Price Trends over the Product Life Cycle and the Optimal Inflation Target*

Klaus Adam
University of Oxford, Nuffield College and CEPR

Henning Weber
Deutsche Bundesbank

PRELIMINARY AND INCOMPLETE – DO NOT CIRCULATE

February 15, 2019

Abstract

We present a sticky price model featuring a product life cycle and rich amounts of heterogeneity that allow capturing salient features of micro price data. We analytically derive the optimal steady-state inflation rate and show how the optimal inflation target trades off price and mark-up distortions across different product categories. We then devise an empirical approach that permits identifying the optimal inflation target from relative price trends over the product life cycle. Using the micro price data underlying the U.K consumer price index, we show that the relative price of products decreases over the life cycle for almost all expenditure categories and that this causes positive inflation targets to be optimal. The optimal U.K. inflation target is estimated to range between 2.6% and 3.2% in the year 2016. The optimal target has steadily increased over the years 1996-2016. We explain how changes in relative price trends contributed to this increase.

JEL Class. No.: E31,

*We thank our discussant Yuriy Gorodnichenko, and seminar participants at the Dutch Central Bank (DNB), Goethe University in Frankfurt, and University of Oxford. Special thanks go to our research assistant Michal Marencak. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem. Any remaining errors are our own.

1 Introduction

A defining feature of modern economies is the high rate at which new products get introduced into the market place.[1] This feature is prevalent in many micro data sets, but is routinely abstracted from in the monetary policy literature. The relative neglect of the product life cycle is somewhat surprising, but - as this paper shows - not innocuous from the perspective of monetary policy design: features of the product life cycle turn out to be important for determining the optimal inflation rate that a welfare maximizing central bank should target.

To make our point, we present a model that incorporates a product life cycle and rich forms of product heterogeneity into an otherwise standard sticky price economy. The incorporated features allow capturing key characteristics of micro price behavior, as observed - for instance - in the micro price data underlying the construction of the consumer price indices. In particular, our model features many (heterogeneous) expenditure categories, each of which consists of many individual products that get continuously replaced over time. Product prices are subject to (heterogeneous degrees of) price rigidity and individual products characterized by idiosyncratic product quality and productivity components. The latter gives rise to (efficient) dispersion of product prices. The economy also features heterogenous growth rates in productivity and product quality across expenditure categories and in an extension we incorporate sales (or temporary) prices.

Despite the richness of the model, we can derive a closed-form expression for the optimal steady-state inflation rate, i.e., for the inflation target that a welfare-maximizing central bank would adopt. Analytical aggregation is partly feasible because we abstain from explicitly modeling the product replacement process, instead treat it as an exogenous (albeit heterogeneous) stochastic process. In fact, the precise economic forces driving product replacement are not important for our results, as long as these forces are independent of the steady state inflation rate pursued by the central bank.[2]

We show that the optimal inflation target depends on the quality and productivity trends that are present at the product level, as these determine whether a product should become cheaper or more expensive over its lifetime relative to the average product it is competing with. Since there are many heterogeneous relative price trends present in the economy, the optimal inflation target is a weighted average of these trends, where the weights are themselves nonlinear functions of the fundamental parameters characterizing the economy (the various degrees of price stickiness, product entry and exit rates, product demand elasticity, etc.).

[1] See Nakamura and Steinsson (**?**) and Broda and Weinstein (**?**) for evidence on product turnover.

[2] A number of potential forces driving product replacement dynamics naturally satisfy the independence requirement: product replacement could be driven by changing consumer tastes that cause some products to fall out of fashion and others to become fashionable; alternatively, replacement could be driven by negative productivity shocks that cause the producer of an existing product to discontinue production and have the next best producer enter the market with a new product.

While it is generally difficult to empirically operationalize the non-linear expression for the optimal inflation target that we derive from the model, we show that to first-order of accuracy only two features of heterogeneity matter for the optimal inflation rate: (1) heterogeneity in productivity and quality growth across expenditure categories and (2) heterogeneity in expenditure weights across expenditure categories.[3] All remaining dimensions of heterogeneity, e.g., the heterogeneity in price stickiness as emphasized in the sticky price literature (e.g., Aoki (**?**), Benigno (2004)), turn out to generate only second-order effects for the optimal inflation rate.

We then show how the first-order result can be made operational from an empirical point of view and how one can estimate the optimal inflation target using micro price data sets only. We thus show how to make direct normative inference from the regularly available micro data sets that underlie the construction of the consumer price indices. These data sets come together with direct estimates of the expenditure weights, thereby covering ingredient (2) mentioned above. In addition, we show that the trend in *relative* prices over the lifetime of a product is sufficient to identify the inflation-relevant trends in productivity and quality, i.e., ingredient (1) mentioned above. This insight is key and holds because price rigidities and historically suboptimal inflation rates distort only the *level* of relative prices but leave its *trend* unaltered. The observed relative price trends are thus efficient, which means they can be used to identify structural features of the economy.

We apply our approach to the micro price data that underlies the construction of the consumer price index in the United Kingdom. The data is provided by the Office of National Statistics (ONS) and covers the period 1996-2016. Over this period, the data set features more than 1200 expenditure items (or expenditure categories) and contains close to 29 million individual price observations.

We document that for the vast majority of expenditure items, the price of an individual product relative to the average item price declines over the lifetime of the product. That is, new products tend to be initially expensive, but become cheaper over their lifetime in relative terms. We furthermore show that this downward trend in relative prices has significantly accelerated over the past two decades: the expenditure items that dropped out of the consumption basket displayed smaller relative price declines than the average expenditure item, and newly entering items displayed above average relative price declines.

The negative price trends we document imply that the optimal inflation target is positive: in the year 2016, it ranges between 2.6% and 3.2% for the U.K. economy. And since negative price trends have accelerated over time, the inflation target has increased by around 1.2% over the period 1996 to 2016.

To understand why negative relative price trends make positive inflation optimal, first recall that the theory tells us that the observed relative price trends in the micro price data are the same as the ones one would observe in a setting

[3] As we explain in the main text, there is a third but quantitatively irrelevant element that matters to first order of approximation: heterogeneity in the steady-state real growth rates of (quality-adjusted) output across expenditure categories.

with flexible prices, i.e., the documented relative price trends are efficient. The question of finding the optimal inflation rate then amounts to determining how to best implement these relative price declines in a setting with sticky prices. The decline could be implemented in many alternative ways, but these turn out *not* to be welfare equivalent in the presence of nominal rigidities.

To illustrate this point, consider two alternative approaches for implementing declining relative prices. One approach lets newly entering products charge a fixed nominal price P, independently of the entry data, and subsequently lets them cut the nominal price at some constant rate over the product lifetime. With constant product entry and exit, the product price distribution and thus the average product price is constant over time, so that there is zero inflation in this item category, but all individual prices decline over their respective lifetimes. This setting requires constant adjustments of existing prices. Yet, when prices are rigid, these price adjustments happen inefficiently. An alternative - and as we show preferable - approach is to have the nominal prices of existing products to remain constant over time. One can nevertheless implement a decline in relative prices, simply by having newly entering products charge a higher (constant) price than the average existing product. This way, relative prices decline because the average product price, i.e., the item price level, keeps rising over time: there is positive inflation. Since this setting requires no adjustment of individual product prices, it is generally preferred when prices are sticky.

To the extent that the strength of relative price declines varies across expenditure items, as we observe in the data, the optimal inflation rate also varies across different expenditure items. It is thus generally not possible to implement the efficient decline in relative prices with perfectly constant nominal product prices in all expenditure items. The optimal weights in our theoretical result trade-off the relative price distortion in different expenditure items against each other and deliver the aggregate inflation rate that creates the least welfare-detrimental price distortions.

The paper also tests a core prediction of our theory using the available micro price data. Specifically, the theory predicts that the degree of excess price dispersion, i.e., the price dispersion over and above the level justified by heterogenous product qualities and heterogeneous productivities, is increasing in the gap between the optimal inflation rate and the actual inflation rate at the item level. We test whether this is in fact the case in the U.K. micro price data set and find strong and robust empirical support. This is comforting, as this channel is the main margin through which the optimal inflation rate is predicted to improve the distribution of relative prices and thus welfare.

The literature discussing the role of the product life cycle is relatively sparse. An early literature presented theoretical models of the evolution of firm entry, exit and product innovation, but abstracted from nominal rigidities and monetary issues, e.g. Shleifer (**?**) and Klepper (**?**). Broda and Weinstein (**?**) present empirical evidence on product creation and destruction for an important consumer good segment and quantify the quality bias in consumer price indices. Nakamura and Steinsson (**?**) present evidence for the BLS consumer and produce price data sets. Argente, Lee and Moreira (**?**) provide empirical evidence on how firms grow through the introduction of new products and Argente and Yeh (?) determine to what extent product replacement and perpetual demand learning by firms contributes to monetary non-neutrality. To the best of our knowledge, the latter paper is the only one incorporating a product life cycle into a setting with nominal rigidities. We are not aware of any paper discussing the product life cycle in connection with monetary policy design.

The monetary policy literature has considered settings with endogenous firm entry and exit (Bergin and Corsetti (2008), Bilbiie et al. (2008) and Bilbiie, Fujiwara and Ghironi (2014)), which could be re-interpreted as models of endogenous product entry and exit.[4] These papers study a complementary setup in which monetary policy affects the entry decisions of firms/products, but abstract from firm/product heterogeneity. Heterogeneity is key for our results.[5]

The paper is also related to the optimal inflation literature, see Schmitt-Grohé and Uribe (2010) for an overview. This literature has identified a number of complementary economic forces affecting the optimal rate of inflation. Concerns about an occasionally binding lower bound constraint on nominal interest rates, for instance, tend to generate a force towards positive inflation on average (Adam and Billi (2006, 2007), Coibion, Gorodnichenko and Wieland (2012)). The same tends to be true when wages are downwardly rigid (Carlsson and Westermark (2016), Benigno and Ricci (2011), Carlsson Kim and Ruge-Murcia (2009)). Conversely, the optimal inflation rate tends to be negative when taking into account cash distortions, as cash generates a force making the Friedman rule optimal (Kahn, King and Wolman (2003)).

