Learning By Doing in New Firms and the Optimal Rate of Inflation

Henning Weber*

February 15, 2015

– PRELIMINARY –

Abstract

Empirical data suggest that new firms tend to grow faster than incumbent firms in terms of their productivity. A sticky-price model with learning-by-doing in new firms fits this data and predicts that for plausible calibrations, the optimal long-run inflation rate is positive and between 0.5% and 1.5% per year. A positive long-run inflation rate helps the fast-growing new firms to align their real price with their idiosyncratic productivity growth. In contrast, the standard sticky-price model without learning-by-doing in new firms predicts an optimal long-run inflation rate near zero. In a two-sector model with learning-by-doing in new firms, the policy tradeoff that arises between new and incumbent firms is considerably more severe than the policy tradeoff that arises between economic sectors.

JEL Code: E31, E32, E52, E61.

Keywords: Optimal monetary policy, heterogenous firms, firm-level productivity growth, firm entry and exit.

*Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany, Tel: +49 069 9566 2217, henning.weber@bundesbank.de. This paper represents the author's personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank, the Eurosystem, or its staff. I particularly thank Larry Christiano, Mirko Wiederholt, Maik Wolters, Ctirad Slavik, and the participants in the Frankfurt Seminar in Economics at Goethe University, for useful comments. Any errors are mine.

1 Motivation

Empirical data suggest that productivity growth rates in new firms tend to exceed productivity growth rates in incumbent firms. One prominent explanation for this data is learning by doing, i.e., the idea that firms accumulate knowledge as by-product of producing goods and services. Since firms' learning curves flatten out in the course of time, new firms see their productivity grow faster than incumbent firms (Bahk and Gort (1993)). Another prominent explanation for new firms growing faster than incumbent firms is embodied productivity growth, i.e., the idea that best-practice technology is embodied in new firms and expands at a faster rate than the technology installed in incumbent firms (Jensen, McGuckin, and Stiroh (2001)).

In economies in which firms set their price based on marginal costs, firm-level productivity growth affects a firm's price setting. Ideally, a profit-maximizing firm sets its nominal price in a way that guarantees that its real price exceeds its real marginal costs by some markup. In this ideal case, the real price of new firms with relatively fast firm-level productivity growth falls over time, because firm-level productivity growth reduces firm-level marginal costs. However, when firms adjust their nominal price only infrequently, as empirical data suggest, new firms will find it difficult to reduce their real price in line with their firm-level productivity growth.

A positive long-run inflation rate erodes firms' real prices and, therefore, can help new firms to reduce their real price over time, and a declining real price helps them to avoid distortions that otherwise would arise from infrequent nominal price adjustment. In contrast to new firms, however, a positive long-run inflation rate can affect incumbent firms adversely. In the case in which incumbent firms expand their productivity less than new firms, incumbent firms prefer to reduce their real price less than new firms, or to even increase it. The purpose of this paper is to quantify the long-run inflation rate that optimally resolves the tradeoff between new and incumbent firms, and to identify the factors that shift the optimal long-run inflation rate in favor of either new or incumbent firms.

Using a basic New Keynesian model that incorporates firm-level productivity growth and new firms that expand at a faster rate than incumbent firms, I find that the optimal long-run inflation rate is between 0.5% and 1.5% per year. The model is calibrated to firm-level data on the U.S. economy. The positive optimal long-run inflation rate arises from the learning-by-doing dynamics in new firms. In a version of the model without learning by doing, the optimal long-run inflation rate is between -1.1% and -0.6% per year. This negative optimal long-run inflation rate arises from embodied productivity growth, which makes the technology of incumbent firms to become obsolete at a faster rate than without embodied growth. Therefore, to preserve their profit-maximizing markup, incumbent firms prefer to increase their real price, and a negative long-run inflation rate helps incumbent firms to achieve this when nominal prices are sticky.

In the calibrated model, learning by doing dominates the role of embodied productivity growth for the optimal long-run inflation rate and makes this rate positive. Learning by doing triggers large and rapid changes in firm-level productivity, whereas embodied productivity growth triggers only small and gradual changes in firm-level productivity. Consequently, the price distortions in new firms from learning by doing are larger than the price distortions in incumbent firms from embodied productivity growth and, therefore, the optimal long-run inflation rate is geared towards new firms. Another factor that emphasizes the role of new firms for the optimal long-run inflation rate and, therefore, makes this rate more positive, is a large market share of new firms. Interestingly, a factor that does not influence the optimal long-run inflation rate much is the degree of firms' price stickiness.

Firm-level productivity growth can differ among new and incumbent firms, but it can also differ among economic sectors. For example, while embodied productivity growth accounts for about two thirds of total productivity growth in manufacturing (Sakellaris and Wilson (2004)), embodied productivity growth accounts for basically all productivity growth in retail trade (Foster, Haltiwanger, and Krizan (2006)). Another important difference across sectors is the degree of price stickiness (Bils and Klenow (2004)). Therefore, in order to obtain a reliable estimate of the optimal long-run inflation rate, I also consider a two-sector model, which varies firm-level productivity growth and the degree of price stickiness across economic sectors.

In this model, each sector has its own optimal long-run inflation rate, and this creates a policy tradeoff for the government when it selects the aggregate long-run inflation rate. This tradeoff between sectors arises in addition to the tradeoff within each sector between new and incumbent firms. To resolve the tradeoff between sectors, the government tilts the optimal aggregate long-run inflation rate towards the optimal long-run inflation rate in the sector with the more sticky prices, because thereby it shifts the price adjustment to the sector with the more flexible prices, where it is least distortive. In the calibrated model, however, the tradeoff between new and incumbent firms within a sector is considerably more severe than the tradeoff between sectors.

The literature on the optimal long-run inflation rate has not yet examined the role of firm-level productivity growth, and this paper contributes to close this gap in the literature.[1] Three related papers examine the role of firm-level and sectoral factors, but their mechanisms and results differ from the ones in this paper. Namely, Wolman (2011) examines the role of sectoral productivity growth and finds that the government weighs the sector with stickier prices more heavily and that mild deflation is socially optimal. Schmitt-Grohe and Uribe (2012) examine quality bias in the officially measured inflation rate in one and two-sector models and find that price stability is optimal if non-quality adjusted prices are sticky. Finally, Janiak and Monteiro (2011) examine entry and exit of heterogenous firms in a flexible-price model with a cash-in-advance constraint and find that the long-run inflation rate affects the level of aggregate productivity.

A main finding in this paper is that firm-level productivity growth can justify a positive optimal long-run inflation rate. Recently, Billi (2011) and Coibion, Gorodnichenko, and Wieland (2012), among others, show that the zero lower bound on nominal interest rates can also justify a positive optimal long-run inflation rate. Incorporating firm-level productivity growth into the analysis suggests that the welfare costs of pursuing a positive long-run inflation rate are smaller than what the zero-lower-bound literature estimates.

This paper is also related to the literature on the role of firm entry and exit for optimal monetary policy. Bergin and Corsetti (2008), Bilbiie, Ghironi, and Melitz (2008), Faia (2012), and Bilbiie, Fujiwara, and Ghironi (2011) analyze optimal monetary policy in models with firm entry and exit and sticky prices. However, while these authors use models with aggregate productivity growth and homogenous firms, I use a model with firm-level productivity growth and heterogenous firms.

This paper continues as follows. Section 2 describes the one-sector model. Section 3 contains the calibration of this model, and Section 4 derives the optimal long-run inflation rate. Section 5 extends the model to two sectors and incorporates sectoral asymmetries. Section 6 derives the optimal long-run inflation rate in the two-sector model, and Section 7 concludes.

[1] However, the literature on the optimal long-run inflation rate, which is reviewed in Schmitt-Groh and Uribe (2010), has examined a long list of factors, and I leave many of them out of my analysis in order to focus it on a lean model. Among these factors are monetary and transaction frictions (e.g., Friedman (1969), Aruoba and Schorfheide (2011)), downwardly rigid nominal wages (e.g., Kim and Ruge-Murcia (2009)), or a positive trend growth rate in aggregate productivity (e.g., Amano, Moran, Murchison, and Rennison (2009)).

2 Model

This section describes a monetary model with firm-level productivity growth and with exogenous firm entry and exit. The model features sticky nominal prices and represents a cashless economy without aggregate uncertainty.

2.1 Firms

In order to set up the model, I index firms by $j \in [0,1]$ and let each firm produce a single product variety. The technology of firm j needs labor ℓ_{jt} as the sole input to produce output y_{jt}:

$$y_{jt} = \hat{q}^{t-s_{jt}} \left(\frac{\hat{a}^t \hat{g}^{s_{jt}}}{1+\overline{\lambda}\lambda^{s_{jt}}} \right) \ell_{jt} \ . \tag{1}$$

The integer variable $s_{jt} = 0, 1, 2, \ldots$ indicates the firm's age at time t, and parameters $\hat{a}, \hat{q}, \hat{g}, \overline{\lambda}$, and λ determine how firm-level productivity evolves over the lifetime of the firm. Specifically, $\overline{\lambda} \geq 0$ and $\lambda \in [0,1]$ capture scope and speed of learning by doing, respectively. When $\overline{\lambda}$ is large, the firm begins production with a low level of productivity and experiences large productivity gains over time. Furthermore, when λ is close to zero, the firm learns quickly and, therefore, productivity gains realize fast. Learning by doing yields large increments to learning when the firm is young and diminishing increments to learning when the firm ages. Reasons for (post-entry) learning are managers' accumulating experiences, workers' learning by doing, or economies of scale.[2] Hornstein and Krusell (1996) use a related model of learning by doing.

In equation (1), $\hat{a} \geq 1$ denotes the growth rate in the firm's productivity component that is common to all firms. Furthermore, the term $\hat{q}^{t-s_{jt}}$ indexes the initial level of productivity in a new firm to the firm's date of market entry. Thus, $\hat{q} > 0$ is the growth rate in productivity embodied in new firms. Only new firms use the new technology, i.e., incumbent firms cannot retool. Finally, $\hat{g} > 0$ determines the growth rate in the productivity component that is specific to incumbent firms once the increments to learning approach zero.

I rearrange the technology of firm j in equation (1) according to

$$y_{jt} \quad = \quad \left(\frac{a^t g^{s_{jt}}}{1+\overline{\lambda}\lambda^{s_{jt}}} \right) \ell_{jt} \ ,$$

[2] Cooper and Johri (2002), Foster, Haltiwanger, and Krizan (2006), Rogers, Helmers, and Koch (2010), and Suarez and De Jorge (2012), among others, provide direct evidence for learning by doing and for that new firms grow faster than incumbent firms. This evidence refers to firms, plants, products, product lines, or establishments, and I use the term "firm" to collectively refer to these units.

