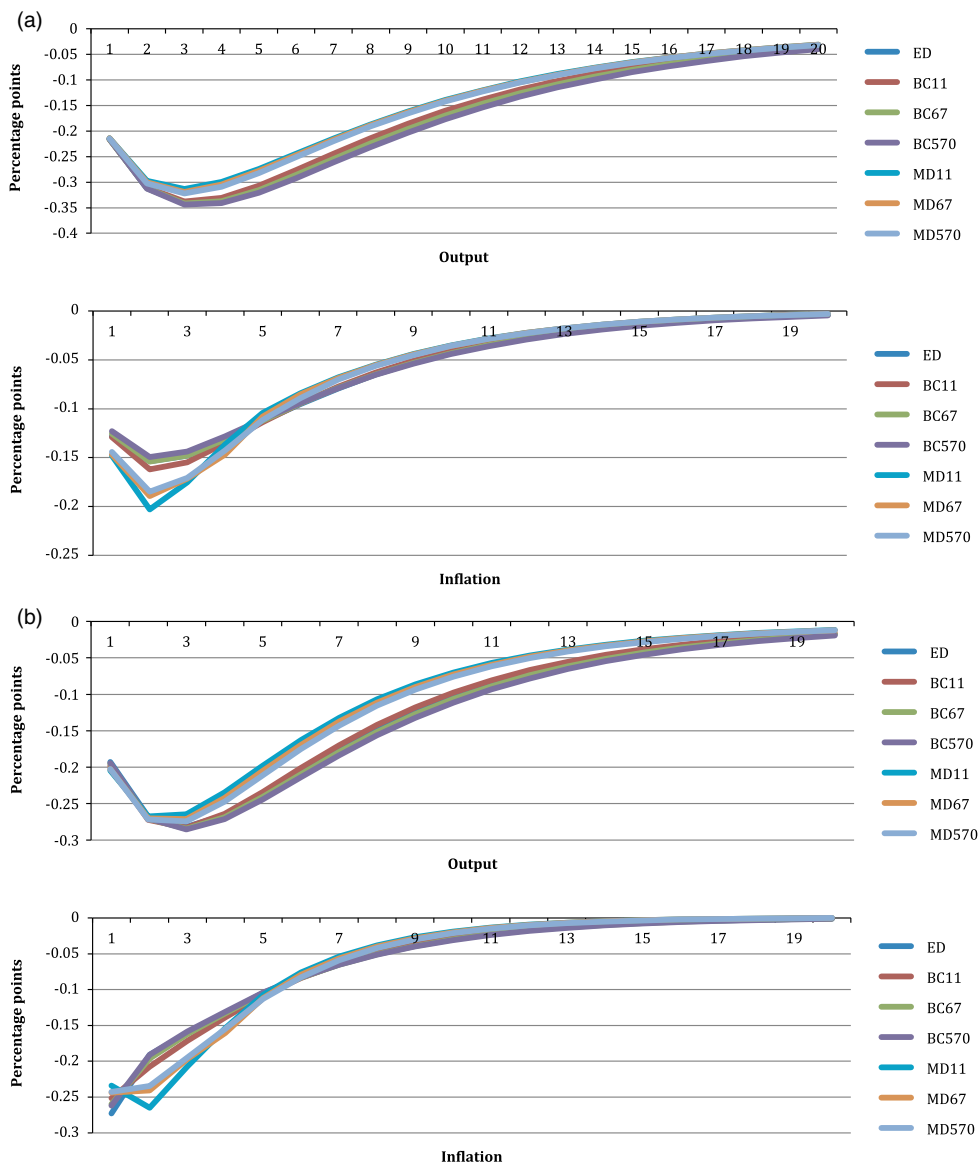[27]The IRFs for the monthly data are given in Figure A1 in the Appendix.

Figure 4. (a) Impulse response function (IRF) for monetary shock in a quarterly Smets–Wouters (SW) model with Generalized Taylor (GT) pricing. (b) IRF for monetary shock in a quarterly SW model with GC pricing

cases that give an *ARD* below 5% are output with GC pricing for monthly data. Out of the remaining 21 panels, 17 are less than 20% and 3 more than 30% (the remaining one at 23%).