The paper is structured as follows. The next section presents the micro price data set that we use in our empirical analysis. Section 3 introduces the sticky price model featuring a product life cyle and section 4 derives the closed-form result for the optimal inflation target. Section 5 explains how to estimate the optimal inflation target from micro price data and section 6 explains how this approach is robust against imperfect quality adjustment of product prices. Section 7 presents our baseline estimation results and section 8 offers various robustness checks. Finally, section 9 shows that the model-implied relationship between inflation and excess price dispersion is present in the micro price data. A conclusion briefly summarizes.

2 U.K. Micro Price Data

We employ the data on product price quotes that the U.K.'s Office for National Statistics (ONS) collects each month to compile the official consumer price index (CPI).[6]

[To be added . . .]

[4] Broda and Weinstein (?) emphasize that product entry and exit dynamics differ considerably from firm or establishment entry and exit dynamics.

[5] In the absence of product heterogeneity, which is a degenerate case of our setting, the optimal steady state inflation rate is always zero

[6] See Office for National Statistics (2014) for extensive data documentation.

3 A Sticky Price Model with a Product Life Cycle

This section introduces a sticky price model featuring a product life cycle. To capture key elements of the behavior of micro price data, the model incorporates considerable amounts of heterogeneity across products and expenditure items, thereby generalizing the setup of Adam and Weber (?) along two key dimensions. First, it considers a setting with multiple expenditure items, each of which is populated by a continuum of heterogeneous products. Expenditure items are allowed to have different degrees of price stickiness, product-entry and exit rates, and productivity and quality trends. Important heterogeneity in terms of price adjustment frequencies and product substitution rates has been documented by Nakamura and Steinsson (?). Likewise, heterogeneity in productivity and quality trends is key for being able to capture the heterogeneity in relative price trends that we uncover in our empirical section later on. Second, we augment the model with idiosyncratic elements to product quality and productivity. This allows the model to generate realistic amounts of price dispersion across products.

Due to the large amounts of heterogeneity, monetary policy is unable to implement the efficient distribution for relative prices when product prices are sticky. Instead, monetary policy must trade-off the price and mark-up distortions generated across the product prices in different items, such that the resulting equilibrium is least distortionary in welfare terms. We derive in the next section a closed-form expression for the optimal steady-state inflation rate, which we interpret as the optimal inflation target. Doing so requires analytically aggregating the non-linear sticky price model.

3.1 Demand Side and Production Side

The demand side of the model is rather parsimonious and consists of a representative consumer with balanced growth consistent preferences over an aggregate consumption good C_t and hours worked L_t, described by

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{[C_t V(L_t)]^{1-\sigma} - 1}{1-\sigma} \right),$$

where $\beta \in (0,1)$ is a discount factor and $\sigma > 0$.[7] The household faces the flow budget constraint

$$C_t + K_{t+1} + \frac{B_t}{P_t} = (r_t + 1 - d)K_t + \frac{W_t}{P_t} L_t + \int_0^1 \frac{\Theta_{jt}}{P_t} \, \mathrm{dj} + \frac{B_{t-1}}{P_t}(1 + i_{t-1}) - T_t,$$

where K_{t+1} denotes the capital stock, B_t nominal government bond holdings, P_t the nominal price of the aggregate consumption good, i_{t-1} the nominal interest

[7] We assume $\sigma > 0$ and that $V(\cdot)$ is such that period utility is strictly concave in (C_t, L_t) and that Inada conditions are satisfied.

rate, W_t the nominal wage rate, r_t the real rental rate of capital, d the depreciation rate of capital, Θ_{jt} nominal profits from ownership of firm j, and T_t lump sum taxes. Household borrowing is subject to a no-Ponzi scheme constraint. The first-order conditions characterizing optimal household behavior are entirely standard.To insure existence of a well-defined balanced growth path, we assume throughout the paper that $\beta < (aq)^{\phi\sigma}$.

The aggregate consumption good C_t is made up of Z_t different consumption items (in the language of the ONS). A consumption item is a product category, e.g., "Flatscreen TV, 30-inch display" or "CD-player, portable", which itself contains a range of individual products. Letting C_{zt} $(z = 1, ..., Z_t)$ denote consumption of item z, we have

$$C_t = \prod_{z=1}^{Z_t} (C_{zt})^{\psi_{zt}} ,$$

where $\psi_{zt} \geq 0$ denotes the expenditure weight for item z at time t and $\sum_{z=1}^{Z_t} \psi_{zt} = 1$. We allow the set of items Z_t and the expenditure weights ψ_{zt} to be time-varying, so as to capture the fact that ONS regularly drops and adds items to its consumption basket and adjusts the expenditure weights over time. For simplicity, we interpret item entry and exit or changing expenditure weights for items as simply being due to changing consumer tastes. Obviously, item substitution could be due to a variety of other factors, such as increased competition from a different item, e.g., flash-drive devices becoming increasingly competitive relative to portable CD players and thus leading to the exit of the latter. For simplicity, we do not model competition across items explicitly, instead take changes in the item structure as exogenous. Changes in the item structure is a slow process in the U.K. with on average only about 0.5% of items leaving the sample every month.

Every item contains a large number of differentiated products. To capture this fact, item level consumption C_{zt} is a Dixit-Stiglitz aggregate of individual products $j \in [0, 1]$, so that

$$C_{zt} = \left(\int_0^1 \left(Q_{jzt} \widetilde{C}_{jzt} \right)^{\frac{\theta-1}{\theta}} \mathrm{dj} \right)^{\frac{\theta}{\theta-1}} ,$$

where $\widetilde{C}_{jzt}$ denotes the consumed units of product j in item z in period t, Q_{jzt} the quality level of the product and $\theta > 1$ the elasticity of substitution between products. The aggregation assumes that consumption goods with higher product quality deliver proportionately higher consumption services relative to consumption goods with a lower quality level. This a standard approach for modeling the quality content of goods, see Schmitt-Grohe and Uribe (2012). In line with the evidence in micro price data, there will be constant churning of products j at the level of each expenditure item z. In particular, as we describe further below, there is constant entry and exit of products, so that the product composition at the item level is also evolving over time. In the U.K. data, product entry and exit is - unlike item entry and exit discussed above - a fast moving process: on average about 8% of products leave the sample every month. The source of the product replacement in the model can again be interpreted as reflecting changing consumption preferences, although alternative interpretations are possible.[8]

In equilibrium, the quantity of products $\widetilde{C}_{jzt}$ consumed must be equal to the quantity $\widetilde{Y}_{jzt}$ produced. Individual products are produced using a Cobb-Douglas production function

$$\widetilde{Y}_{jzt} = A_{zt} G_{jzt} \left(K_{zjt}\right)^{1-\frac{1}{\phi}} \left(L_{zjt}\right)^{\frac{1}{\phi}} \tag{1}$$

where A_{zt} denotes the level of productivity common to all producers of products in item z and G_{jzt} a product-specific productivity factor that captures idiosyncratic productivity componenents, as well as productivity dynamics associated with experience accumulation in the manufacturing of the product. The variables K_{jzt} and L_{jzt}, respectively, denote the capital and labor inputs into production.

Let $\delta_z \in (0, 1)$ denote the exogenous and idiosyncratic probability that a product j in item z drops out and is replaced by a new product in period t. This may happen because (1) consumers simply no longer demand this product, or (2) the producers of this product receive a sufficiently negative productivity shock that causes the product to become uncompetitive and being replaced by a new product, or (3) a newly available product is in quality-adjusted terms simply more attractive, as will become clear below. Whatever is the precise cause for product turnover, we assume that it can be described by an idiosyncratic and exogenous Poisson process with arrival rate δ_z.

For simplicity, we assign to the newly entering product the same product index j as to the exiting product. Let s_{jzt} denote the age of product j in item z at time t, with $s_{jzt} = 0$ in the period of entry. Given this definition, the period $t - s_{jzt}$ denotes - at any time t - the date at which product j that is available in period t entered into the economy. Initial productivity component at the time of product entry is independently drawn from some distribution Ξ_z^G, that is

$$\epsilon_{jzt}^G \sim \Xi_z^G, \tag{2}$$

with the property $E[\epsilon_{jzt}^G] = 1$ and $\epsilon_{jzt}^G > 0$. The idiosyncratic productivity remains constant throughout the lifetime of the product. We denote the experience-related productivity component of firm j with age s_{jzt} as

$$G_{jzt} \quad = \quad \overline{G}_{jzt} \cdot \epsilon_{jzt}^G, \tag{3}$$

where $\overline{G}_{jzt}$ evolves over time according to

$$\overline{G}_{jzt} = \begin{cases} 1 & \text{for } s_{jzt} = 0, \\ g_{zt} \overline{G}_{jz,t-1} & \text{otherwise,} \end{cases} \tag{4}$$

[8]A specific model of a portable CD players, for instance, may simply fall out of fashion, with a new model becoming fashionable instead.

and

$$g_{zt} \quad = \quad g_z \epsilon^g_{zt}, \tag{5}$$

where $g_z > 0$ denotes the average growth rate of this productivity component and captures the average rate of experience accumulation in the production of products in item z. The disturbance ϵ^g_{zt} is an arbitrary stationary process satisfying $E \ln \epsilon^g_{zt} = 0$. The common level of productivity for item z evolves according to

$$\begin{aligned} A_{zt} &= a_{zt} A_{zt-1} \\ a_{zt} &= a_z \epsilon^a_{zt}, \end{aligned}$$

where $a_z > 0$ denotes the average productivity growth rate and ϵ^a_{zt} is an arbitrary stationary processes satisfying $E \ln \epsilon^a_{zt} = 0$. While accumulated experience G_{jzt} associated with product j in item z is lost upon exit of the product, the growth rate in the common productivity level A_{zt} allows for permanent productivity gains in item z.