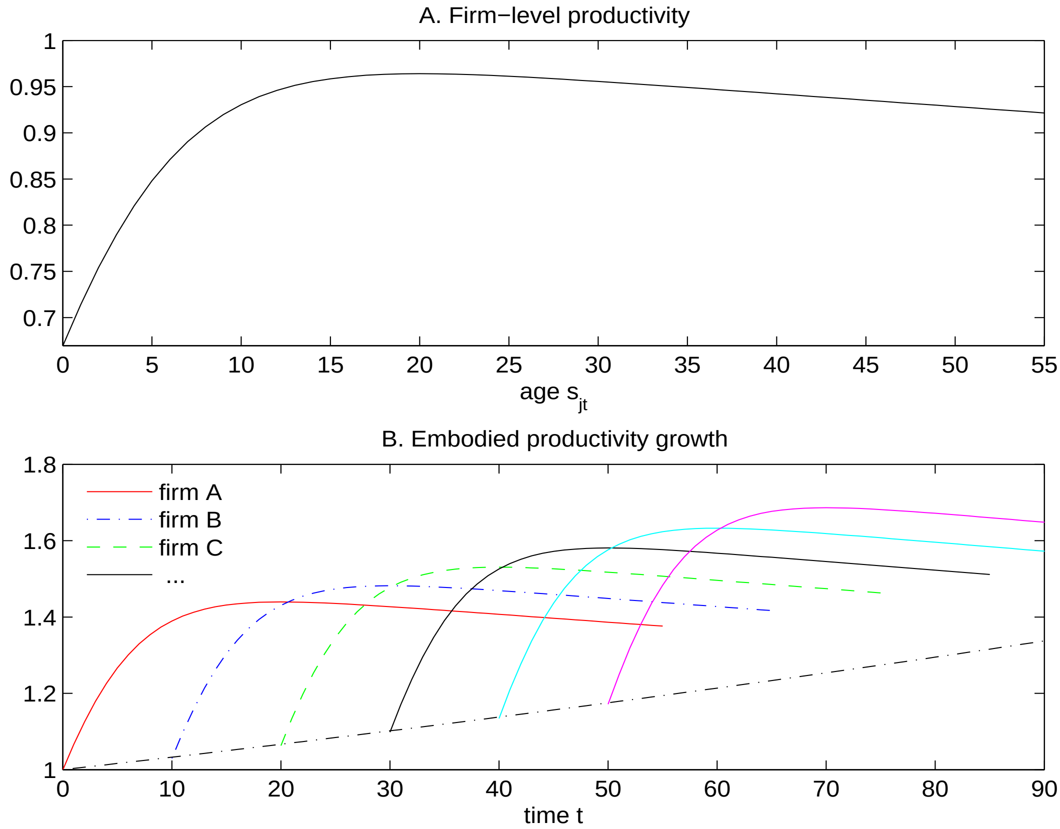

Figure 1: Panel A shows (detrended) firm-level productivity, $(\hat{g}/\hat{q})^{s_{jt}}/(1+\overline{\lambda}\lambda^{s_{jt}})$, over the lifetime of firm j. Panel B shows the role of embodied productivity growth (dashed line), $\hat{q}$, for the level of productivity, $\hat{q}^t$ $(\hat{g}/\hat{q})^{s_{jt}}/(1+\overline{\lambda}\lambda^{s_{jt}})$, in various new firms j. The calibration of parameters is described in Section 3.

with $a = \hat{a}\hat{q}$ and $g = \hat{g}/\hat{q}$. Here, a denotes aggregate productivity growth, which arises from both embodied and common productivity growth, and g denotes the rate at which incumbent firms become obsolete relative to new firms. In the special case when g equals unity and $\overline{\lambda}$ equals zero, all firms are equally productive, as in the basic New Keynesian model derived in, e.g., Woodford (2003).

Panel A of Figure 1 illustrates the role of learning by doing and incumbent and embodied productivity growth for (detrended) firm-level productivity, $(\hat{g}/\hat{q})^{s_{jt}}/(1+\overline{\lambda}\lambda^{s_{jt}})$. When the firm is young, its productivity grows fast as a result of learning by doing. When the firm ages, learning by doing fades away, and the firm's productivity starts to decline relative to the productivity in new firms as a result of embodied productivity growth. Panel B also illustrates the role of embodied productivity growth for the level of productivity in new firms.

Firms enter and exit the economy continuously. At the beginning of a period, $\delta \in [0, 1)$ new firms enter the economy, while at the end of a period, δ firms exit the economy.[3] The exit of firms occurs randomly and, therefore, firms with various levels of productivity are equally exposed to exit. In reality, a firm with high productivity may exit because a major shift in consumer taste occurs, a new regulation is passed, or product liability legislation is changed; because a new firm crowds the established firm out of the market by supplying a close substitute; or because the established firm starts exporting and stops selling at home.

When a new firm enters the economy, it sets a price for its product. In subsequent periods, the firm resets its price with probability $(1-\alpha)$, $\alpha \in [0,1)$, each period until exit. Firm j sets its nominal price P_{jt} to solve

$$\max_{P_{jt}} \sum_{i=0}^{\infty} \kappa^i \; \Omega_{t,t+i} \left[P_{jt} - W_{t+i} \left(\frac{1+\overline{\lambda}\lambda^{s_{jt+i}}}{a^{t+i} g^{s_{jt+i}}} \right) \right] y_{jt+i} \qquad \text{s.t.} \qquad y_{jt+i} = \left(\frac{P_{jt}}{P_{t+i}} \right)^{-\theta} y_{t+i} \; . \tag{2}$$

$\Omega_{t,t+i}$ discounts nominal payoffs, and $\kappa = \alpha(1-\delta)$ is the probability to produce at current prices in the next period. When the firm sets its price, it anticipates the evolution of its productivity. The constraint in the firm's problem is the household demand for product j, derived below, and P_t, W_t, and y_t denote the aggregate price level, the nominal wage, and the aggregate output, respectively. Wages are identical across firms because firms hire labor in a perfectly competitive labor market, as in, e.g., Melitz (2003).

The optimal price of firm j equates the discounted sum of marginal revenues to the discounted sum of marginal costs. I rearrange this condition to obtain:

$$\begin{aligned} \left(\frac{P^{\star}_{jt}}{P_t} \right) g^{s_{jt}} &= \frac{\theta}{\theta-1} \left(\frac{N_t + \overline{\lambda}\lambda^{s_{jt}} N_{\lambda t}}{D_t} \right) \; , \\ N_t &= w_t/a^t + (\kappa/g)\beta\pi^{\theta}_{t+1} N_{t+1} \; , \\ N_{\lambda t} &= w_t/a^t + (\kappa\lambda/g)\beta\pi^{\theta}_{t+1} N_{\lambda t+1} \; , \\ D_t &= 1 + \kappa\beta\pi^{\theta-1}_{t+1} D_{t+1} \; . \end{aligned} \tag{3}$$

The (gross) inflation rate is the change in the aggregate price level, $\pi_t = P_t/P_{t-1}$, and N_t, $N_{\lambda t}$, and D_t denote auxiliary variables. The optimal price of firm j, $P^{\star}_{jt}$, depends on firm-level productivity, which depends on the firm's age, s_{jt}. Consequently, because firms of different age

[3] If I were to consider endogenous firm entry, as in, e.g., Bilbiie, Fujiwara, and Ghironi (2011), the number of firms evolves according to $N_t = (1-\delta)[N_{t-1} + N_{Et-1}]$. In this case, the steady-state fraction of new over all firms, N_E/N, also depends on only the exit rate, $N_E/N = \delta/(1-\delta)$. This suggests that endogenous firm entry adds little to my results on the optimal long-run inflation rate derived for exogenous firm entry.

adjust their price in a given period, adjusting firms set various optimal prices.

2.2 Household

The representative household maximizes discounted lifetime utility:

$$\max_{\{\ell_t, c_{jt}, Q_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left[u(c_t) - h(\ell_t)\right] \ , \quad 0 < \beta < 1 \ , \tag{4}$$

where c_t is aggregate consumption, and ℓ_t is aggregate labor. The functional form of period utility is $u(c) = \log(c)$ and $h(\ell) = \eta_L \ell^{1+\nu}/(1+\nu)$. The household is subject to the budget constraint

$$\Omega_{t,t+1} Q_{t+1} + \int_0^1 P_{jt} c_{jt} \ dj \quad \leq Q_t + (1-\tau_L) W_t \ell_t + V_t + T_t \ . \tag{5}$$

It selects a financial portfolio of nominal claims with payoff Q_{t+1}. The price of this portfolio at date t is $\Omega_{t,t+1} Q_{t+1}$, where $\Omega_{t,t+1}$ is the unique discount factor, to be determined by complete financial markets. The household consumes and it receives $(1-\tau_L) W_t \ell_t$ as labor income net of taxes. While the labor income tax τ_L is not essential for the main results, it will facilitate characterizing them analytically. The household also receives profits V_t from the ownership of firms and a lump-sum transfer T_t from the government. Terminal conditions (not shown) require household solvency. The household's preference for intermediate products is $c_t = (\int_0^1 c_{jt}^{\frac{\theta-1}{\theta}} \ dj)^{\frac{\theta}{\theta-1}}$, with $\theta > 1$. The household's optimization yields the product demand, $c_{jt}/c_t = (P_{jt}/P_t)^{-\theta}$, the cost-minimal price of aggregate consumption, $P_t = (\int_0^1 P_{jt}^{1-\theta} \ dj)^{\frac{1}{1-\theta}}$, and $P_t c_t = \int_0^1 P_{jt} c_{jt} \ dj$.

2.3 Equilibrium and the balanced growth path

In equilibrium in the decentralized economy, firms set prices according to equation (2); the household maximizes lifetime utility (4) subject to the budget constraint (5) and the definition of aggregate consumption c_t; product markets clear at $y_{jt} = c_{jt}$; the labor market clears at $\ell_t = \int_0^1 \ell_{jt} \ dj$; financial markets clear at $Q_t = 0$; the resource constraint $y_t = c_t$ holds; and the government sets τ_L, ensures $T_t = \tau_L W_t \ell_t$, and sets the nominal short-term interest rate i_t, which is the payoff to a one-period nominal bond, $(1+i_t)^{-1} = \beta \Omega_{t,t+1}$, in order to control π_t.

In equilibrium, aggregate variables will grow at constant rates, because there are no aggregate shocks to perturb the balanced growth path of the economy. I assume no population growth at the balanced growth path such that ℓ_t is constant, and that the government maintains a constant long-run inflation rate.

2.4 Aggregation

Firms differ from one another in two dimensions, namely, in the level of their productivity and in the length of their price spell. Differences in the first dimension arise from firm entry and from assuming that firm-level productivity grows over the lifetime of a firm, whereas differences in the second dimension arise from staggered pricing of firms.

To aggregate firms' prices to the aggregate price level, I replace firm index j by two new indices, n and k, each representing one dimension of heterogeneity, and denote the current price of firm j as[4]

$$P_{jt} = P^{\star}_{t-(n+k),t-k} \ , \quad n = 0, 1, 2, \ldots, \ k = 0, 1, 2, \ldots \ .$$

The first subscript, $t-(n+k)$, indicates the date of market entry. The second subscript, $t-k$, indicates the date of the last price change. Thus, index k denotes the length of the price spell, and index n denotes the time between market entry and last price change.

The price level P_t comprises the prices of all cohorts of firms. For the moment, I consider the cohort that entered $s \geq 0$ periods ago, at date $t-s$, and normalize its mass to unity. At date t, the weighted average price of this cohort, $\Lambda_t(s)$, is

$$\Lambda_t(s) = (1-\alpha)\sum_{k=0}^{s-1} \alpha^k (P^{\star}_{t-s,t-k})^{1-\theta} + \alpha^s (P^{\star}_{t-s,t-s})^{1-\theta} \ , \tag{6}$$

if $s \geq 1$, and $\Lambda_t(s) = (P^{\star}_{t,t})^{1-\theta}$ if $s = 0$. Upon entry ($s = 0$), all firms in a cohort s set the same optimal price. At subsequent dates ($s \geq 1$), some firms change their prices, while others keep their price, and therefore the price distribution of the cohort s fans out.

At date t, the mass of cohort s is equal to $(1-\delta)^s\delta$ because firm exit diminishes the cohort's mass over time. Summing over all cohorts s yields the unit mass of firms that underlies the price level, $1 = \sum_{s=0}^{\infty}(1-\delta)^s\delta$. Thus, the price level $P_t^{1-\theta} = \int_0^1 P_{jt}^{1-\theta}\, dj$ is equal to the sum over

[4] This approach is related to Dotsey, King, and Wolman (1999). Unlike in my approach, however, they consider a finite-dimensional state vector of prices and firms with homogenous productivity.

cohort prices $\Lambda_t(s)$, each weighted by the mass $(1-\delta)^s\delta$ of its cohort:

$$P_t^{1-\theta} = \sum_{s=0}^{\infty}(1-\delta)^s\delta\Lambda_t(s) \ . \tag{7}$$

In order to rearrange this equation, I consider the optimal prices of two firms, denoted j and j', at the same date $t-k$. Both firms adjust their price at this date. However, while firm j is a new firm with age $s_{jt-k} = 0$, firm j' is a firm with age $s_{j't-k} = n$. Relating the pricing equations (3) for both firms to one another yields:

$$P^{\star}_{t-k,t-k} \quad = \quad g^n \left(\frac{1+\overline{\lambda}(N_\lambda/N)}{1+\lambda^n\overline{\lambda}(N_\lambda/N)}\right) P^{\star}_{t-(n+k),t-k} \ , \tag{8}$$

where I have used that N_t and $N_{\lambda t}$ are constant at the balanced growth path. The equation states that optimal prices of firms with different age are proportional to one another. Proportionality corresponds to the differential, in terms of expected discounted marginal costs, between incumbent and new firm. Marginal costs differ across firms with different age because these firms maintain different levels of productivity.