If we take the results from the simulations of both the QT model and the SW model, we can see that small differences in the distribution of durations need not matter that much when it comes to the behaviour of the macroeconomic model in terms of IRFs. Within the three Bernoulli–Calvo distributions, the resultant IRFs BC11, BC67 and BC570 are fairly similar across a range of possibilities (output inflation, GT and GC pricing,

TABLE 6

*The average relative difference AD in IRF from SW model*

| SW | GT | | | GC | | |
|---|---|---|---|---|---|---|
| ARD | BC11 | BC67 | BC570 | BC11 | BC67 | BC570 |
| Output quarterly | 2.64% | 0.87% | 3.54% | 3.71% | 1.28% | 3.58% |
| Inflation quarterly | 3.10% | 0.91% | 3.28% | 6.12% | 2.06% | 2.34% |
| Output monthly | 2.96% | 1.21% | 4.20% | 1.11% | 0.07% | 0.46% |
| Inflation monthly | 8.71% | 4.35% | 13.87% | 4.56% | 1.38% | 1.60% |

*Notes*: 'BC11', 'BC67' and 'BC570' denote the IRFs implied by the cross-sectional distributions derived from the 'Bernoulli–Calvo distribution' at different aggregate level, corresponding to 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively.

TABLE 7

*The average relative difference AD in IRF from SW model*

| SW | GT | | | GC | | |
|---|---|---|---|---|---|---|
| ARD | MD11 | MD67 | MD570 | MD11 | MD67 | MD570 |
| Output quarterly | 12.33% | 11.33% | 10.39% | 18.94% | 16.73% | 14.90% |
| Inflation quarterly | 16.07% | 14.71% | 12.67% | 22.68% | 18.92% | 16.90% |
| Output monthly | 12.23% | 11.67% | 11.01% | 4.41% | 4.12% | 3.73% |
| Inflation monthly | 33.37% | 33.91% | 32.37% | 14.71% | 13.19% | 12.06% |

*Notes*: 'MD11', 'MD67' and 'MD570' denote the IRFs corresponding to the cross-sectional distributions derived from the 'minimum distribution' at different levels of aggregation, corresponding to 11 COICOP categories, 67 disaggregated COICOP categories and 570 items respectively.

monthly and quarterly). The distribution closest to the estimated distribution (570 item) does not necessarily give the best performance in terms of IRF: BC11 performs better than BC570.

However when we look at distributions that are substantially different, the resultant IRFs will also differ. The three minimum distributions have means and medians that are much lower than the estimated distribution. The resultant IRFs are also very different in almost all cases. The poor performance of MD11, 67 and 570 suggests that this is not a good calibration: using it as a prior in a Bayesian estimation as in Carvalho *et al.* (2015) is perhaps not a good place to start. Using the Bernoulli–Calvo distribution ensures that we have a result closer to the estimated distribution, at least for UK data.

## V. Conclusion

In this paper, we asked the question what can the sectoral data on the frequency of price-change tell us? On the theoretical level, sectoral frequencies tell us what expected duration of a price-spell is. This is of some interest, but from a macroeconomic perspective we are more interested in the behaviour of the economic agents setting prices – the *cross-sectional* distribution is much more informative. Unfortunately the frequency itself says little about the cross-sectional distribution: to uncover this we need to make additional assumptions.

However, we are able to say what the theoretical minimum cross-sectional mean duration is consistent with an observed frequency: it is the mean duration of a price-spell which occurs when all price-spells in the sector have the same or almost the same length. However, the cross-sectional mean can be much longer: intuitively, firms that have prices that last for a long time do not reset their prices very often. These firms contribute little to the monthly frequency but can add a lot to the cross-sectional mean.

When we look at the UK data using an estimated hazard function, we find that the UK data is a long way from the distribution implied by the theoretical minimum. Looking at different levels of disaggregation, we find that whilst the minimum theoretical mean duration is around 5.5–6.7 months, the mean estimated from the data is almost twice as long as that at 11 months. We also look at the aggregate data using the hypothesis that the sectoral frequencies are generated by a Bernoulli–Calvo distribution. Under this assumption, the cross-sectional mean is much larger than the minimum, and gets closer to the estimated mean as you become more disaggregated, with the most disaggregated having almost exactly the same mean and median as estimated from the data. Whilst the mean and median of the hypothetical Bernoulli–Calvo distribution can be close to the estimated values, the shape differs in two distinct ways: first, there is no 12-month spike, secondly there are not enough flexible prices.