It now remains to describe the product quality dynamics at the item level. Let $\widetilde{P}_{jzt}$ denote the price at which one unit of output $\widetilde{Y}_{jzt}$ is sold at time t. The price $\widetilde{P}_{jzt}$ will be set by the producer optimally subject to price-adjustment frictions. The quality-adjusted price of the product is then given by

$$P_{jzt} = \frac{\widetilde{P}_{jzt}}{Q_{jzt}}, \tag{6}$$

where Q_{jzt} denotes the product-specific quality level that is assumed constant over the lifetime of the product.[9] The quality level of products j entering in period t is given by

$$Q_{jzt} = Q_{zt} \cdot \epsilon^Q_{jzt}, \tag{7}$$

where ϵ^Q_{jzt} captures an idiosyncratic quality component and Q_{zt} a common item-level quality trend. The idiosyncratic component is an independent draw from some distribution Ξ^Q_z, that is

$$\epsilon^Q_{jzt} \sim \Xi^Q_z,$$

with the properties $E[\epsilon^Q_{jzt}] = 1$ and $\epsilon^Q_{jzt} > 0$. The common quality component evolves according to

$$Q_{zt} \quad = \quad q_{zt} Q_{zt-1} \tag{8}$$
$$q_{zt} \quad = \quad q_z \epsilon^q_{zt}, \tag{9}$$

where q_z denotes the average quality progress in item z and ϵ^q_{zt} a random component of quality growth, which is an arbitrary stationary process satisfying

[9]Within the model, we interpret a new quality level of the same product as being a new product. Since we assign the same index j to the exiting and newly entering product, Q_{jzt} must have a time subscript, as the entering and exiting product can have different quality levels.

$E \ln \epsilon_{zt}^{q} = 0$. Since the quality of a product remains constant over its lifetime, we have

$$Q_{jzt} = Q_{jzt-s_{jzt}} \text{ for all } (j, z, t),$$

where s_{jzt} denotes the age of product j in item z at time t.

Using the quality-adjusted product price (6), we can then define the quality-adjusted price level for item z according to

$$P_{zt} = \left(\int_0^1 \left(\frac{\widetilde{P}_{jzt}}{Q_{jzt}} \right)^{1-\theta} \mathrm{dj} \right)^{\frac{1}{1-\theta}}. \tag{10}$$

Aggregation across items z delivers the quality-adjusted overall price level

$$P_t = \prod_{z=1}^{Z_t} \left(\frac{P_{zt}}{\psi_{zt}} \right)^{\psi_{zt}}.$$

The optimal inflation rate will be defined in terms of this (perfectly) quality-adjusted price level, i.e., inflation is defined as

$$\Pi_t = \frac{P_t}{P_{t-1}}.$$

We discuss the effects of imperfect quality adjustment of the price level for our results in section 6.

Optimal product demand by consumers and market clearing in all products implies that product demand satisfies

$$Y_{jzt} = Y_{zt} \left(\frac{P_{jzt}}{P_{zt}} \right)^{-\theta}, \tag{11}$$

where

$$Y_{jzt} \equiv Q_{jzt} \widetilde{Y}_{jzt} \tag{12}$$

denotes output in constant quality units.

3.2 Optimal Price Setting

We now consider the price setting problem of product producers. We assume that the price of a product can be chosen freely at the time of product entry, but that price adjustments at the product level are subsequently subject to Calvo-type adjustment frictions. In particular, let $\alpha_z \in [0, 1)$ denote the time-invariant idiosyncratic probability that the price of some product j in item z can not be adjusted in any given period. Since product quality is constant over the product lifetime (new qualities are essentially treated as new products), the price setting problem can - without loos of generality - be formulated as a problem of choosing the quality-adjusted product price P_{jzt}. Furthermore, let W_t denote the nominal wage and r_t the real rental rate of capital. The factor input mix (K_{jzt}, L_{jzt}) is then optimally chosen so as to minimize production costs $K_{jzt}P_t r_t + L_{jzt}W_t$ subject to the constraints imposed by the production function (1). Let

$$I_{jzt} \equiv Y_{jzt}/(A_{zt}Q_{jz,t-s_{jt}}G_{jzt})$$

denote the units of factor inputs $(K_{jzt}^{1-\frac{1}{\phi}} L_{jzt}^{\frac{1}{\phi}})$ required to produce Y_{jzt} units of (quality-adjusted) output. We show in appendix ?? that cost minimization implies that the (nominal) marginal costs of I_{zjt} are given by

$$MC_t = \left(\frac{W_t}{1/\phi}\right)^{\frac{1}{\phi}} \left(\frac{P_t r_t}{1-1/\phi}\right)^{1-\frac{1}{\phi}}. \tag{13}$$

We can then express the optimal price-setting problem of a firm j that can optimize its price in period t as

$$\max_{P_{jzt}} E_t \sum_{i=0}^{\infty} (\alpha_z(1-\delta_z))^i \frac{\Omega_{t,t+i}}{P_{t+i}} \left[(1+\tau)P_{jzt}Y_{jzt+i} - MC_{t+i}I_{jzt+i}\right] \tag{14}$$

$$s.t. \quad I_{jzt+i} = Y_{jzt+i}/(A_{zt+i}Q_{zt+i}\mathcal{Q}_{jzt+i})$$

$$Y_{jzt+i} = \psi_z \left(\frac{P_{jzt}}{P_{zt+i}}\right)^{-\theta} \left(\frac{P_{zt+i}}{P_{t+i}}\right)^{-1} Y_{t+i},$$

where $\mathcal{Q}_{jzt} = Q_{jz,t-s_{jzt}}G_{jzt}/Q_{zt}$, $\Omega_{t,t+i}$ denotes the representative household's discount factor between periods t and $t+i$, and τ denotes a sales tax/subsidy. The first constraint in the previous problem captures the firm's technology (1); the second constraint captures consumers' optimal product demand function (11). Appendix ?? shows that the optimal price $P_{jzt}^{\star}$ satisfies

$$\frac{P_{jzt}^{\star}}{P_t}\left(\frac{Q_{jzt-s_{jzt}}G_{jzt}}{Q_{zt}}\right) = \left(\frac{\theta}{\theta-1}\frac{1}{1+\tau}\right)\frac{N_{zt}}{D_{zt}}, \tag{15}$$

where the item-level variables N_{zt} and D_{zt} are independent of the firm index j and defined in the appendix. The previous equation shows that the optimal relative reset price of a firm $(P_{jzt}^{\star}/P_t)$ depends only on how its own quality-adjusted productivity $(A_{zt}Q_{jzt-s_{jzt}}G_{jzt})$ relates to the average quality-adjusted productivity of newly entering products $(A_{zt}Q_{zt})$ and on item level variables.

4 The Optimal Inflation Target: Theory

This section presents our main analytic result about the optimal inflation target. The optimal target is defined as the inflation rate that maximizes steady state utility. The idea underlying this approach is that economic shocks generate only temporary deviations of the optimal inflation rate from its steady-state value, so that the average inflation rate that a welfare-maximizing central bank should target is in fact the optimal steady-state inflation rate.[10]

[10]This holds true to a first-order approximation in the aggregate shocks. Higher-order terms can cause the average inflation rate to differ from the steady state value. In sticky price

A steady state is a situation without aggregate shocks and without item turnover, in which turnover at the product level continues to operate:

Definition 1 *A steady state is a situation with a fixed set of items* $Z_t = Z$*, constant expenditure weights* $\psi_{zt} = \psi_z$*, no aggregate item-level disturbances* $(g_{zt} = g_z, q_{zt} = q_z, a_{zt} = a_t)$*, and a constant (but potentially suboptimal) inflation rate* Π*. The following idiosyncratic shocks continue to operate in a steady state: product entry and exit shocks, shocks to price adjustment opportunities, and product-specific shocks to product quality and productivity that realize at the time of product entry.*

The following proposition states our main result:

Proposition 1 *Consider a steady state with a Pigouvian output subsidy that eliminates steady-state mark-ups under flexible prices* $(\theta/((1+\tau)(\theta-1)) = 1)$*. Consider the limit* $\beta(\gamma^e)^{1-\sigma} \to 1$*, where* $\gamma^e = \prod_{z=1}^{Z}(a_z q_z)^{\psi_z \phi}$ *denotes the growth rate of the aggregate economy (in quality-adjusted terms).*

The welfare maximizing steady-state inflation rate $\Pi^\star$ *is given by*

$$\Pi^\star = \sum_{z=1}^{Z} \omega_z \frac{\gamma_z^e}{\gamma^e} \frac{g_z}{q_z}, \tag{16}$$

where γ_z^e/γ^e *is the output growth rate of item* z *relative to the aggregate growth rate, given by*

$$\frac{\gamma_z^e}{\gamma_z} = \frac{a_z q_z}{\prod_{z=1}^{Z}(a_z q_z)^{\psi_z}},$$

and the item weights $\omega_z \geq 0$ *are given by*

$$\omega_z = \frac{\tilde{\omega}_z}{\sum_{z=1}^{Z} \tilde{\omega}_z},$$

with

$$\tilde{\omega}_z = \frac{\psi_z \theta \alpha_z (1-\delta_z)(\gamma^e/\gamma_z^e)^\theta (\Pi^\star)^\theta (q_z/g_z)}{[1-\alpha_z(1-\delta_z)(\gamma^e/\gamma_z^e)^\theta (\Pi^\star)^\theta (q_z/g_z)]\ [1-\alpha_z(1-\delta_z)(\gamma^e/\gamma_z^e)^{\theta-1}(\Pi^\star)^{\theta-1}]}.$$

Equation (16) shows that the optimal inflation target is a doubly-weighted average of the item-level terms g_z/q_z. We start by discussing the item-level terms g_z/q_z and consider thereafter the weights ω_z and γ_z^e/γ^e.