Combining equations (6), (7), and (8), and defining the optimal real price of a new firm as $p^\star = P^\star_{t,t}/P_t$ yields the long-run inflation rate as function of $p^\star$:

$$1 \quad = \quad \{\delta + (1-\alpha)\overline{m}\}\,(p^\star)^{1-\theta} + \kappa\pi^{\theta-1} \ . \tag{9}$$

The term in curly brackets indicates that the δ firms that are new at date t maintain the optimal real price $p^\star$. Furthermore, there is a fraction $(1-\alpha)$ of incumbent firms that reset prices optimally, and optimal prices of incumbent firms are equal to the optimal price of new firms after accounting for the differential in marginal costs between new and incumbent firms. Parameter $\overline{m}$, which is equal to

$$\overline{m} \quad = \quad \delta\sum_{n=0}^{\infty}(1-\delta)^{n+1}g^{(n+1)(\theta-1)}\left(\frac{1+\overline{\lambda}(N_\lambda/N)}{1+\lambda^{n+1}\overline{\lambda}(N_\lambda/N)}\right)^{\theta-1} \ , \tag{10}$$

accounts for this differential, as it is a weighted sum of differentials in marginal costs between new firms and incumbent firms of all ages.

Aggregation also involves combining the technology of firms to the aggregate technology. To this end, I combine the technology of firms, labor-market clearing, and product demand, and this yields:

$$y \;=\; \ell/\Delta \;, \tag{11}$$

where aggregate output in the steady state is equal to aggregate output at the balanced growth path divided by aggregate productivity growth, $y = y_t/a^t$. Thus, economic growth is exogenous and arises from common and embodied productivity growth, $\hat{a}$ and $\hat{q}$, respectively.[5] Furthermore, $1/\Delta$ is the endogenous steady-state level of productivity:

$$\Delta = \int_0^1 \left(\frac{1 + \overline{\lambda}\lambda^{s_{jt}}}{g^{s_{jt}}} \right) \left[\frac{P_{jt}}{P_t} \right]^{-\theta} dj \;,$$

which is constant. The term in round brackets, which is absent in the basic New Keynesian model, shows that both, the rate at which incumbent firms become obsolete relative to new firms, $\hat{g}/\hat{q}$, and learning by doing affect the steady-state level of productivity. For example, a large scope for learning by doing ($\overline{\lambda}$ large) depresses the initial level of productivity in new firms and, thereby, $1/\Delta$. The term in square brackets, which also occurs in the basic New Keynesian model, shows the effect of cross-sectional price dispersion.[6] Price dispersion implies that the household consumes an uneven distribution of products, substituting expensive for less expensive products, and this reduces $1/\Delta$. The appendix derives Δ as a function of π.

2.5 Decentralized relative to planned economy

I represent the decentralized economy relative to the planned economy. This representation illustrates the nature of distortions in the decentralized economy and whether or not the government faces a policy tradeoff when it selects the optimal long-run inflation rate. The decentralized economy comprises the aggregate technology (11) and the household's optimality condition $h_\ell(\ell_t)/u_c(y_t) = (1 - \tau_L)w_t$, and rearranging these equations yields:

$$y = R(\pi)\frac{\ell}{\Delta_e} \;, \qquad \frac{h_\ell(\ell)}{u_c(y)} \left(\frac{\mu(\pi)}{1 - \tau_L} \right) = \frac{1}{\Delta_e} \;. \tag{12}$$

[5]While $\hat{g}$ affects firm-level productivity growth, it does not affect aggregate productivity growth as a result of the interaction between firm entry and non-selective firm exit.

[6]With firm-level productivity growth, price dispersion arises not only from staggered pricing, but also from firm-level productivity. Therefore, prices will differ from one another even if they are fully flexible; this is distinct from the price dispersion in Yun (2005), which arises exclusively from staggered pricing. This consequence of firm-level productivity growth helps to improve the model's fit to the large amount of price dispersion observed in micro data.

Table 1: Calibration

Parameter	Value	Description
a	$1.0207^{1/4}$	Aggregate productivity growth
$\hat{a}$	$1.0077^{1/4}$	Common productivity growth
$\hat{g}$	$1.0070^{1/4}$	Incumbent productivity growth
$\hat{q}$	$1.0130^{1/4}$	Embodied productivity growth
λ	0.8085	Speed of learning
$\overline{\lambda}$	0.4933	Scope of learning
δ	0.0180	Firm turnover rate
α	0.5092	Probability to not adjust price
θ	3.8	Steady-state markup
ν	0.25	Labor supply elasticity
β	0.995	Discount factor

Notes: See main text for explanation.

Here, I define the relative price distortion, $R(\pi) = \Delta_e/\Delta$, the markup, $\mu(\pi) = 1/(w\Delta_e)$, and the markup distortion, $\mu(\pi)/(1-\tau_L)$. While the relative price distortion arises only from staggered pricing of firms, the markup distortion arises also from monopolistic competition among firms. The appendix derives $R(\pi)$ and $\mu(\pi)$.

Parameter Δ_e derives from the solution of the planned economy that comprises two equations, which are similar to those characterizing the decentralized economy:

$$y^e = \frac{\ell^e}{\Delta_e}\ , \qquad \frac{h_\ell(\ell^e)}{u_c(y^e)} = \frac{1}{\Delta_e}\ . \tag{13}$$

The first equation shows the aggregate technology in the planned economy. The second equation shows that the planner equates the marginal rate of substituting labor for consumption to the marginal rate of transformation. Furthermore, the planner resolves an important tradeoff at the firm level: while some firms can produce a given amount of a product with less labor than other firms, the household prefers to consume an even distribution of all products. Parameter $1/\Delta^e = (\int_0^1 [g^{s_{jt}}/(1+\overline{\lambda}\lambda^{s_{jt}})]^{(\theta-1)} dj)^{1/(\theta-1)}$ arises from resolving this tradeoff optimally and indicates the efficient amount of output dispersion. Finally, comparing equations (12) and (13) to one another shows that decentralized and planned economy coincide when relative price distortion and markup distortion are equal to unity.

3 Calibration

Parameter values pertaining to firm-level productivity growth affect firms' pricing and, therefore, are likely to matter for the optimal long-run inflation rate. Jensen, McGuckin, and Stiroh (2001) estimate both embodied and incumbent productivity growth in labor productivity of U.S. manufacturing plants, controlling for aggregate productivity growth by time effects. They estimate that the initial productivity level of plants that enter in 1992 is 46.8% higher than the initial productivity level of plants that enter in 1963 (see β_{92} in their Table 2). This estimate implies a rate $\hat{q}$ of embodied productivity growth equal to 1.3% per year. To estimate the rate of incumbent productivity growth, the authors use a sample in which a cohort contains only plants that survive the entire sample period. Based on this sample, they estimate that the 1967 cohort shows a 18.7 % productivity gain in 1992 relative to 1967 (see λ_5 in their Table 3). This estimate implies a rate $\hat{g}$ of incumbent productivity growth equal to 0.7% per year.

The rate a of aggregate productivity growth is set to 2.07% per year, based on estimates in Basu, Fernald, and Kimball (2006) for non-durable manufacturing (see their Table 2). This rate implies that embodied productivity growth accounts for about two-thirds of aggregate growth, which is consistent with results in, e.g., Sakellaris and Wilson (2004). Also, using $a = \hat{a}\hat{q}$, the rate $\hat{a}$ of common productivity growth is equal to 0.77% per year.

The scope of learning by doing, $\overline{\lambda}$, is set equal to 0.49. This value yields a progress ratio, i.e., the total increase in a firm's productivity from learning, $1 + \bar{\lambda}$, that falls well into the range of industry estimates obtained in Jovanovic and Nyarko (1995) (see their Table 1). This value also yields an employment size of new firms (aged one year or younger) relative to incumbent firms equal to 60%. According to Miranda, Klimek, and Jarmin (2004), p.10, however, the relative size of new U.S. manufacturing firms (aged five years or younger) is equal to 35%.[7] Thus, the calibrated model overestimates the relative size of new firms. This is plausible because there are certainly further determinants of the relative size of new firms in addition to learning by doing; for instance, Foster, Haltiwanger, and Syverson (2012) emphasize firms' learning about demand.

The speed of learning by doing, λ, is set so that it takes two years for a firm to close 75% of its learning gap, i.e., the firm's productivity level after learning is completed minus its initial productivity level, $1-1/(1+\overline{\lambda})$. This speed of learning is larger than the (slightly above) one year

[7] To obtain meaningful results when I impose a relative size of new firms aged five years or younger, as in Miranda, Klimek, and Jarmin (2004), I have to reduce the speed of learning to unreasonbly low numbers.

used in Hornstein and Krusell (1996), but it is smaller than the three years used in Yorukoglu (1998). This speed of learning may also be compared to case studies, which provide compelling evidence for learning by doing.[8] However, case studies usually examine well-defined learning tasks (assembling a car, say), whereas examining the introduction of new products requires taking a broader perspective on learning (one that also involves organisational learning, say).

Taking one quarter as the time period in the model, I calibrate the remaining parameters as follows. I set the rate of firm turnover δ to 7% per year so that the fraction of new firms (aged five years or younger) over all firms is equal to 30%, as reported in Miranda, Klimek, and Jarmin (2004), p.9.[9] Also, I set the probability α for a firm not to adjust its price so that it yields a median price duration in the truncated price distribution equal to two quarters, as in Wolman (2011). Furthermore, I set θ equal to 3.8, which yields a static markup of 36%, as in Bilbiie, Ghironi, and Melitz (2012).[10] Finally, I set the discount factor β to 0.995, which implies a 4% annual real interest rate after accounting for aggregate productivity growth, and I set the labor-supply elasticity ν equal to 0.25 and η_L equal to 3. Table 1 summarizes the calibration.

4 The optimal long-run inflation rate

The policy problem that I solve to derive the optimal long-run inflation rate consists of the government that uses a restricted set of policy instruments, i.e., the long-run inflation rate and the labor income tax, to maximize steady-state welfare.

4.1 The model without learning by doing

I begin to compute the optimal long-run inflation rate in the model without learning by doing. This model corresponds to an economy with only firms that realize no further increments in learning, and allows me to derive the following result analytically.

Proposition 1: In the model without learning by doing, in which $\overline{\lambda} = 0$, the optimal long-run

[8] E.g., Levitt, List, and Syverson (2012) examine learning by doing in an automobile assembly plant, and refer to numerous further case studies.

[9] To convert annual into quarterly rates, I solve $0.07 = \delta \sum_{s=0}^{3} (1-\delta)^s$ to account for firm exit throughout the year.

[10] Estimated models with fixed costs in production or firm entry imply even smaller values of θ and, thus, even larger static markups. For instance, Smets and Wouters (2007) estimate a value of θ equal to 2.67 in a model with fixed costs in production, and Lewis and Poilly (2012) estimate a value of θ equal to 2.6 in a model with firm entry.

inflation rate that maximizes steady-state welfare is:

$$\pi^\star = \hat{g}/\hat{q} \ ,$$

and the optimal labor income tax is equal to $\tau_L^\star = -1/(\theta - 1)$. With the optimal policy, the decentralized economy (12) coincides with the planned economy (13) and, therefore, is first best.