However, the hypothetical distributions are similar to the estimated distribution in a significant manner: unlike the minimum duration distribution, they also have a long fat tail of long durations, which can generate substantial persistence in the response of output and inflation to monetary policy shocks (Dixon and Kara, 2011). Hence, although we do find that whilst the Bernoulli–Calvo distribution hypothesis differs from the estimated distribution both at the aggregate and sectoral levels, it can nonetheless 'work' at the aggregate level when it is used to calibrate DSGE models. However, the level of disaggregation at which the frequencies are measured seems to have a major effect. In our examples, the intermediate level of disaggregation yielded the best IRFs, better than the most disaggregated. This suggests that when we do not have reliable hazard function estimates or access to the price microdata, we can use the disaggregated frequency data to calibrate DSGE models. The theoretical minimum distribution is a long way from the estimated distribution in the case of the UK data. However, the Bernoulli–Calvo distribution seems to provide a calibration closer to the estimated one.

# References

Alvarez, L. and Burriel, P., (2010). 'Is a Calvo price setting model consistent with individual price data?', *The B.E. Journal of Macroeconomics, Berkeley Electronic Press*, Vol. 10, pp. 13.

Alvarez, L., Dhyne, E., Hoeberichts, M., Kwapil, C., Le Bihan, H., Lunnemann, P., Martins, F., Sabbatini, R., Stahl, H., Vermeulen, P. and Vilmunen, J. (2006). 'Sticky prices in the euro area: a summary of new micro-evidence', *Journal of the European Economic Association*, Vol. 4, pp. 575–584.

Alvarez, L. and Hernando, I (2006). 'Price setting behaviour in Spain: stylised facts using consumer price micro data', *Economic Modelling*, Vol. 23, pp. 699–716.

Alvarez, L. and Lippi, F. (2014). 'Price setting with menu cost for multiproduct firms', *Econometrica*, Vol. 82, pp. 89–135.

Ascari, G. (2003). 'Price/wage staggering and persistence: a unifying framework', *Journal of Economic Surveys, Wiley Blackwell*, Vol. 17, pp. 511–540.

Baharad, E. and Eden, B. (2004). 'Price rigidity and price dispersion: evidence from micro data', *Review of Economic Dynamics*, Vol. 7, pp. 613–641.

Baudry, L., Le Bihan, H, Sevestre, P. and Tarrieu, S. (2007). 'What do thirteen million price records have to say about consumer price rigidity?', *Oxford Bulletin of Economic Statistics*, Vol. 69, pp. 139–183.

Best, D. J. and Rayner, J. C. W. (2003). 'Tests of fit for the geometric distribution', *Communications in Statistics-Simulation and Computation*, Vol. 32, pp. 1065–1078.

Bils, M. and Klenow, P (2004). 'Some evidence on the importance of sticky prices', *Journal of Political Economy*, Vol. 112, pp. 947–985.

Bunn, P. and Ellis, C. (2009). 'Price setting behavior in the UK', *Bank of England Quarterly Bulletin*, Vol. 49, pp. 28–36.

Bunn, P. and Ellis, C. (2012). 'Examining the behaviour of individual UK consumer prices', *Economic Journal*, Vol. 122, pp. F35–F55.

Carvalho, C., Dam, N. and Lee, J. (2015). *Real Rigidities and the Cross-sectional Distribution of Price Stickiness: Evidence from Micro and Macro Data Combined*, Danmarks Nationalbank working paper 93.

Chari, V. V., Kehoe, P. J. and McGrattan, E. R. (2000). 'Sticky price models of the business cycle: can the contract multiplier solve the persistence problem?', *Econometrica*, Vol. 68, pp. 1151–1179.

Coenen, G., Levin, A.T. and Christoffel, K. (2007). 'Identifying the influences of nominal and real rigidities in aggregate price-setting behavior', *Journal of Monetary Economics*, Vol. 54, pp. 2439–2466.