Items with $g_z > q_z$ generate a force towards positive inflation ($\Pi^* > 1$), while items with $g_z < q_z$ generate a force towards deflation ($\Pi^* < 1$). To understand why this the case, abstract for a moment from quality progress

models, however, higher-order terms tend to be quantitatively small, as a long as the lower bound on nominal rates is not binding.

($q_z = 1$) and suppose $g_z > 1$. Productivity then increases with the lifetime of the product, so that old products should become increasingly cheaper relative to newly entering products. In the presence of price setting frictions, this relative price decline of old products is best implemented by having new products charge higher prices, i.e., by positive amounts of inflation, rather than by having old products continuously adjust prices downward. This is so because price cuts cannot be synchronized across firms due to Calvo frictions and thereby give rise to inefficient price dispersion.[11] Now consider the polar case without age-dependent productivity ($g_z = 1$) and positive quality progress ($q_z > 1$). New products can then be produced at increasingly higher quality, without having to use more inputs into their production. New products should thus become cheaper (in quality-adjusted terms), relative to old products. Again, in the presence of price setting frictions, this is best achieved by having new products charge lower prices, i.e., via deflation, rather than by having old product increase prices.

We now consider the optimal nonlinear weighting scheme for the item-level terms g_z/q_z. A first set of weights, γ_z^e/γ_z, captures the (quality-adjusted) output growth in item z relative to the growth rate of the aggregate economy. This leads to an overweighting of items with fast output growth and an underweighting of item with slow growth. The second set of weights, ω_z, are rather nonlinear functions of item-level fundamentals, i.e., of the expenditure weight ψ_z, the product turnover rate δ_z, the price stickiness α_z, and the demand elasticity θ. The weights also depend on the item-level terms q_z/g_z and on the optimal inflation rate Π^* itself. Admittedly, the optimal item weights ω_z are generally difficult to interpret. They imply that in the special case where an item features no price stickiness ($\alpha_z = 0$), the optimal item weight is zero ($\omega_z = 0$), in line with the insights provided in Aoki (**?**). Similarly, in the special case without item heterogeneity, the result delivers the steady-state result derived in Adam and Weber (**?**).

Note that the proposition assumes an appropriate output subsidy and $\beta(\gamma^e)^{1-\sigma} \approx 1$. It follows from the proof of proposition 1 that when these conditions fail to hold, the inflation target (16) still minimizes the welfare losses stemming from relative price distortions, i.e., it still achieves the smallest possible distortions of productive efficiency. Yet, it then fails to also minimize the welfare effects associated with mark-up distortions.[12]

5 Optimal Inflation Target: Empirical Approach

This section explains how one can estimate the optimal inflation target. Unfortunately, it is rather difficult to empirically operationalize the nonlinear analytic

[11] With menu cost frictions, continous price cuts would be equally undesirable because price adjustment is costly.

[12] In fact, mark-up distortions are generally minimized by a different steady-state inflation rate, whenever $\beta(\gamma^e)^{1-\sigma} \ll 1$. If mark-up distortions can be eliminated by setting appropriate item-specific output subsidies, then the optimal inflation target (16) remains optimal even if $\beta(\gamma^e)^{1-\sigma} \ll 1$.

expression for the optimal inflation target provided in proposition 1. Doing so would require empirically identifying all structural parameters in a coherent way. To make progress, we consider a first-order approximation to the optimal inflation rate and show how one can empirically identify all elements that matter for the optimal inflation rate to first order using official micro price data set only. We leave an exploration of the effects of higher-order terms to future research.

The subsequent lemma provides a first-order approximation to the optimal inflation rate. It shows - rather surprisingly - that only a few dimensions of heterogeneity matter to first order:

Lemma 1 *The optimal steady-state inflation rate in the multi-sector economy is equal to*

$$\Pi^* = \sum_{z=1}^{Z} \psi_z \frac{\gamma_z^e}{\gamma^e} \frac{g_z}{q_z} + O(2), \tag{17}$$

where $O(2)$ denotes a second-order approximation error and where the approximation to equation (16) has been taken around a point, in which $\frac{g_z}{q_z}\frac{\gamma_z^e}{\gamma^e}$ and $\alpha_z(1-\delta_z)(\gamma^e/\gamma_z^e)^{\theta-1}$ are constant across sectors $z = 1, \ldots Z$.

Three dimensions of heterogeneity are relevant to first order: heterogeneity in expenditure weights ψ_z, heterogeneity in growth rates γ_z^e/γ^e and heterogeneity in the item-level terms g_z/q_z.

The first two dimensions can be readily identified from official micro price data sets. In particular, the expenditure weights are trivially part of these data sets. The heterogeneity in growth rates can be identified using the model-implied relationship $\gamma_z^e/\gamma^e = \Pi/\Pi_z$, which allows using the sample means of these inflation rates to estimate γ_z^e/γ^e. Heterogeneity along this dimension is, however, not quantitatively relevant for the estimated optimal inflation rate.

Identifying the item-level terms g_z/q_z is more challenging and explaining how to this can be achieved is a main contribution of the present paper. The following proposition states our main insight:

Proposition 2 *Consider a stochastic sticky price economy with a stationary (and potentially suboptimal) inflation rate Π_t. Let $T_{jz}^\star$ denote the set of periods in which the price of product j in item z can be adjusted and let s_{jzt} denote the product age. The optimal reset price $P_{jzt}^\star$ in adjustment periods, defined in equation (15), satisfies*

$$\ln \frac{P_{jzt}^\star}{P_{zt}} = f_{jz}^\star - \ln\left(\frac{g_z}{q_z}\right) \cdot s_{jzt} + u_{jzt}^\star, \qquad \textit{for all } t \in T_{jz}^\star, \tag{18}$$

where the residual satisfies $E[u_{jzt}^\star] = 0$.

Proof. From equation (**??**), the optimal reset price is

$$\frac{P_{jzt}^\star}{P_{zt}} \left(\frac{Q_{jz,t-s_{jzt}} G_{jzt}}{Q_{zt}}\right) = \left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\right) \frac{n_{zt}}{p_{zt} d_{zt}}.$$

Taking the natural logarithm of the previous equation yields

$$\ln \frac{P^\star_{jzt}}{P_{zt}} = \ln\left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\right) - \ln\left(\frac{Q_{jz,t-s_{jzt}} G_{jzt}}{Q_{zt}}\right) + \ln\left(\frac{n_{zt}}{p_{zt} d_{zt}}\right), \tag{19}$$

where

$$\ln\left(\frac{Q_{jz,t-s_{jzt}} G_{jzt}}{Q_{zt}}\right) = \ln(\epsilon^G_{jzt}\epsilon^Q_{jzt}) + \ln\left(\frac{Q_{z,t-s_{jzt}} \overline{G}_{jzt}}{Q_{zt}}\right),$$

which follows from equations (3) and (7) and where $\ln(\epsilon^G_{jzt}\epsilon^Q_{jzt})$ denote the idiosyncratic product-fixed effects. Using equations (4) and (8), we can write the previous equation as

$$\ln\left(\frac{Q_{jz,t-s_{jzt}} G_{jzt}}{Q_{zt}}\right) = \ln(\epsilon^G_{jzt}\epsilon^Q_{jzt}) + \ln\left(\frac{g_z}{q_z}\right) \cdot s_{jzt} + \sum_{i=t-s_{jzt}+1}^{t} \ln\left(\frac{\epsilon^g_{zi}}{\epsilon^q_{zi}}\right). \tag{20}$$

Substituting the previous equation into equation (19) yields

$$\begin{aligned} \ln \frac{P^\star_{jzt}}{P_{zt}} =& \ln\left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\right) - \ln(\epsilon^G_{jzt}\epsilon^Q_{jzt}) - \ln\left(\frac{g_z}{q_z}\right) \cdot s_{jzt} \\ &+ \ln\left(\frac{n_{zt}}{p_{zt} d_{zt}}\right) - \sum_{i=t-s_{jzt}+1}^{t} \ln\left(\frac{\epsilon^g_{zi}}{\epsilon^q_{zi}}\right). \end{aligned}$$

Defining

$$\begin{aligned} f^\star_{jz} &\equiv \ln\left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\frac{n_z}{p_z d_z}\right) - \ln(\epsilon^G_{jzt}\epsilon^Q_{jzt}) \\ u^\star_{jzt} &\equiv \ln\left(\frac{n_{zt}}{p_{zt} d_{zt}}\frac{p_z d_z}{n_z}\right) - \sum_{i=t-s_{jzt}+1}^{t} \left(\ln \epsilon^g_{zi} - \ln \epsilon^q_{zi}\right), \end{aligned}$$

we obtain equation (18). ■

The proposition shows that the optimal reset price displays a time trend at the rate g_z/q_z, which is our parameter of interest. This may appear surprising, given that the proposition considers a setting with sticky prices and potentially suboptimal inflation rates. To understand why this is the case, we note that with flexible prices ($\alpha_z = 0$), the same relative price trend would emerge. Specifically, the proof of proposition 2 implies that flexible prices satisfy

$$\ln \frac{P^f_{jzt}}{P_{zt}} = f^f_{jz} - \ln\left(\frac{g_z}{q_z}\right) \cdot s_{jzt} + u^f_{jzt}. \tag{21}$$

Price stickiness and suboptimal inflation rates thus only distort the level of relative reset prices, i.e., the intercept term f^f_{jz} and the residual u^f_{jzt}, but leave the time trend of relative prices invariant. This invariance property is key for being able to identify the structural parameters g_z/q_z. In fact, it implies that the relative price trends that are present in micro price data are efficient, i.e., reflect economic fundamentals.