My calibration implies that in the model without learning by doing, $\pi^\star$ is negative and equal to -0.59% per year. The key equation to understand this first result is the firms' pricing equation (3). For a new firm without learning by doing and denoting $w = w_t/a^t$, this equation can be rearranged as

$$0 = \sum_{s=0}^{\infty} (\kappa\beta\pi^\theta)^s \left[\frac{p^\star}{\pi^s} - \frac{\theta}{\theta-1} \frac{w}{(\hat{g}/\hat{q})^s} \right] \ . \tag{14}$$

Square brackets contain the difference between the (constrained) optimal real price, $p^\star/\pi^s$, and the desired real price, $\frac{\theta}{\theta-1}\frac{w}{(\hat{g}/\hat{q})^s}$, i.e., static markup times firm-level marginal costs. Firm-level marginal costs increase over time because the firm's technology becomes obsolete at rate $\hat{g}/\hat{q}$ relative to the technology embodied in the new firms to come. Accordingly, when the firm can adjust its nominal price only infrequently, it will also find it difficult to adjust its real price in line with its real marginal costs.[11]

The negative optimal long-run inflation rate, $\pi^\star = \hat{g}/\hat{q}$, helps the firm to increase its real price over time and, thereby, to continuously align it with its increasing real marginal costs.[12] Consequently, the firm has no reason to actually change its nominal price, and this prevents distorted relative prices. Furthermore, the firm continuously maintains the static markup and, therefore, the optimal labor income tax remedies the markup distortion. This is shown in Figure 2, which contains the distortions $\mu(\pi)/(1-\tau_L^\star)$ and $R(\pi)$ and ticks $\pi^\star$ (bold lines). Evidently, the optimizing government faces no policy tradeoff, because it can eliminate both distortions and, therefore, achieves the same allocation as in the planned economy.

Panel A in Figure 2 also shows that for π either sufficiently above or below $\hat{g}/\hat{q}$ (ticked),

[11]Proposition 1 does not hinge on assuming time-dependent pricing, since other pricing assumptions also will not interfere with the government's ability to reconcile decentralized and planned economy.

[12]This negative inflation rate arises from new firms. In the model without learning by doing, they set their nominal price to below the average price because their productivity exceeds the average level of productivity. In contrast, incumbent firms do not create any inflation because they keep their nominal price constant.

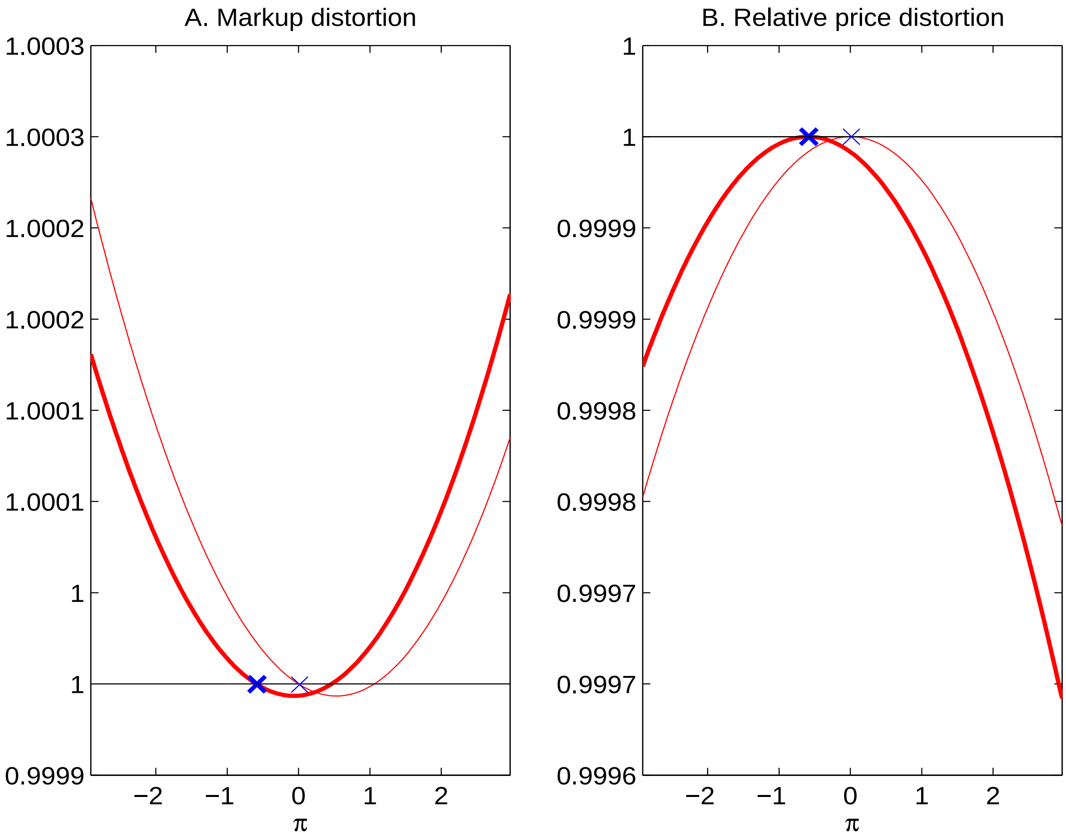

Figure 2: Distortions in the model without learning by doing, but with incumbent and embodied productivity growth (bold lines). Panel A shows the markup distortion $\mu(\pi)/(1-\tau_L^\star)$, and Panel B shows the relative price distortion $R(\pi)$. In both panels, π is the annualized net inflation rate. Thin lines indicate the distortions in the basic New Keynesian model.

the markup $\mu(\pi)$ exceeds the static markup, which is equal to $(1-\tau_L^\star)$. When π is sufficiently above $\hat{g}/\hat{q}$, adjusting firms set a higher nominal price than otherwise, because they anticipate π to excessively erode their real price over time and, thereby, to compress their markup to below the static markup. The higher prices of adjusting firms elevate $\mu(\pi)$. Further, when π is below $\hat{g}/\hat{q}$, firms that cannot adjust their price see their marginal costs decline at a faster rate than their real price, and this also elevates $\mu(\pi)$.[13]

Similarly, Panel B in Figure 2 shows that relative prices are distorted when π deviates from $\hat{g}/\hat{q}$. In this case, firms do not manage to continuously realize the static markup and, therefore, adjust their price whenever they can. This disperses relative prices because only a subset of firms adjust their price in each period. Panels A and B also contain the distortions in the basic New Keynesian model, with $\overline{\lambda}=0$ and $\hat{g}/\hat{q}=1$ (thin lines). They show that in the model without learning by doing, firm-level productivity growth merely re-centers the distortions.

[13]King and Wolman (1999) describe similar effects in the basic New Keynesian model.

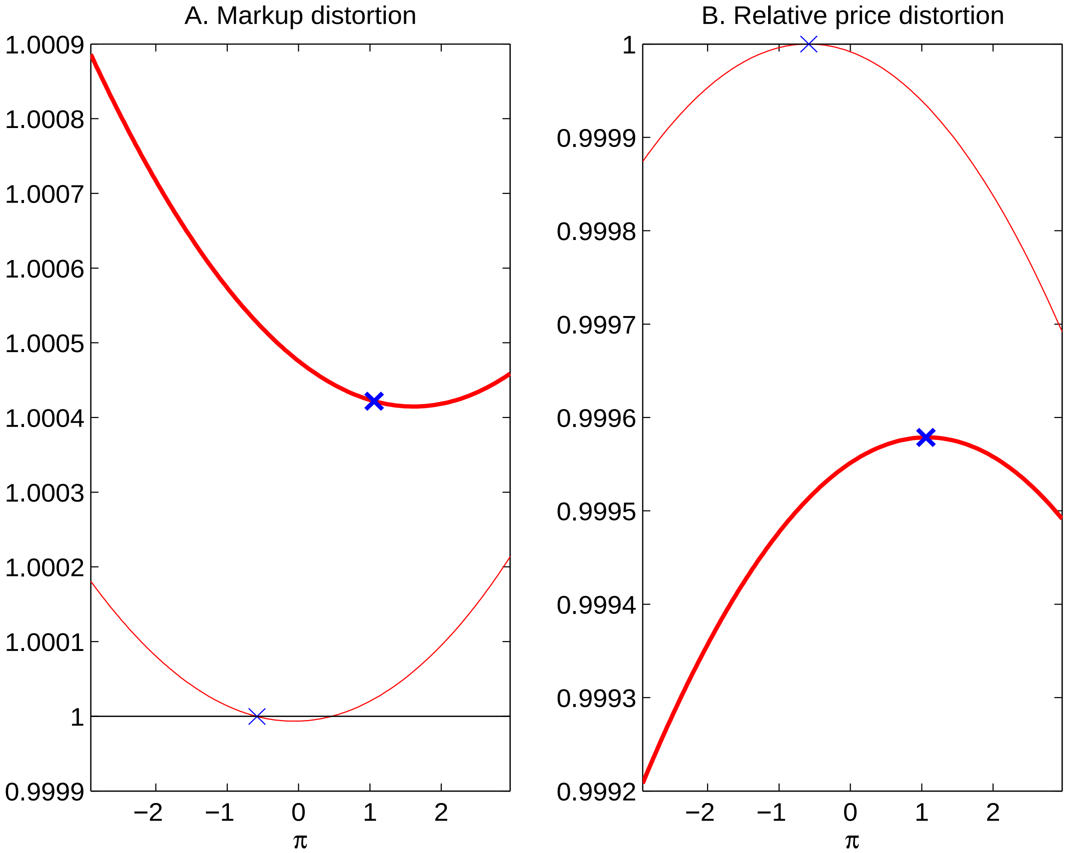

Figure 3: Distortions in the model with learning by doing and incumbent and embodied productivity growth (bold lines). Panel A shows the markup distortion $\mu(\pi)/(1-\tau_L^\star)$, and Panel B shows the relative price distortion $R(\pi)$. In both panels, π is the annualized net inflation rate. Thin lines indicate the distortions in the model without learning by doing.

4.2 The model with learning by doing

In the model with learning by doing, the optimal long-run inflation rate is positive and equal to 1.07% per year, which is the second result in this paper. Thus, incorporating learning by doing into the model increases the optimal long-run inflation rate by 1.66 percentage points, i.e., from −0.59% to 1.07%. With learning by doing, however, $\pi^\star$ can no longer be computed analytically. Thus, I compute it numerically by maximizing steady-state welfare keeping $\tau_L^\star = -1/(\theta - 1)$.

To explain this substantial increase in the optimal long-run inflation rate, Figure 3 plots the distortions $\mu(\pi)/(1-\tau_L^\star)$ and $R(\pi)$ in the models with learning by doing (bold lines) and without learning by doing (thin lines). In contrast to the model without learning by doing, no long-run inflation rate can eliminate the distortions in the model with learning by doing. In this model, thus, the government faces a policy tradeoff, and the figure shows that this tradeoff is resolved optimally by a positive, instead of a negative, optimal long-run inflation rate (ticked).

The policy tradeoff arises because new firms, which experience learning-by-doing dynamics, prefer a positive long-run inflation rate. However, new firms coexist with incumbent firms, which no longer learn and thus prefer a negative long-run inflation rate, as shown in Proposition 1. Thus, the government can use the long-run inflation rate to help either new or incumbent firms, but it cannot help both of them at the same time.

To illustrate this, I revisit pricing equation (3) of a new firm with learning by doing:

$$0 = \sum_{s=0}^{\infty} (\kappa\beta\pi^{\theta})^s \left[\frac{p^\star}{\pi^s} - \frac{\theta}{\theta - 1} \left(\frac{1 + \overline{\lambda}\lambda^s}{(\hat{g}/\hat{q})^s} \right) w \right] . \tag{15}$$

The firm sees its real marginal costs decline over its lifetime s, as a result of learning by doing, which is governed by $\overline{\lambda}$ and λ^s. When the firm can adjust its nominal price only infrequently, it will find it difficult to reduce its real price, $p^\star/\pi^s$, in order to align it with its declining real marginal costs. The positive optimal long-run inflation rate helps the firm to reduce its real price over time. However, when learning-by-doing dynamics fade away as the firm ages, the fact that the firm's technology becomes obsolete at rate $\hat{g}/\hat{q}$ relative to the technology of newer firms begins to dominate the firm's pricing. In this situation, the firm prefers a negative long-run inflation rate, as shown in Proposition 1.