Dixon, H. and Kara, E. (2006). 'How to compare Taylor and Calvo contracts: a comment on Michael Kiley', *Journal of Money, Credit and Banking*, Vol. 38, pp. 1119–1126.

Dixon, H. and Kara, E. (2010). 'Can we explain inflation persistence in a way that is consistent with the micro-evidence on nominal rigidity', *Journal of Money, Credit and Banking*, Vol. 42, pp. 151–170.

Dixon, H. and Kara, E. (2011). 'Contract length heterogeneity and the persistence of monetary shocks in a dynamic Generalized Taylor economy', *European Economic Review*, Vol. 55, pp. 280–292.

Dixon, H. (2012). 'A unified framework for using micro–data to compare dynamic time-dependent price-setting models', *BE Journal of Macroeconomics* (Contributions), Vol. 12, pp. 1–45.

Dixon, H. and Le Bihan, H. (2012). 'Generalized Taylor and Generalized Calvo price and wage-setting: micro evidence with macro implications', *Economic Journal*, Vol. 122, pp. 532–554.

Edge, R. (2002). 'The equivalence of wage and price staggering in monetary business cycle models', *Review of Economic Dynamics*, Vol. 5, pp. 559–585.

Fougere, D., Le Bihan, H. and Sevestre, P. (2007). 'Heterogeneity in price stickiness: a microeconometric investigation', *Journal of Business & Economic Statistics*, Vol. 25, pp. 247–264.

Guerrieri, L. (2006). 'The inflation persistence of staggered contracts', *Journal of Money, Credit and Banking*, Vol. 38, pp. 484–494.

Kara, E. (2010). 'Optimal monetary policy in the Generalized Taylor economy', *Journal of Economic Dynamics and Control*, Vol. 34, pp. 2023–2037.

Kehoe, P. and Midrigan, V. (2015). 'Prices are sticky after all', *Journal of Monetary Economics*, Vol. 75, pp. 35–53.

Klenow, P. and Kryvtsov, O. (2008). 'State-dependent vs. time-dependent pricing: does it matter for recent U.S. inflation?', *The Quarterly Journal of Economics*, Vol. 123, pp. 863–904.

Mankiw, G. and Reis, R. (2002). 'Sticky information versus sticky prices: a proposal to replace the new Keynesian Phillips curve', *Quarterly Journal of Economics*, Vol. 117, pp. 1295–1328.

Mash, R. (2003). *New Keynesian Microfoundations Revisited: A Calvo-Taylor-Rule-of-Thumb Model and Optimal Monetary Policy Delegation*, Oxford University Discussion Paper 174.

Mash, R. (2004). *Optimising Microfoundations for Inflation Persistence*, Oxford University Discussion Paper 183.

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

Paustian, M. and von Hagen, J. (2008). *The Role of Contracting Schemes for Assessing the Welfare Costs of Nominal rigidities*, CFPR Discussion Papers 6726.

Sheedy, K. (2010). 'Intrinsic inflation persistence', *Journal of Monetary Economics*, Vol. 57, pp. 1049–1061.

Smets, F. and Wouters, R. (2003). 'An estimated dynamic stochastic general equilibrium model of the euro area', *Journal of the European Economic Association*, Vol. 1, pp. 1123–1175.

Taylor, J. B. (1993). *Macroeconomic Policy in a World Economy*, Norton, New York, NY.

Taylor, J.B. (2016). 'The staying power of staggered wage and price setting models in macroeconomics', Chapter 25 in Taylor J. B. and Uhlig H. (eds), *Handbook of Macroeconomics*, *Vol. II*, Dublin: Elsevier B.V, pp. 2009–2042.

Walsh, C. (2003). *Monetary Theory and Policy*, 2nd Edition, MIT Press, Boston, MA.

Wolman, A. (1999). 'Sticky prices, marginal cost and the behaviour of inflation', *Federal reserve bank of Richmond quarterly*, Vol. 85, pp. 29–47.

Woodford, M. (2003). '*Interest and Prices: Foundations of a Theory of Monetary Policy*', Princeton University Press, Princeton, NJ.

## Supporting Information

Additional supporting information may be found in the online version of this article:

**Appendix:** Monthly calibration for SW model.