This invariance of relative price trends is by no means special to the Calvo setting. It would similarly be present, if price rigidities were driven by menu costs instead. Menu costs generate inaction bands around the frictionless optimal relative price. If this frictionless price is trending over time, then reset prices must display the same time trend.[13]

Equation (18) thus provides a simple approach for empirically identifying g_z/q_z.[14] This is the case, despite the fact that the true product age s_{jzt} is typically not observed. As is easily seen, using the number of months since the product has been included into the price data set as the 'age' regressor, instead of the true product age, affects only the estimated intercept term, but leaves the coefficient of interest in front of the 'age' term unchanged. One can thus estimate the parameter of interest even without observing the true product age.

Since the complete price path of a product consists of a sequence of reset prices, which are adjusted when price-adjustment opportunities arise, the complete price path of the product displays the very same time trend as the path of reset prices, i.e., we have

$$\ln \frac{P_{jzt}}{P_{zt}} = f_{jz} - \ln\left(\frac{g_z}{q_z}\right) \cdot s_{jzt} + u_{jzt}, \tag{22}$$

for some alternative intercept term and residuals that again satisfy $E[u_{jzt}] = 0$.[15] In our empirical approach, we shall use equation (22) as our baseline equation, but we will also consider estimates based on the reset prices in equation (18) in a robustness exercise.

6 Imperfect Quality Adjustment

The estimation approach and the underlying theory developed in the previous sections assume that statistical agencies perfectly adjust prices for quality. This is clearly an idealized assumption, as a number of studies show that quality adjustment is far from perfect (Broda and Weinstein (**?**), Aghion et al. (**?**)).

This section shows that failure to adjust prices for quality levels leads to biases in the slope coefficients identified in regression (18), i.e., the coefficient in front of the age trend will then *not* be equal to g_z/q_z, unlike in the case with perfect quality adjustment. While this may be a source of concern, we show below that the optimal inflation target computed according to lemma 1 based on these biased slope estimates then nevertheless delivers the welfare

[13] This assumes - as is standard in menu cost models - that the width of the inaction bands does not itself display a time trend.

[14] This is true even though the residual $u^{\star}_{jzt}$ can potentially contain a unit root, see the proof of proposition 2. Since product lives tend to be short, the asymptotics of interest are the ones where the number of products gets large, not the ones where the lifetime of the products get large. Non-stationarity is thus not an issue for consistency and asymptotic normality of OLS estimates.

[15] Appendix **??** derives the properties of the modified residual in the previous equation.

maximizing target for the imperfectly quality-adjusted price index. In other words, the approach developed in the previous sections works perfectly well, even if quality adjustment is imperfect, as is likely to be the case in practice.

To make this point most forcefully, we consider an extreme setting in which the statistical agency makes no quality adjustments whatsoever. The item price level is thus computed using not-quality-adjusted prices $\widetilde{P}_{jzt}$ and given by

$$\widetilde{P}_{zt} = \left(\int_0^1 (\widetilde{P}_{jzt})^{1-\theta} \mathrm{dj} \right)^{\frac{1}{1-\theta}}, \tag{23}$$

with the associated item-level inflation rate given by

$$\widetilde{\Pi}_{zt} = \widetilde{P}_{zt} / \widetilde{P}_{zt-1}.$$

Appendix **??** derives the recursive law of motion for the not-quality adjusted item price level and shows that in steady state the following holds:

$$\widetilde{\Pi}_z = q_z \Pi_z. \tag{24}$$

The item-level inflation rate without quality adjustment $\widetilde{\Pi}_z$ thus exceeds the quality-adjusted inflation rate whenever there is quality growth ($q_z > 1$). Similarly, the aggregate steady-state inflation rate without quality adjustment $\widetilde{\Pi}$ is given by

$$\ln \widetilde{\Pi} = \ln \Pi + \sum_{z=1}^{Z} \psi_z \ln q_z, \tag{25}$$

and exceeds the quality-adjusted rate by a weighted average of item-level quality growth rates. This feature is well understood in the literature.

The key new observation in this section is that running the relative price-trend regressions (18) using not-quality-adjusted prices and price levels, leads to a particular bias in the time-trend coefficient, relative to a setting with perfect quality adjustment:

Proposition 3 *Consider a steady state with a potentially suboptimal inflation rate* Π*, and let* $T^\star_{jz}$ *denote the set of periods in which the price of product* j *in item* z *can be adjusted. The not-quality-adjusted optimal reset price* $\widetilde{P}^\star_{jzt}$ *relative to the not-quality-adjusted item price level* $\widetilde{P}_{zt}$ *satisfies*

$$\ln \frac{\widetilde{P}^\star_{jzt}}{\widetilde{P}_{zt}} = \widetilde{f}^\star_{jz} - \ln(g_z) \cdot s_{jzt}, \qquad \textit{for all } t \in T^\star_{jz}. \tag{26}$$

where s_{jzt} *denotes the age of product* j *in item* z *at time* t*.*

Proof. Consider a steady state and use equation (6) to replace in equation (**??**) the quality adjusted reset price $P^\star_{jzt}$ by $\widetilde{P}^\star_{jzt}/Q_{jzt}$. This yields

$$\frac{\widetilde{P}^\star_{jzt}}{\widetilde{P}_{zt}} \frac{\widetilde{P}_{zt}}{P_{zt}} \frac{1}{Q_{jz,t-s_{jzt}}} \left(\frac{Q_{jz,t-s_{jzt}} G_{jzt}}{Q_{zt}} \right) = \left(\frac{1}{1+\tau} \frac{\theta}{\theta-1} \right) \frac{n_z}{p_z d_z},$$

where we also augmented the left-hand side by $\widetilde{P}_{zt}$. Taking the natural logarithm of the previous equation and using equation (20) to substitute for $\ln(Q_{jz,t-s_{jzt}} G_{jzt}/Q_{zt})$ in the steady state yields

$$\ln \frac{\widetilde{P}^{\star}_{jzt}}{\widetilde{P}_{zt}} = \ln\left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\frac{n_z}{p_z d_z}\right) + \ln\left(\frac{Q_{jz,t-s_{jzt}}}{\epsilon^G_{jzt}\epsilon^Q_{jzt}}\right) - \ln\left(\frac{g_z}{q_z}\right)\cdot s_{jzt} + \ln\left(\frac{P_{zt}}{\widetilde{P}_{zt}}\right). \tag{27}$$

Steady-state relative item price levels evolve as

$$\begin{aligned}\ln(P_{zt}/\widetilde{P}_{zt}) &= (t+1)\ln(\Pi_z/\widetilde{\Pi}_z) + \ln(P_{z,-1}/\widetilde{P}_{z,-1})\\ &= -(t+1)\ln q_z + \ln(P_{z,-1}/\widetilde{P}_{z,-1})\\ &= -s_{jzt}\ln q_z - (t - s_{jzt} + 1)\ln q_z,\end{aligned}$$

where the second line follows from equation (24) and the third line assumes the initial condition $P_{z,-1}/\widetilde{P}_{z,-1} = 1$, without loss of generality. Using the previous equation to substitute for the ratio of price levels in equation (27) yields

$$\begin{aligned}\ln \frac{\widetilde{P}^{\star}_{jzt}}{\widetilde{P}_{zt}} = \ln\left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\frac{n_z}{p_z d_z}\right) + \ln\left(\frac{Q_{jz,t-s_{jzt}}}{\epsilon^G_{jzt}\epsilon^Q_{jzt}}\right) - (t - s_{jzt} + 1)\ln q_z\\ - \ln\left(\frac{g_z}{q_z}\right)\cdot s_{jzt} - \ln(q_z)\cdot s_{jzt},\end{aligned} \tag{28}$$

Defining the product-fixed effect as

$$\widetilde{f}^{\star}_{jz} \equiv \ln\left(\frac{1}{1+\tau}\frac{\theta}{\theta-1}\frac{n_z}{p_z d_z}\right) + \ln\left(\frac{Q_{jz,t-s_{jzt}}}{\epsilon^G_{jzt}\epsilon^Q_{jzt}}\right) - (t - s_{jzt} + 1)\ln q_z$$

shows that equation (28) is equivalent to equation (26) in the proposition. ■

In the presence of quality progress ($q_z > 1$), the regression estimates are thus also upwardly distorted by the amount of quality progress and given and given by g_z instead of g_z/q_z. Therefore, when computing the optimal inflation target for the not-quality-adjusted inflation rate using the distorted regression coefficients, one unwittingly implements the target rate for the quality-adjusted rate of inflation.

To formally shows this, note that the optimal inflation target from lemma 1 can alternatively be expressed as[16]

$$\ln \Pi^{\star} = \sum_{z=1}^{Z} \psi_z \ln\left(R_z \frac{\gamma^e_z}{\gamma^e}\right) + O(2). \tag{29}$$

where $\ln R_z$ is the regression coefficient on time trend in equation (18).[17] Using the distorted regression coefficients from proposition 3 and equation (29), we

[16] See appendix ?? for a derivation

[17] We have $R_z = g_z/q_z$ for perfect quality adjustment and $R_z = g_z$ in the absence of quality adjustment.

arrive at an optimal inflation target of

$$\ln \widetilde{\Pi}^{\star} = \sum_{z=1}^{Z} \psi_z \ln \left(g_z \frac{\gamma_z^e}{\gamma^e} \right)$$

for the not-quality-adjusted price index. From equation (25) then follows

$$\ln \Pi = \ln \widetilde{\Pi}^{\star} - \sum_{z=1}^{Z} \psi_z \ln q_z = \sum_{z=1}^{Z} \psi_z \ln \left(\frac{g_z}{q_z} \frac{\gamma_z^e}{\gamma^e} \right) = \ln \Pi^{\star},$$

which shows that the inflation target in terms of quality-adjusted prices is optimal. Imperfect quality adjustment is thus not a source of concern for the approach developed in this paper.