The role of the positive optimal long-run inflation rate for the firm's pricing is shown in Panel A in Figure 4, which contains a sample path of the firm's real price (bars) and the evolution of firm-level marginal costs (solid line) over the lifetime of the firm. Panel B shows the corresponding plot for the case without inflation. When the firm is young, the positive long-run inflation rate in Panel A reduces the gap between the firm's real price and its marginal costs. However, when the firm ages, the positive long-run inflation rate widens this gap. While a positive long-run inflation rate erodes the firm's real price continuously, the zero long-run inflation rate in Panel B implies that the firm's real price changes only at those times when the firm's nominal price changes.

But why does the optimizing government weigh new firms more than incumbent firms and, therefore, selects a positive long-run inflation rate? Figure 4 shows that learning by doing triggers rapid changes in marginal costs of new firms, whereas incumbent and embodied productivity growth trigger only gradual changes in marginal costs of incumbent firms. Accordingly, nominally sticky prices, which prevent a firm from aligning its real price with its marginal costs, constrain new firms more than incumbent firms. Therefore, the optimal long-run inflation rate is geared towards helping new firms, and this holds true despite the fact that learning by doing dominates

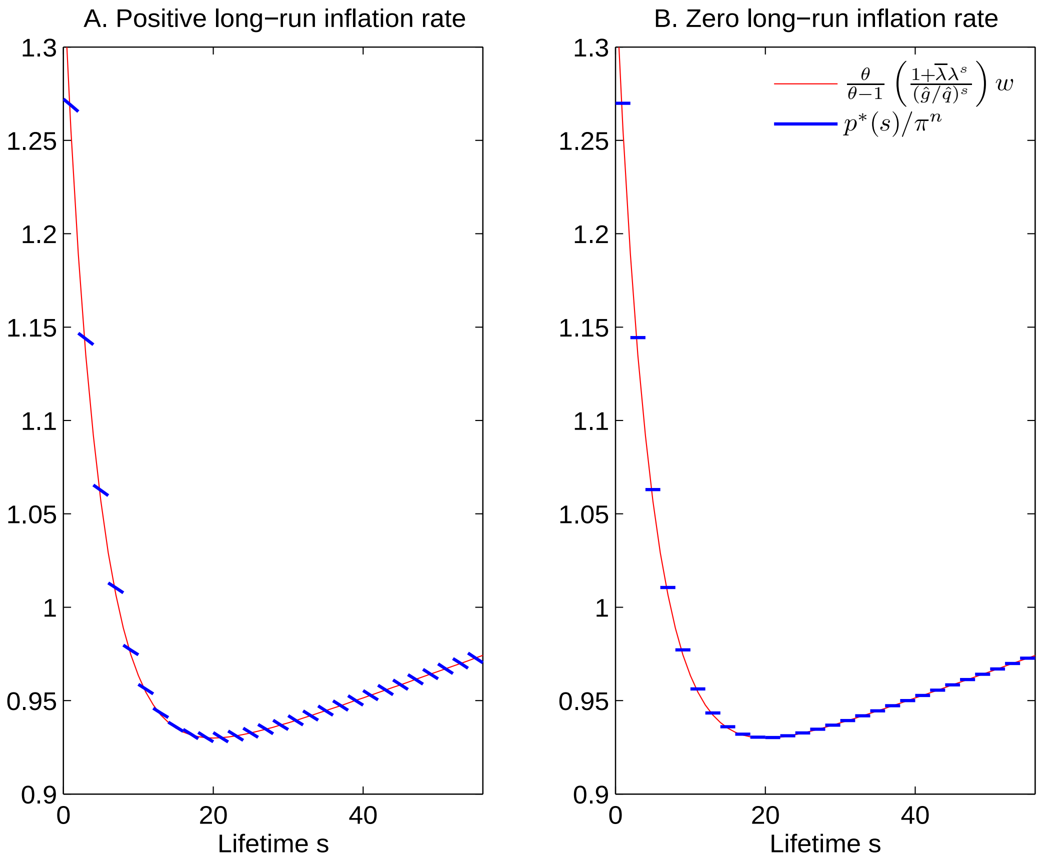

Figure 4: Firm-level marginal costs (solid line) and a sample path of a firm's real price (bars) over the lifetime of the firm. s denotes the age of the firm. Panel A shows the case when the long-run inflation rate is 1.07% per year. Panel B shows the case when the long-run inflation rate is zero. The firm's nominal price is assumed to adjust every two quarters.

the dynamics of firm-level marginal costs for only a relatively short period, about one third, of the lifetime of the average firm.

4.3 Robustness

Calibration of model parameters is crucial to quantify the optimal long-run inflation rate, and I now explore how robust the optimal long-run inflation rate is with respect to changing $\hat{g}/\hat{q}, \lambda, \overline{\lambda}, \alpha, \delta$, and θ. Each panel in Figure 5 shows the optimal long-run inflation rate when one parameter is varied keeping other parameters fix at their benchmark value, except for $\overline{\lambda}$ and λ, which (unless they are varied independently) are set to keep the relative size of new firms at 60% and the speed of learning at two years, respectively, as in Section 3. I find that parameters governing learning by doing, firm turnover, and the price elasticity of product demand are important to determine $\pi^\star$, whereas other parameters, including the one for price stickiness, are less important.

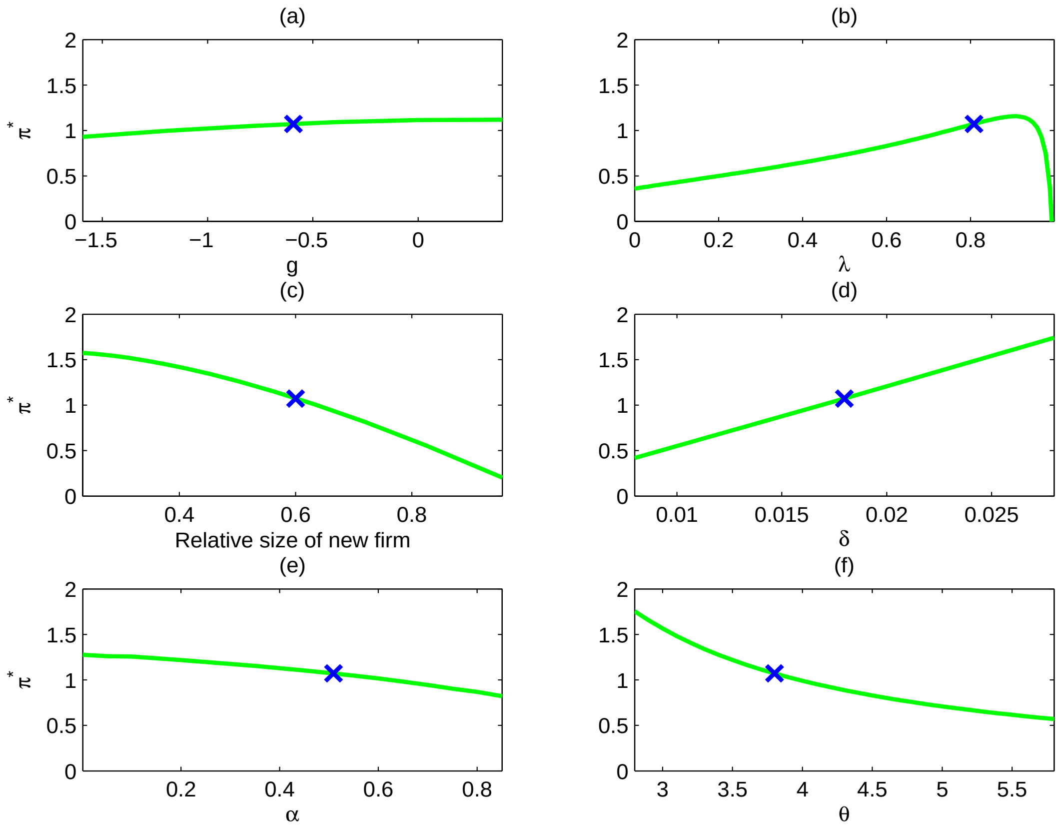

Figure 5: Robustness of the optimal long-run inflation rate with respect to various parameters. $\pi^\star$ and g are annualized net growth rates. In Panel (b), $\overline{\lambda}$ is kept fix at its benchmark value in order to compute $\pi^\star$.

Panel (a) in Figure 5 shows that $\pi^\star$ is fairly insensitive to the rate $\hat{g}/\hat{q}$, at which incumbent firms become obsolete relative to new firms. Qualitatively, reducing this rate to below its benchmark value (ticked) also reduces $\pi^\star$, because when $\hat{g}/\hat{q}$ is reduced, incumbent firms prefer to increase their real price at a faster rate, and this reduces the optimal long-run inflation rate.

Panel (b) shows that $\pi^\star$ increases initially and falls subsequently when the speed of learning is reduced, which corresponds to increasing λ from zero to unity. While $\pi^\star$ is positive already for instantaneous learning ($\lambda = 0$), reducing the speed of learning lengthens firms' learning period and, thereby, increases $\pi^\star$. However, reducing the speed of learning further increasingly flattens firms' learning curve and, therefore, reduces $\pi^\star$. In the boundary case without learning ($\lambda = 1$), $\pi^\star$ converges to $\hat{g}/\hat{q}$ (not shown). Furthermore, Panel (c) shows that $\pi^\star$ increases when the employment size of new relative to incumbent firms falls. In the model, reducing the relative size of new firms corresponds to increasing the scope of learning $\overline{\lambda}$. This amplifies the learning dynamics in new firms and, thereby, increases $\pi^\star$.

Panels (d) and (f) show that $\pi^\star$ increases when the rate δ of firm turnover increases and the price elasticity θ of product demand falls, respectively. High firm turnover increases the portion of new firms in the market and thereby the market share of new firms. Thus, the optimizing government attributes more weight to new versus incumbent firms and, hence, increases $\pi^\star$. Furthermore, price-inelastic product demand implies that households find it difficult to substitute away from the relatively expensive products of new firms, and this also preserves the market share of new firms with similar effects for $\pi^\star$.

Finally, Panel (e) shows that $\pi^\star$ decreases only moderately when the amount of price stickiness is increased.[14] Prices of new firms are more flexible than prices of incumbent firms, because a firm sets its price in its first period unconstrained, whereas this firm is subject to a sticky price with likelihood α in each subsequent period. When prices are flexible (α small), this asymmetry between new and incumbent firms is quantitatively unimportant. However, when prices are sticky, this asymmetry matters and tends to increase the amount of price stickiness in incumbent firms relative to new firms. This shifts $\pi^\star$ in favor of incumbent firms and, therefore, reduces it. The finding that $\pi^\star$ is nevertheless fairly insensitive to the amount of price stickiness differs from the related literature, in which price stickiness often is a core determinant of $\pi^\star$.

5 Sectoral asymmetries

Differences in firm-level productivity growth between new and incumbent firms coexist with differences in magnitude and composition of productivity growth between economic sectors. For example, productivity growth in manufacturing (Goods) is about 2% per year, whereas productivity growth in retail trade (Services) is about 1% per year. As another example, embodied productivity growth accounts for about two-thirds of sectoral productivity growth in Goods, whereas it accounts for basically all sectoral productivity growth in Services. A further difference between economic sectors is the amount of price stickiness.

To refine my estimate of the optimal long-run inflation rate, I incorporate these sectoral asymmetries into my analysis by extending it to a two-sector model. Firms in one sector differ

[14] Nevertheless, increasing the amount of price stickiness substantially increases the welfare costs of price stickiness. One measure of these costs is the fraction ϵ of steady-state consumption that the household in the planned economy is willing to give up to be as well off as in the decentralized economy with optimal policy, i.e., $u(y^e(1-\epsilon)) - h(\ell^e) = u(y) - h(\ell)$, which yields $\epsilon = 1 - \exp(\log(y) - h(\ell) - [\log(y^e) - h(\ell^e)])$. This measure increases from 0.003%, when the median price duration is 1.1 quarters ($\alpha = 0.09$), to 0.415%, when the median price duration is 5 quarters ($\alpha = 0.81$).

from the firms in the other sector in terms of their productivity growth, their degree of price stickiness, and their likelihood to survive. Such asymmetries are not only a realistic feature, but the literature has also shown that they can imply important policy tradeoffs.