7 Optimal U.K. Target - Baseline Results

We now use the approach developed in the previous sections to estimate the optimal inflation target for the U.K. economy. The top panel of figure 1 presents our baseline estimate. The baseline estimation approach uses all available price observations to estimate the item level relative price trends g_z/q_z, see equation (21), and then uses the expenditure items present at any considered date, the corresponding ONS expenditure weights, as well as the estimated values for γ_z/γ, to compute the optimal inflation target at this date according to lemma 1.

Figure 1 shows that the optimal inflation target is significantly positive, has steadily increased by about 1.2% points over the period 1996-2016, and stands at approximately 2.6% in 2016. While there is a steady upward trend in the optimal inflation target over time, the figure does not imply that the Bank of England should have continuously revised its inflation target upward in line with the estimates shown in the graph. If target adjustments are economically costly, say because they require costly reputation building, then the optimal adjustments to the target would happen in a much more lumpy form than suggested by figure 1. This said, the overall increase in the estimated target of more than one percentage point is substantial.

The optimal inflation target in figure 1 deviates from zero in a quantitative significant way because there are strong relative price trends present at the item-level. Panel B in figure 1 depicts the distribution of item-level optimal inflation rates $\Pi_z^* \equiv g_z/q_z$ for all items present over the period 1996-2016. The panel depicts the distribution once in expenditure weighted form (blue bars) and once using item frequencies (red line). Both of these distribution show that the optimal inflation rate is positive for the vast majority of items, which shows that relative prices fall over the product life cycle in almost all expenditure items.

Panel B in figure 1 also highlights that the aggregate inflation result is not driven by outliers, instead there is a large mass of items for which the item-level

Figure 1: Optimal U.K. Target - Baseline Results

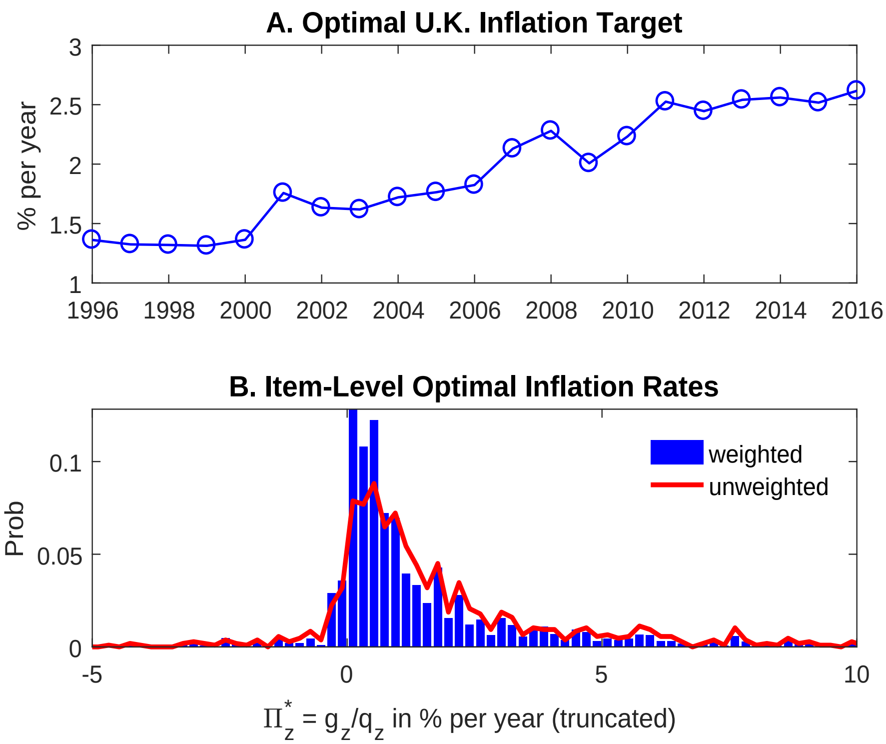

Notes:

optimal inflation rate is close to the estimated optimal inflation target. There is, however, considerable heterogeneity in relative price trends in the economy, with some expenditure items displaying very large amounts of relative price declines over the product life cycle.

Table 1 lists the 20 items with the largest item level inflation rate, i.e., the largest relative price declines, which have an expenditure weight of more than 0.2%. The table contains products that are either subject to a rapid loss of appeal over time, such as entertainment material or apparel, and products where supply increases rapidly over time, such as consumer electronics or household durables.

To understand the source of the upward trend in the optimal inflation target in figure 1, we perform a dynamic Olley-Pakes decomposition, following the approach of Melitz and Polanec (**?**). Specifically, we decompose the increase in the inflation target relative to the base year 1996 into three components: the effect of newly added items up to the year of interest, the effect of items that

Table 1: Items with Largest Absolute Optimal Item-Level Inflation Rates

Item Description	$\Pi_z^\star$ (in % per year)	Weight (in %)
Positive Optimal Inflation Rates		
WOMENS SHORT SLEEVE FORMAL TOP	76.02	0.22
PRE-RECORDED DVD (FILM)	53.91	0.32
WOMENS DRESS -CASUAL/FORMAL	44.51	0.26
WOMEN S DRESS-CASUAL 1	34.32	0.35
COMPUTER GAME TOP 20 CHART	27.70	0.62
LADY S SCARF	25.29	0.34
WOMENS TROUSERS-CASUAL	20.38	0.24
WOMENS CASUAL TROUSERS 2	18.31	0.21
WOMENS BLOUSE/SHIRT 1	17.57	0.20
WOMENS FORMAL JACKET	16.34	0.21
WOMENS SKIRT: WORK/FORMAL	14.71	0.23
DIGITAL CAMCORDER RS 2011	14.55	0.28
FLAT PANEL TV	11.62	0.25
WOMENS SUIT	9.83	0.35
PRE-RECORDED DVD TOP 20	8.86	0.46
MENS SHOES TRAINERS	8.51	0.36
WOMENS TROUSERS-FORMAL	7.67	0.33
MIDI HIFI (INC TWIN DECK CASS)	7.47	0.23
HIFI - 2008	7.22	0.28
WASHING MACHINE - 2005	5.90	0.25
Negative Optimal Inflation Rates		
DISPOSABLE RAZORS PACK OF 10	-0.30	0.25
CIGARETTES 5	-0.33	0.51
MILK SEMI-PER 2 PINTS/1.136 L	-0.34	0.52
AUTOMATIC WASHING MACHINE 2009	-0.35	0.32
PRIV RENTD UNFURNISHD PROPERTY	-0.41	2.05
18 CT GOLD GEMSTONE RING	-0.43	0.29
WASHING MACHINE - 2006	-0.44	0.24
AIR FRESHENER AEROSOL SPRAY	-0.48	0.21
COOKED HAM PREPACKED/SLICED	-0.84	0.34
BANANAS-PER KG	-1.05	0.24
LEISURE CENTRE ANNUAL M SHIP	-1.32	0.33
APPLES -DESSERT-PER LB	-1.40	0.22
WASHING MACH NO DRYER MAX 1800	-1.46	0.34
WASHING MACHINE - 2008	-1.79	0.32
CAMCORDER-8MM OR VHS-C	-2.29	0.33
WIDESCREEN TV - 2005	-2.48	0.62
HIFI - 2007	-3.18	0.31
PHOTO FRAME	-3.29	0.21
FRIDGE/FREEZER - 2007	-3.32	0.24
FRIDGE/FREEZER - 2005	-4.19	0.21

Notes: The table reports only items with expenditure weight $> 0.2\%$.

have exited up to the year of interest, and the effects of changing expenditure weights among continuing items up to the year of interest.

The results of this decomposition are depicted in figure 2. The bottom panel of the figure shows the number of continuing, exiting and entering items, respectively, at any given date relative to the base year 1996. The top panel decomposes the total increase in the inflation target (the solid blue line) into the three elements mentioned before. It shows that all of these individual elements have contributed to the increase in the optimal inflation target. The largest upward force comes from newly entering items, which display (on average) a larger rate of relative price decline than the items present in 1996. The second largest upward force comes from exiting items: exiting items displayed a rate of relative price decline that was on average below the one displayed by items that were present in 1996. Finally, a small positive force is due to a reshuffling of expenditure weights among the set of continuing items towards items displaying a larger rate of price.

Figure 2: Decomposing the upward trend in the optimal inflation target

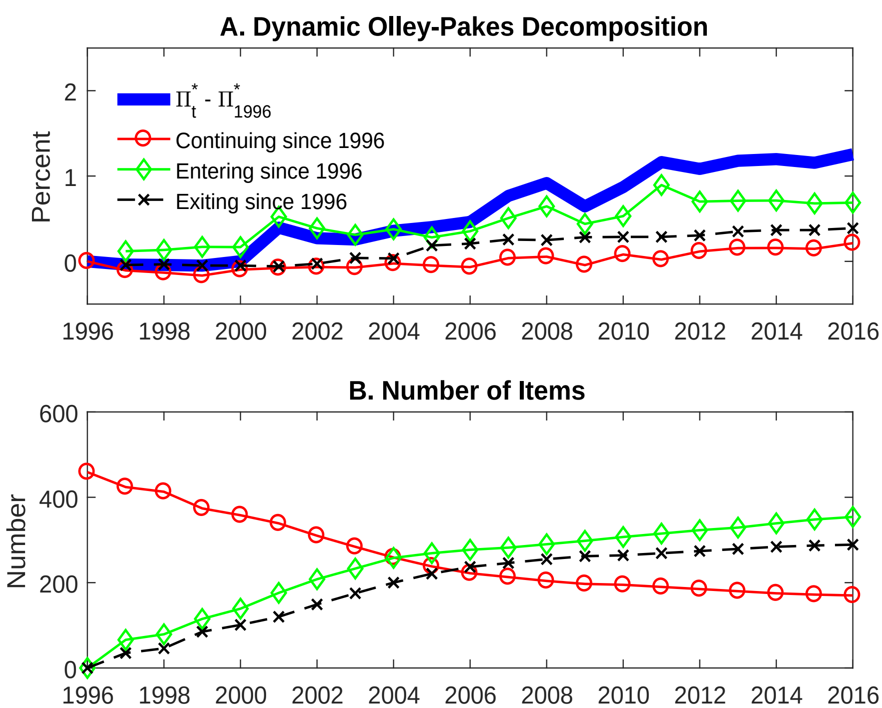

Notes:

Figure 3 compares the expenditure-weighted distribution of item-level inflation rates in 1996 and 2016. It shows how item entry and exit, as well as expenditure reweighting among continuing items have shifted the distribution of optimal inflation rates towards the right. The figure makes it clear that there was a sizable shift in the center of the distribution and that results are not driven by outliers.

Figure 3: Item-level optimal inflation rates: 2016 versus 1996, expenditure-weighted distributions

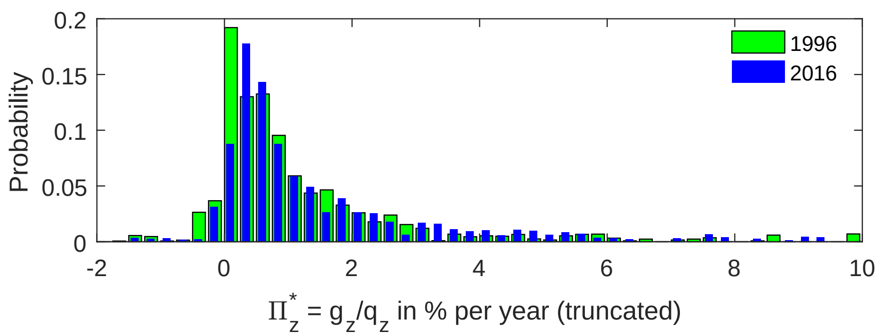

Notes:

8 Optimal U.K. Inflation - Robustness Checks

This section explores the robustness of the baseline results along two dimensions. Section 8.1 shows that results are robust towards using only reset prices, rather than all prices in the regressions. Section 8.2 shows that alternative approaches for dealing with sales prices tend to generate somewhat higher estimates for the optimal inflation target.

8.1 Using Reset Prices in the Price Trend Estimation

The baseline approach uses price regressions that incorporate all price observations available in our sample. Section 5 argued that - according to the theory - the price trend can alternatively be recovered using reset prices only (proposition 2). We thus rerun our relative price regressions using only price observations for which the monthly price deviated from the previous month's price, which is our empirical proxy for capturing a price adjustment opportunity. This naturally leads to a much smaller number of price observations. Yet, as figure 4 shows, the inflation target recovered via this alternative approach differs only in quantitatively minor ways from the baseline results. We find this reassuring, at the results can alternatively be interpreted as a test of an overidentifying restriction imposed by the underlying price setting model.

Figure 4: Optimal inflation target: baseline versus reset price based estimation

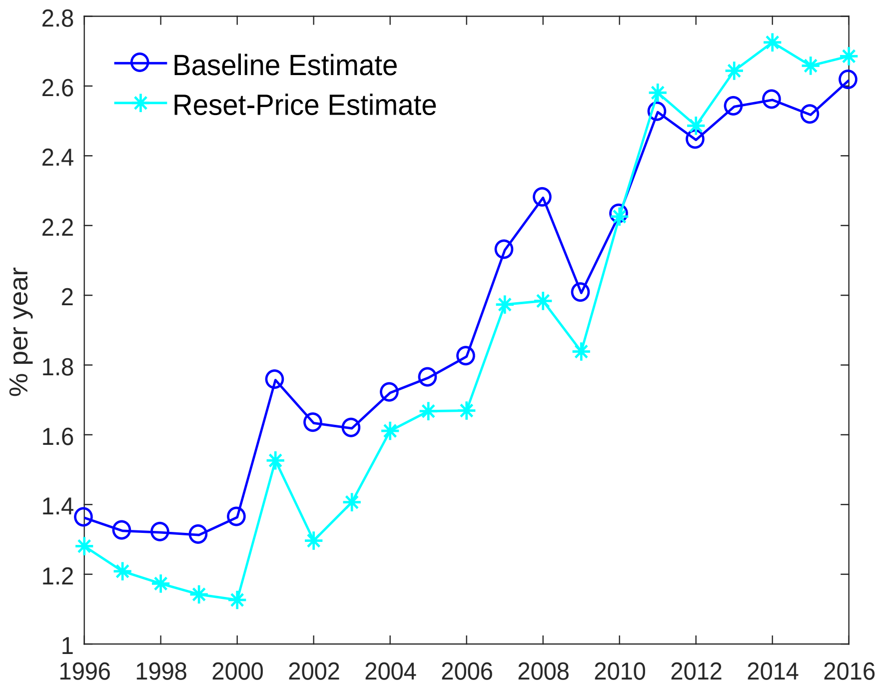

Notes:

8.2 Alternative Treatment of Sales Prices

An important feature of micro price data is that it features many short-lived price changes that are subsequently reversed. These typically take the form of temporary price reductions (sales), but also the form of occasional temporary price increases. While the sticky price framework outlined in the previous sections does not allow for such temporary price changes, it can be augmented to do so along the lines of Kehoe and Midrigan (**?**). As we explain below, the empirical estimation approach works even in the presence of sales.

Suppose that firms choose a regular list price P^L_{jzt}, which is subject to the same price adjustment frictions as the prices in the pure Calvo model presented thus far. After learning about the adjustment opportunity for the list price, a share $\alpha_T \in [0,1)$ of producers gets to choose freely a temporary price P^T_{jzt} at which they can sell the product in the current period. The temporary price is valid for one period only and does not affect the list price. Furthermore, absent further temporary price adjustment opportunities in the next period, prices revert to the list price in the next period. With this setup, the optimal temporary price $P^{T\star}_{jzt}$ is equal to the static optimal price in the period, i.e.,

$$\frac{P^{T\star}_{jzt}}{P_t}\left(\frac{Q_{jzt-s_{jt}}G_{jzt}}{Q_{zt}}\right) = \left(\frac{\theta}{\theta-1}\frac{1}{1+\tau}\right)\frac{MC_t}{P_t A_{zt} Q_{zt}}, \tag{30}$$

and the optimal list prices are given by the optimal price (15) from the model without temporary prices. Equation (30) shows that amongst the set of temporary prices in item z, prices are set such that they inversely reflect relative productivities and qualities. Furthermore, the temporary prices (30) will display the same price trends as the optimal list prices (15), so that the introduction of temporary price leaves the model-implied estimation approach unaltered. In particular, the optimal temporary price $P^{T\star}_{jzt}$, as defined in equation (30), satisfies

$$\ln\frac{P^{T\star}_{jzt}}{P_{zt}} = \widetilde{f}_{jz} - \ln\left(\frac{g_z}{q_z}\right)\cdot s_{jzt} + \widetilde{u}_{jzt} \qquad \text{for all } t \in \widetilde{T}^{\star}_{jz},$$

where $\widetilde{T}^{\star}_{jz}$ denotes the set of periods for which temporary prices can be charged for the product j in item z. In light of this result, the baseline estimation approach included all prices - including sales prices - in its regression.

Despite the previous observation, sales prices can make a difference for the estimated relative price trends. Sales prices might, for instance, not be evenly distributed over the product life cycle, unlike assumed in the theoretical setup sketched in the previous paragraph. If sales happen predominantly at the end of the product life spell, our baseline regressions would probably overestimate relative price declines and thereby overestimate the optimal inflation target. Conversely, if sales prices happen predominantly at the time of product introduction, or are larger at the time of introduction, then the opposite would be true. It is - therefore - of interest to investigate the robustness of our baseline results towards using alternative approaches for treating sales prices in the data.

Figure 5 displays the baseline estimate of the optimal inflation target together with various alternative estimates for the optimal inflation target. A first approach (baseline w/o sales prices) uses the sales flag from the ONS data sample to exclude all sales prices from regression (21).[18] The figure shows that the optimal inflation target increases by around 0.3% per year as a result. A quantitatively similar result is obtained, if only the so-called "regular prices" are used in the regression (Kehoe-Midrigan, regular prices only), where regular prices are defined according to the regular price filter of Kehoe and Midrigan (2015).

[18] A sales flag is an indicator variable that the price collector records, whenever she/he finds the product to be on sale. In this and subsequent robustness checks, we always recompute the item price levels after excluding or adjusting sales prices.

Rather than simply excluding sales prices, one can adjust sales prices based on various adjustment techniques and continue using them in the regressions. Figure 5 reports the outcome when making adjustments based on the filters A and B from Nakamura and Steinsson (2008) and the regular price filter of Kehoe-Midrigan (Kehoe-Midrigan, filtered prices). The outcome across these filtering approaches varies quite substantially. While the Nakamura-Steinsson filter B leads to only small adjustments relative to the baseline estimation, filter A leads to adjustments of the same order of magnitude as when dropping sales prices from the regression. The largest upward revision of the inflation target is observed for the regular price filter of Kehoe and Midrigan: the inflation target is on average about 0.56% higher than the baseline estimate.