5.1 Firms and household

As stated above the model now has two sectors, $z = 1, 2$, and each sector contains many firms that produce intermediate products. Firms in a sector z enter and exit continuously at the rate $\delta_z \in [0, 1)$, and exiting firms are drawn randomly. Firm $j \in [0, 1]$ in a sector z uses the technology $y_{zjt} = (a_z^t g_z^{s_{zjt}} \ell_{zjt})/(1 + \overline{\lambda}_z \lambda_z^{s_{zjt}})$, with $a_z = \hat{a}_z \hat{q}_z$ and $g_z = \hat{g}_z/\hat{q}_z$. Here, a_z denotes sectoral productivity growth, $\hat{a}_z$, $\hat{q}_z$, and $\hat{g}_z$ denote common, embodied, and incumbent productivity growth in sector z, respectively, and $\overline{\lambda}_z$ and λ_z denote scope and speed of learning in sector z, respectively. Firm j's pricing problem is analogous to the one in equation (2), after incorporating the sectoral asymmetries, one of which is the probability to produce tomorrow at current prices, $\kappa_z = \alpha_z(1 - \delta_z)$. Further, firm j hires labor ℓ_{zjt} in an economy-wide, competitive labor market.

The household uses the preference $c_t = c_{1t}^{\psi}\, c_{2t}^{1-\psi}$, with $\psi \in (0, 1)$, to combine consumption in a sector z, c_{zt}, to aggregate consumption c_t. The corresponding price level equals $P_t = (P_{1t}/\psi)^{\psi}\,(P_{2t}/(1-\psi))^{1-\psi}$. The household also uses the preference $c_{zt} = (\int_0^1 c_{zjt}^{\frac{\theta-1}{\theta}}\, dj)^{\frac{\theta}{\theta-1}}$, with $\theta > 1$, to combine the intermediate products to c_{zt}. The corresponding price level in a sector z equals $P_{zt} = (\int_0^1 P_{zjt}^{1-\theta}\, dj)^{\frac{1}{1-\theta}}$. The intertemporal problem that the household solves corresponds to the one described in Section 2.2.

5.2 Decentralized relative to planned economy

Along the lines of the one-sector model, I represent the decentralized economy with two sectors relative to the planned economy with two sectors. To this end, I let $p_z^{\star} = P_{z,t,t}^{\star}/(\eta_z^t P_t)$ denote the relative price of a new firm in a sector z, $p_z = P_{zt}/(\eta_z^t P_t)$ the relative price in a sector z, and $\pi_z = P_{zt}/P_{zt-1}$ the inflation rate in this sector. The appendix shows that these variables are constant. Parameters $\eta_1 = (a_2/a_1)^{1-\psi}$ and $\eta_2 = (a_1/a_2)^{\psi}$ represent trends in P_{zt}/P_t, which depend on the sectoral productivity growth differential a_2/a_1. I also let $\pi = P_t/P_{t-1}$ denote the long-run inflation rate, which I assume is constant.

Like in the one-sector model, the decentralized economy with two sectors consists of the aggregate technology, the intratemporal household optimality condition, and two aggregate distortions that are indexed by the long-run inflation rate:

$$y = R(\pi)\frac{\ell}{\Delta_e} \ , \qquad \frac{h_\ell(\ell)}{u_c(y)}\left(\frac{\mu(\pi)}{1-\tau_L}\right) = \frac{1}{\Delta_e} \ . \tag{16}$$

$R(\pi)$ denotes the aggregate relative price distortion, $\mu(\pi)$ denotes the aggregate markup, and $\mu(\pi)/(1-\tau_L)$ denotes the aggregate markup distortion. When both aggregate distortions are equal to unity, decentralized and planned economy coincide with one another, as in the one-sector model.

The aggregate distortions are functions of the sectoral relative price distortion, $\rho_z(\pi)$, and the sectoral markup, $\mu_z(\pi)$, with $z = 1, 2$:

$$R(\pi) = \left[\psi\left(\frac{\mu_2(\pi)}{\mu_1(\pi)}\right)^{1-\psi}\rho_1(\pi)^{-1} + (1-\psi)\left(\frac{\mu_1(\pi)}{\mu_2(\pi)}\right)^{\psi}\rho_2(\pi)^{-1}\right]^{-1} , \tag{17}$$

$$\mu(\pi) = \mu_1(\pi)^{\psi}\mu_2(\pi)^{1-\psi} \ . \tag{18}$$

The aggregate markup is a weighted geometric mean of the sectoral markup. This is defined as $\mu_z(\pi) = p_z/(w\Delta_z^e)$, indicating with $1/\Delta_z^e$ the efficient amount of output dispersion in a sector z. Furthermore, the aggregate relative price distortion is a weighted mean of the sectoral relative price distortion defined as $\rho_z(\pi) = \Delta_z^e/\Delta_z$. The weights depend on ψ and $1-\psi$ and on the ratio of sectoral markups, which is proportional to the relative price p_2/p_1. Thus, uneven sectoral markups distort the relative price of sectoral consumption and, therefore, the allocation of the household's expenditure across sectors. This source of price dispersion is absent in the one-sector model. Functions $\mu_z(\pi)$ and $\rho_z(\pi)$ are derived in the appendix.

5.3 Calibration of sectoral asymmetries

Productivity growth in the Goods sector 2 is calibrated as the one-sector model in Table 1. To calibrate productivity growth in the Services sector 1, I use evidence from firm-level data on the U.S. retail trade industry in Foster, Haltiwanger, and Krizan (2006). They estimate that labor productivity in this industry grows by 11.43% between 1987 and 1997, and that this increase arises almost exclusively from productivity growth embodied in new firms. These estimates imply that the rates a_1 and $\hat{q}_1$ of sectoral and embodied productivity growth, respectively, are equal to 1.08% per year both. Also, using $a_1 = \hat{a}_1\hat{q}_1$, the rate $\hat{a}_1$ of common productivity growth is equal to zero percent. Furthermore, I set the rate $\hat{g}_1$ of incumbent productivity growth equal to 0.18%

Table 2: Calibration of sectoral asymmetries

Parameter	Value	Description
a_1	$1.0108^{1/4}$	Aggregate productivity growth
$\hat{a}_1$	1	Common productivity growth
$\hat{g}_1$	$1.0018^{1/4}$	Incumbent productivity growth
$\hat{q}_1$	$1.0108^{1/4}$	Embodied productivity growth
λ_1	0.8157	Speed of learning
$\overline{\lambda}_1$	0.3677	Scope of learning
δ_1	0.0315	Firm turnover rate
α_1	0.6883	Probability to not adjust price
ψ	0.60	Relative size of Services sector

Notes: See main text for explanation.

per year, based on the within share in the productivity decomposition in Foster, Haltiwanger, and Krizan (2006) (see their Table 3).

The scope of learning in Services, $\overline{\lambda}_1$, is set equal to 0.37 such that the relative size of new firms in Services is larger by a factor 1.25 than the relative size of new firms in Goods. This factor is taken from Miranda, Klimek, and Jarmin (2004), p.10, and yields in the model a relative size of new firms (aged one year or younger) in Services equal to 75%. Furthermore, the speed of learning in Services, λ_1, is set so that it takes two years for a firm to close 75% of its learning gap, as in the Goods sector.

The size ψ of the Services sector relative to the Goods sector is set equal to 60%, as in Wolman (2011). The rate of firm turnover in the Service sector, δ_1, is set to 12% per year so that the fraction of new firms (aged five years or younger) over all firms is equal to 47%, as reported in Miranda, Klimek, and Jarmin (2004), p.9. Thus, the economy-wide firm turnover rate, $\psi\delta_1 + (1-\psi)\delta_2$, is equal to 10% per year, as in Bilbiie, Ghironi, and Melitz (2012). Further, the probability α_1 for a firm not to adjust its price is set to obtain a median price duration in the truncated price distribution equal to three quarters, as in Wolman (2011). Thus, as shown in Bils and Klenow (2004), Services prices are stickier than Goods prices. Table 2 summarizes this calibration. All remaining parameters are as in Table 1.

6 The optimal long-run inflation rate with sectoral asymmetries

6.1 The two-sector model without learning by doing

I first derive the optimal long-run inflation rate, which optimizes steady-state welfare, in the special case in which there is no learning by doing and the discount factor β approaches unity.

Optimizing steady-state welfare in this case is equivalent to optimizing only one of the two aggregate distortions in the decentralized equilibrium (16)–(18) because these distortions are inversely equal to one another. Optimizing only one of the two aggregate distortion instead of optimizing steady-state welfare simplifies deriving analytical results. Thus, a third result in this paper follows from minimizing $\mu(\pi)$ (or maximizing $R(\pi)$) and shows how the optimizing government resolves a policy tradeoff that arises between asymmetric sectors.

Proposition 2: In the two-sector model without learning by doing, $\overline{\lambda}_z = 0$, and in which $\beta \rightarrow 1$, the optimal long-run inflation rate solves:

$$0 \;=\; \omega(\pi^\star)\left(\frac{\pi^\star - (g_1/\eta_1)}{(g_1/\eta_1)}\right) + [1-\omega(\pi^\star)]\left(\frac{\pi^\star - (g_2/\eta_2)}{(g_2/\eta_2)}\right) \;, \tag{19}$$

with $\eta_1 = (a_2/a_1)^{(1-\psi)}$ and $\eta_2 = (a_1/a_2)^{\psi}$. The weight fulfills the condition that $\omega(\pi) \in [0,1]$ and depends on the long-run inflation rate:

$$\omega(\pi) \;=\; \left[1 + \left(\frac{1-\psi}{\psi}\right)\left(\frac{\kappa_2}{\kappa_1}\right)\left(\frac{a_1}{a_2}\right)^{\theta-1}\left(\frac{1-\kappa_1(\eta_1\pi)^{\theta}/g_1}{1-\kappa_2(\eta_2\pi)^{\theta}/g_2}\right)\left(\frac{1-\kappa_1(\eta_1\pi)^{\theta-1}}{1-\kappa_2(\eta_2\pi)^{\theta-1}}\right)\right]^{-1} \;,$$

with $\kappa_z = \alpha_z(1-\delta_z)$ and $z = 1, 2$.

Equation (19) shows that in the two-sector model without learning by doing, the government faces a policy tradeoff between a long-run inflation rate equal to either g_1/η_1 or g_2/η_2, and it resolves this tradeoff optimally using $\omega(\pi)$.[15] In contrast to Proposition 1, thus, the optimal long-run inflation rate in Proposition 2 generally does not recover the first-best allocation in the two-sector model without learning by doing. The policy tradeoff arises from a lack of policy instruments that work at the sectoral level.[16] Namely, while the government can use the long-run inflation rate to fully offset the distortions in either sector 1 or sector 2, this instrument is not able to fully offset the distortions in both sectors at the same time.

[15] Alternatively, using equation (19), $\pi^\star$ corresponds to a weighted harmonic mean:

$$\pi^\star = \left(\frac{\omega(\pi^\star)}{g_1/\eta_1} + \frac{1-\omega(\pi^\star)}{g_2/\eta_2}\right)^{-1} \;.$$

[16] One special case in which the policy tradeoff disappears and the decentralized two-sector economy with optimal policy is first best arises when $g_1/\eta_1 = g_2/\eta_2$ or, equivalently, $a_1 g_1 = a_2 g_2$. In this case, Proposition 2 yields $\pi^\star = a_1 g_1$. Accordingly, this case generalizes Proposition 1, for the limit $\beta \rightarrow 1$, to a two-sector model with asymmetric price stickiness and sectoral productivity growth. Another special case in which the policy tradeoff disappears arises when firms in sector 2, say, have flexible prices. In this case, Proposition 2 yields $\pi^\star = g_1/\eta_1$.