Overall, we can conclude that a different treatment of sales prices can lead to considerably higher optimal inflation targets than the ones obtained via our baseline approach.

Figure 5: Optimal inflation target for alternative treatments of sales prices

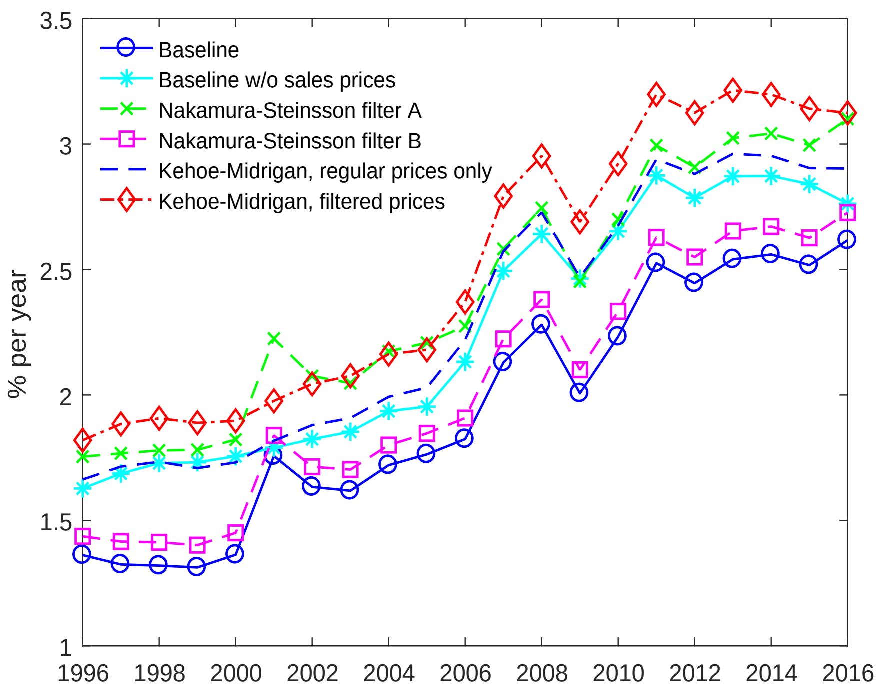

Notes:

9 Inflation and Excess Price Dispersion

The inflation target derived in our theoretical setting minimizes the amount of price distortions in the economy. In particular, the theory predicts that any deviation of the item-level inflation rate Π_z from the optimal item-level inflation rate $\Pi_z^* \equiv g_z/q_z$ gives rise to excess price dispersion, i.e., price dispersion over and above the efficient price dispersion justified by quality and productivity differences across products. We now investigate whether the model-predicted relationship between inflation and excess price dispersion is actually present in U.K. micro price data. In previous work, Nakamura et al. (**?**) argue that detecting such a relationship is challenging in U.S. data.

The theory implies (to a second-order approximation) the following steady-state relationship

$$\ln\left(\frac{\Delta_z}{\Delta_z^e}\right) = c_z \cdot \left(\Pi_z^\star - \Pi_z\right)^2, \tag{31}$$

where $\Delta_z/\Delta_z^e \geq 1$ is a measure of excess price dispersion, with Δ_z^e denoting the efficient amount of price dispersion, $c_z \geq 0$ a coefficient that depends on structural parameters (c_z is strictly positive in the presence of nominal rigidities), Π_z the actual inflation rate and $\Pi_z^\star$ the optimal inflation rate at the item level.[19] For the r.h.s. of equation (31) we can use the historic average item inflation rate $E[\Pi_{zt}]$ as a proxy for Π_z and our baseline estimates for the optimal item-level inflation rate Π_z^*, which is available for more than 1000 expenditure items. Letting σ_z denote the empirical analogue of $\ln\left(\Delta_z/\Delta_z^e\right)$ on the l.h.s. of equation (31), as defined below, we then consider regressions of the form

$$\sigma_z = a + b\left(\Pi_z^* - E[\Pi_{z,t}]\right) + c\left(\Pi_z^* - E[\Pi_{z,t}]\right)^2, \tag{32}$$

and test whether $c > 0$ and $b = 0$, as implied by the price setting theory.[20]

Our baseline approach computes σ_z as suggested by the theory, i.e., we compute for each product j in item z the standard deviation of the price deviation from the relative price trend, i.e., the standard deviation of the residuals in equation (21). We then take the median standard deviation across products in item z as our measure of excess price dispersion σ_z. We shall also explore the robustness of our baseline results towards using alternative measures for σ_z.

Table 2 provides the baseline results when estimating equation (32) using the cross-section of items. The table shows that - as predicted by theory - we have a statistically significant value for $c > 0$, while the estimate for b is

[19] More precisely, we have

$$\frac{\Delta_z}{\Delta_z^e} = \int_0^1 \left(\frac{Q_{zt}}{G_{jzt}Q_{zt-s_{jt}}}\right)\left(\frac{P_{jzt}}{P_{zt}}\right)^{-\theta} dj \Big/ \left(\int_0^1 \left(\frac{Q_{zt}}{G_{jzt}Q_{zt-s_{jt}}}\right)^{1-\theta} \mathrm{dj}\right)^{\frac{1}{1-\theta}}$$

and

$$c_z = \frac{1}{2}\theta\left[\frac{\alpha_z(1-\delta_z)(\Pi_z^\star)^{\theta-1}}{(1-\alpha_z(1-\delta_z)(\Pi_z^\star)^{\theta-1})^2}\frac{1}{(\Pi_z^\star)^2}\right] > 0.$$

[20] The theory also implies $a = 0$, but measurement error in the standard deviation will bias the estimate for σ_z upwards, so that testing this restriction is not very informative.

insignificant at conventional significance levels. We also find that the estimated minimum for the excess price dispersion is not statistically significantly different from zero, i.e., the data is consistent with the notion that the minimum excess price dispersion is reached at $E[\Pi_{z,t}] = \Pi_z^*$, as implied by the theory.

The finding of a statistically significant coefficient $c > 0$ turns out to be very robust to using alternative empirical specifications for σ_z. For instance, we obtain $c > 0$ also when using the mean instead of the median standard deviation of residual, when estimating the residuals using sales filtered data, or when allowing for product specific age trends in the specification of equation (21) to compute the residuals. It thus appears that the relationship between suboptimal inflation and excess price dispersion predicted by the theory is actually present in the data. This is an important finding, because the optimal inflation rate precisely seeks to exploit this relationship for improving the relative price distribution and welfare.

Table 2: Estimates for equation (32)

Coefficient	Estimate	t-Statistic
a	0.0288	34.024
b	-0.0235	-1.3127
c	1.3979	4.7303

Notes:

10 Conclusions

The present paper shows how relative price trends at the product level are informative about the optimal inflation target that a welfare-maximizing central bank should pursue. While our empirical approach allows for a rich set of heterogeneity across products, we have abstracted from a number of features that are worthwhile investigating in future work. Given our focus on consumer products, we have abstracted from intermediated and imported goods, both of which would appear to be relevant for the U.K. economy. In fact, introducing heterogeneity in product trends into sticky price models featuring sectoral input-output structures, e.g., Nakamura and Steinsson (2010), Pasten, Schoenle and Weber (**?**), appears to be an interesting avenue for future research. Likewise, it would seem to be of interest to investigate the potential relevance of nonlinear terms for the optimal inflation target, from which the empirical analysis in the present paper has abstracted.

References

Adam, K., and R. Billi (2007): "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 54, 728–752.

Adam, K., and R. M. Billi (2006): “Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates,” *Journal of Money, Credit and Banking*, 38(7), 1877–1905.

Benigno, P. (2004): “Optimal monetary policy in a currency area,” *Journal of International Economics*, 63(2), 293 – 320.

Benigno, P., and L. A. Ricci (2011): “The Inflation-Output Trade-Off with Downward Wage Rigidities,” *American Economic Review*, 101, 1436–1466.

Bergin, P. R., and G. Corsetti (2008): “The extensive margin and monetary policy,” *Journal of Monetary Economics*, 55(7), 1222–1237.

Bilbiie, F. O., I. Fujiwara, and F. Ghironi (2014): “Optimal monetary policy with endogenous entry and product variety,” *Journal of Monetary Economics*, 64, 1–20.

Bilbiie, F. O., F. Ghironi, and M. J. Melitz (2008): “Monetary Policy and Business Cycles with Endogenous Entry and Product Variety,” in *NBER Macroeconomics Annual 2007, Volume 22*, NBER Chapters, pp. 299–353. National Bureau of Economic Research, Inc.

Carlsson, M., and A. Westermark (2016): “Labor market frictions and optimal steady-state inflation,” *Journal of Monetary Economics*, 78, 67 – 79.

Coibion, O., Y. Gorodnichenko, and J. Wieland (2012): “The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?,” *Review of Economic Studies*, 79(4), 1371–1406.

Kahn, A., R. G. King, and A. L. Wolman (2003): “Optimal Monetary Policy,” *Review of Economic Studies*, 70, 825–860.

Kim, J., and F. J. Ruge-Murcia (2009): “How much inflation is necessary to grease the wheels?,” *Journal of Monetary Economics*, 56(3), 365 – 377.

Nakamura, E., and J. Steinsson (2010): “Monetary Non-neutrality in a Multisector Menu Cost Model,” *Quarterly Journal of Economics*, 125, 961–1013.

Office for National Statistics (2014): *Consumer Price Indices - Technical Manual*. Office of Public Sector Information.

Schmitt-Grohé, S., and M. Uribe (2010): “The Optimal Rate of Inflation,” in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3 of *Handbook of Monetary Economics*, chap. 13, pp. 653–722. Elsevier.

——— (2012): “On Quality Bias and Inflation Targets,” *Journal of Monetary Economics*, 59(4), 393 – 400.