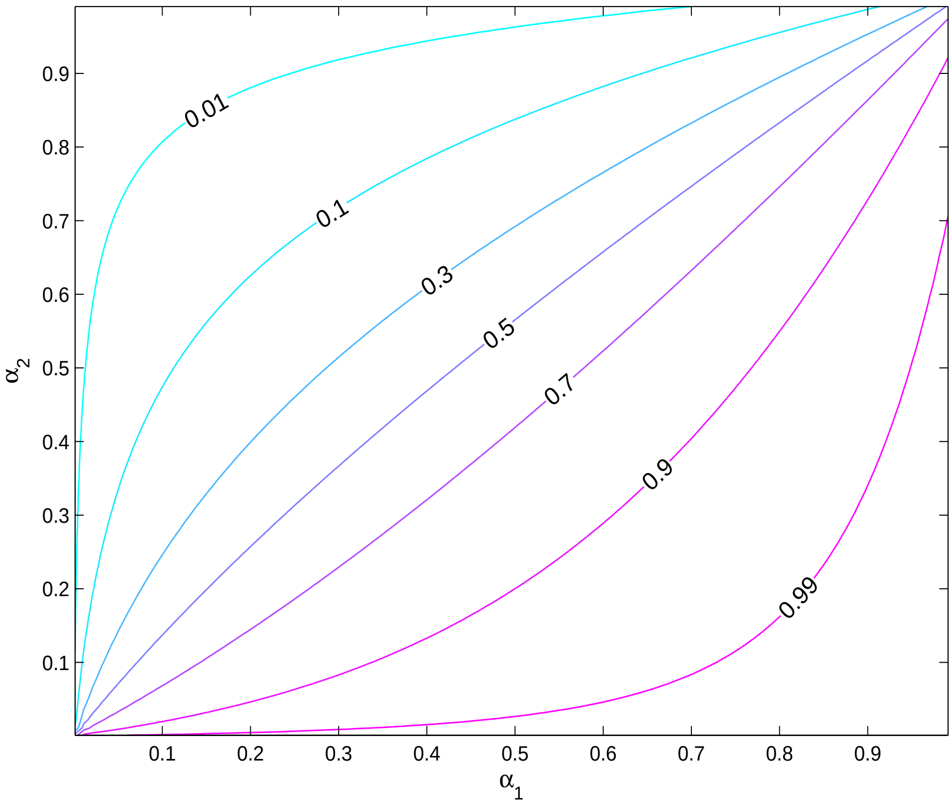

Figure 6: Weight $\omega(\pi^\star)$ as a function of the degree of price stickiness α_1 and α_2 in each sector. Lines indicate the combinations of α_1 and α_2 that yield a particular value of $\omega(\pi^\star)$. The long-run inflation rate is the one in Proposition 2, and the calibration is the one in Section 5.3.

The weight $\omega(\pi)$ in Proposition 2 depends on various sectoral asymmetries, but the predominant asymmetry is the amount of price stickiness in a sector z, which is shown in Figure 6. For a particular value of α_1, reducing the value of α_2 increases the weight on sector 1. The optimizing government weights the sector with the stickier prices more heavily, because thereby it shifts the price adjustment to the sector with the more flexible prices, where it is least distortive. This phenomenon is known as the "stickiness principle" in the literature on the optimal inflation *stabilization* policy (e.g., Aoki (2001), Mankiw and Reis (2003), Benigno (2004), and Eusepi, Hobijn, and Tambalotti (2011)). An important consequence of Proposition 2 is that in a two-sector model with firm-level productivity growth, this principle applies equally to the choice of the optimal *long-run* inflation rate.

The policy tradeoff between sectors in equation (19), which the optimizing government resolves using the stickiness principle, has two sources. One is the rate g_z at which incumbent firms become obsolete relative to new firms in a sector z. To illustrate this, I consider the case

$a_1 = a_2$, in which equation (19) reduces to

$$0 \;=\; \omega(\pi^\star)(\pi^\star - g_1)/g_1 + [1 - \omega(\pi^\star)](\pi^\star - g_2)/g_2 \;.$$

In this case, a natural interpretation of g_z is that it represents the long-run inflation rate that eliminates all distortions in a sector z. This is similar to Proposition 1, in which $\pi^\star = g$ eliminates all distortions in the one-sector model. In the two-sector model, however, no long-run inflation rate is optimal in both sectors at the same time, as long as $g_1 \neq g_2$, and this creates a policy tradeoff. In the case in which this is the only tradeoff, the optimal long-run inflation rate is negative and equal to -0.84% per year. This inflation rate is closer to g_1, which is equal to -0.89% per year, than to g_2, which is equal to -0.59% per year, because Services prices are stickier than Goods prices.

The second source of the policy tradeoff in equation (19) is the differential a_1/a_2 in sectoral productivity growth, incorporated into η_z. To illustrate this, I consider the case $g_1 = g_2 = 1$, in which this equation reduces to

$$0 \;=\; \omega(\pi^\star)(\eta_1\pi^\star - 1) + [1 - \omega(\pi^\star)](\eta_2\pi^\star - 1) \;.$$

Here, $\eta_z\pi$ is equal to the long-run inflation rate π_z in a sector z. Thus, in both sectors, the optimizing government targets π_z equal to zero percent per year. This is optimal because without (effective) firm-level productivity growth, firm-level marginal costs are constant, as in the basic New Keynesian model. In a model with two sectors and sectoral productivity growth, however, the relative price P_{1t}/P_{2t} is trending at rate a_2/a_1, because relative productivity gains in one sector reduce the relative price in this sector, and this triggers the sectoral inflation differential $\pi_1/\pi_2 = a_2/a_1$. Therefore, the optimizing government, which only controls π, cannot achieve $\pi_z = 1$ in both sectors at the same time, and this creates a policy tradeoff.

In the case in which this is the only policy tradeoff, the optimal long-run inflation rate is equal to -0.22% per year. This inflation rate implies that π_1, which is equal to 0.17% per year, is closer to zero than π_2, which is equal to -0.80% per year, because Services prices are stickier than Goods prices. The trending relative price P_{1t}/P_{2t} also constitutes the policy tradeoff analyzed in Wolman (2011). While he uses more general models of how firms set prices than the one used here, he also finds that mild deflation is the optimal policy.

I now return to the case in Proposition 2, which combines both sources of the policy tradeoff

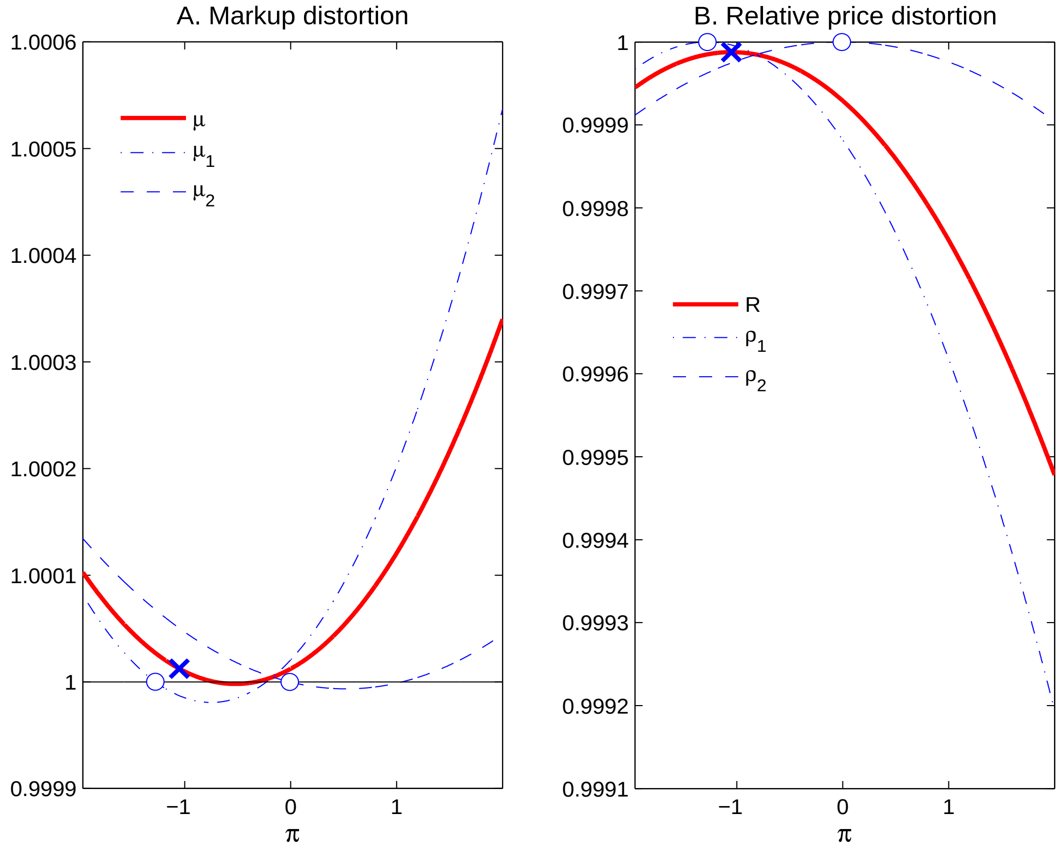

Figure 7: Aggregate distortions (bold lines) and sectoral distortions (thin lines) in the model without learning by doing and $\beta = 0.995$. Panel A shows the aggregate and sectoral markup distortions. Panel B shows the aggregate and sectoral relative price distortions. In both panels, π is the annualized net inflation rate.

between sectors, i.e., the various rates of obsolescence and the trending relative price. The optimal long-run inflation rate in Proposition 2 is equal to -1.05% per year and indicates that both sources of the policy tradeoff work into the same direction and amplify one another. Furthermore, $\pi^\star$ is closer to -1.28%, which minimizes the distortions in sector 1, than to -0.01%, which minimizes the distortions in sector 2, because Services prices are stickier than Goods prices. Figure 7 illustrates this by showing both, aggregate distortions (bold lines) and sectoral distortions (thin lines) in the model without learning by doing and $\beta = 0.995$. Using $\beta = 0.995$ instead of $\beta \to 1$ changes the optimal long-run inflation rate (ticked) and the inflation rates that minimize the sectoral distortions (circles), only marginally.

6.2 The two-sector model with learning by doing

With two sectors and learning by doing, the government's policy tradeoff between sectors coexists with its policy tradeoff between new and incumbent firms within a sector. In this case, I establish

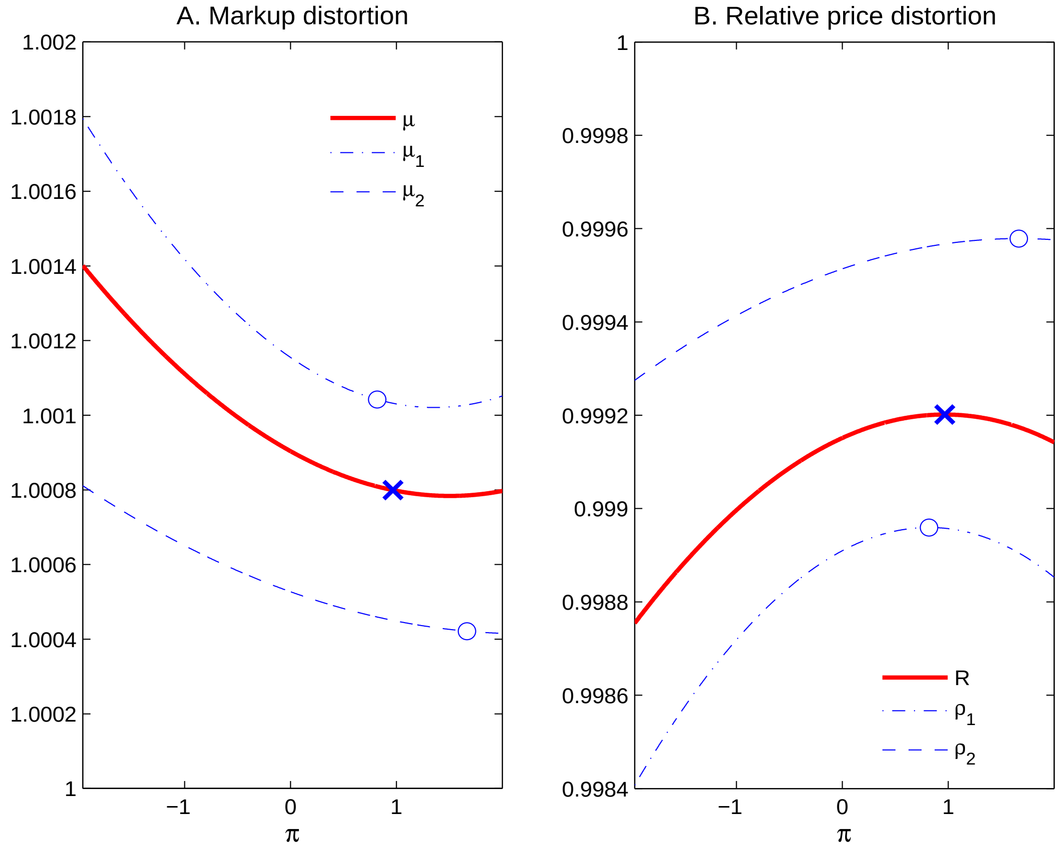

Figure 8: Aggregate distortions (bold lines) and sectoral distortions (thin lines) in the model with learning by doing and $\beta = 0.995$. Panel A shows the aggregate and sectoral markup distortions. Panel B shows the aggregate and sectoral relative price distortions. In both panels, π is the annualized net inflation rate.

a fourth result, namely that the optimal long-run inflation rate is positive and equal to 0.96% per year. Thus, incorporating learning by doing into the two-sector model increases the optimal long-run inflation rate by about two percentage points, i.e., from -1.05% to 0.96%. I compute the optimal long-run inflation rate numerically by maximizing steady-state welfare keeping $\tau_L^\star = -1/(\theta - 1)$ (which I vary below).

This result resembles the result in the one-sector model that when firms are learning by doing, the optimizing government increases its long-run inflation rate by a large amount in order to reduce the distortions that new firms experience from nominally sticky prices. Thus, despite incorporating sectoral asymmetries and, hence, additional policy tradeoffs, into the analysis, the optimizing government still resolves the policy tradeoff between new and incumbent firms in favor of new firms and, therefore, selects a positive long-run inflation rate of around one percent per year.

When it comes to resolving the policy tradeoff between sectors, the optimizing government continues to apply the stickiness principle, as in the model without learning by doing. Figure 8 illustrates this by showing aggregate distortions (bold lines) and sectoral distortions (thin lines) in the model with learning by doing. The optimal long-run inflation rate (ticked) is closer to 0.82% (circle), which minimizes the distortions in the Services sector, than 1.67% (circle), which minimizes the distortions in the Goods sector. This complies with the stickiness principle because Services prices are stickier than Goods prices and, therefore, a suboptimal inflation rate distorts the Goods sector less than the Services sector.

Comparing Figures 7 and 8 also illustrates that the aggregate distortions are considerably larger in the model with than without learning by doing, and the same message follows from comparing the aggregate distortions in terms of their welfare costs.[17] This suggests that the government's policy tradeoff between new and incumbent firms within a sector is considerably more severe than the government's policy tradeoff between sectors and, thus, should take a top priority in monetary policy analysis.

To explore the robustness of my quantitative results, I vary the absolute and relative amount of price stickiness in a sector. Increasing only the duration of Goods prices from two to three quarters increases $\pi^\star$ from 0.96% to 1.07%, because thereby $\pi^\star$ moves closer to the inflation rate that minimizes the distortions in the Goods sector. Furthermore, for corresponding reasons, increasing only the duration of Services prices from three to 4.5 quarters reduces $\pi^\star$ from 0.96% to 0.61%. Finally, increasing the duration of Goods prices from two to three quarters and Services prices from three to 4.5 quarters reduces $\pi^\star$ from 1.07% to 0.70%. Thus, while increasing the duration of prices in only one sector either increases or reduces $\pi^\star$, increasing the duration of prices in both sectors jointly reduces $\pi^\star$, as in the one-sector model (see Panel (e) in Figure 5).

The value of the labor income tax also affects the optimal long-run inflation rate. Namely, when τ_L is set to zero, then $\pi^\star$ increases from 0.96% to 1.08%. The government increases the optimal long-run inflation rate in order to also erode firms' markups to below the static markup because with a zero labor income tax, the static markup represents another distortion in the economy. This finding is in line with the finding in King and Wolman (1999) that the optimal long-run inflation rate is higher in a model with a static markup than in a model, in which the static markup is undone by another policy instrument.

[17] One measure of these costs is the consumption-equivalent reduction ϵ (see footnote 14). In the model with learning by doing and optimal policy, ϵ is equal to 0.0798%, whereas in the model without learning by doing and optimal policy, ϵ is equal to 0.0012%. Thus, the consumption-equivalent reduction is 65 times larger in the model with than without learning by doing.

7 Conclusion

Empirical data suggest that productivity in new firms grows faster than productivity in incumbent firms, as a result of learning by doing and embodied productivity growth. The purpose of this paper is to analyze the consequences of this firm-level productivity growth for macroeconomic policy choices and, particularly, for choosing the optimal long-run inflation rate. My analysis incorporates firm-level productivity growth into a stylized monetary model with sticky prices that admits the heterogenous firms to be aggregated analytically.

My baseline result is that firm-level productivity growth justifies an optimizing government in targeting a positive long-run inflation rate of between 0.5% and 1.5% per year. This inflation rate helps the fast-growing new firms to align their real price with their productivity growth. The baseline result is robust with respect to changing parameter calibrations, and it is also robust with respect to incorporating sectoral asymmetries and, hence, additional policy tradeoffs.

A key difference between the sticky-price model used here and the basic New Keynesian model, which predicts an optimal long-run inflation rate near zero, is the behavior of real marginal costs at the firm level. While they remain constant in the steady state of the basic New Keynesian model, they decline over the lifetime of a firm in the steady state of the model used here. The generic conclusion that can be derived from this difference is that sticky nominal prices alone do not constitute a compelling reason for the government to target a zero long-run inflation rate.

In the wake of the recent financial turmoil, Blanchard, Dell'Ariccia, and Mauro (2010), among others, discuss some of the consequences of raising inflation targets to above their current levels in order to provide central banks with more leeway to cope with large adverse shocks. My results contribute to this discussion by demonstrating that the welfare costs caused by a moderately positive long-run inflation rate derived in, e.g., Coibion, Gorodnichenko, and Wieland (2012), represent conservative estimates if one also accounts for firm-level productivity growth, as I have done here.

There are at least two interesting ways to extend my analysis in future work. First, while my analysis finds a positive optimal long-run inflation rate in a cashless economy, the literature emphasizes that the costs arising from holding money imply a negative optimal long-run inflation rate, and future work could incorporate these costs. Second, in line with the evidence, my analysis emphasizes a supply-side factor, i.e., productivity growth, and shows how it affects firm-level marginal costs and, thereby, the optimal long-run inflation rate. Yet, Foster, Haltiwanger, and Syverson (2008) have recently suggested that demand-side factors are another important distinction across firms, and future work could also analyse these factors.

References

Amano, R., K. Moran, S. Murchison, and A. Rennison (2009): "Trend inflation, wage and price rigidities, and productivity growth," *Journal of Monetary Economics*, 56(3), 353–364.

Aoki, K. (2001): "Optimal monetary policy responses to relative-price changes," *Journal of Monetary Economics*, 48(1), 55–80.

Aruoba, S. B., and F. Schorfheide (2011): "Sticky Prices versus Monetary Frictions: An Estimation of Policy Trade-Offs," *American Economic Journal: Macroeconomics*, 3(1), 60–90.

Bahk, B.-H., and M. Gort (1993): "Decomposing Learning by Doing in New Plants," *Journal of Political Economy*, 101(4), 561–83.

Basu, S., J. G. Fernald, and M. S. Kimball (2006): "Are Technology Improvements Contractionary?," *American Economic Review*, 96(5), 1418–1448.

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

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

Bilbiie, F., F. Ghironi, and M. 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.

Bilbiie, F. O., I. Fujiwara, and F. Ghironi (2011): "Optimal Monetary Policy with Endogenous Entry and Product Variety," NBER Working Papers 17489, National Bureau of Economic Research, Inc.

Bilbiie, F. O., F. Ghironi, and M. J. Melitz (2012): "Endogenous Entry, Product Variety, and Business Cycles," *Journal of Political Economy*, 120(2), 304 – 345.

Billi, R. M. (2011): "Optimal Inflation for the U.S. Economy," *American Economic Journal: Macroeconomics*, 3(3), 29–52.

Bils, M., and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112(5), 947–985.

Blanchard, O., G. Dell'Ariccia, and P. Mauro (2010): "Rethinking Macroeconomic Policy," *Journal of Money, Credit and Banking*, 42(s1), 199–215.

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

Cooper, R., and A. Johri (2002): "Learning-by-doing and aggregate fluctuations," *Journal of Monetary Economics*, 49(8), 1539–1566.

Dotsey, M., R. G. King, and A. L. Wolman (1999): "State-Dependent Pricing And The General Equilibrium Dynamics Of Money And Output," *The Quarterly Journal of Economics*, 114(2), 655–690.

Eusepi, S., B. Hobijn, and A. Tambalotti (2011): "CONDI: A Cost-of-Nominal-Distortions Index," *American Economic Journal: Macroeconomics*, 3(3), 53–91.

Faia, E. (2012): "Oligopolistic competition and optimal monetary policy," *Journal of Economic Dynamics and Control*, 36(11), 1760–1774.

Foster, L., J. Haltiwanger, and C. J. Krizan (2006): "Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s," *The Review of Economics and Statistics*, 88(4), 748–758.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 98(1), 394–425.

Foster, L., J. C. Haltiwanger, and C. Syverson (2012): "The Slow Growth of New Plants: Learning about Demand?," NBER Working Papers 17853, National Bureau of Economic Research, Inc.

Friedman, M. (1969): *The Optimum Quantity of Money*. Macmillan.

Hornstein, A., and P. Krusell (1996): "Can Technology Improvements Cause Productivity Slowdowns?," in *NBER Macroeconomics Annual 1996, Volume 11*, NBER Chapters, pp. 209–276. National Bureau of Economic Research, Inc.

Janiak, A., and P. S. Monteiro (2011): "Inflation and Welfare in Long-Run Equilibrium with Firm Dynamics," *Journal of Money, Credit and Banking*, 43(5), 795–834.

Jensen, J. B., R. H. McGuckin, and K. J. Stiroh (2001): "The Impact Of Vintage And Survival On Productivity: Evidence From Cohorts Of U.S. Manufacturing Plants," *The Review of Economics and Statistics*, 83(2), 323–332.

Jovanovic, B., and Y. Nyarko (1995): "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," *Brookings Papers on Economic Activity: Microeconomics Issue (1)*, pp. 247–299.

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.

King, R., and A. L. Wolman (1999): "What Should the Monetary Authority Do When Prices Are Sticky?," in *Monetary Policy Rules*, NBER Chapters, pp. 349–404. National Bureau of Economic Research, Inc.

Levitt, S. D., J. A. List, and C. Syverson (2012): "Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant," NBER Working Papers 18017, National Bureau of Economic Research, Inc.

Lewis, V., and C. Poilly (2012): "Firm entry, markups and the monetary transmission mechanism," *Journal of Monetary Economics*, 59(7), 670–685.

Mankiw, N. G., and R. Reis (2003): "What Measure of Inflation Should a Central Bank Target?," *Journal of the European Economic Association*, 1(5), 1058–1086.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

Miranda, J., S. Klimek, and R. Jarmin (2004): "Firm Entry and Exit in the U.S. Retail Sector, 1977-1997," Working Papers 04-17, Center for Economic Studies, U.S. Census Bureau.

Rogers, M., C. Helmers, and C. Koch (2010): "Firm growth and firm size," *Applied Economics Letters*, 17(16), 1547–1550.

Sakellaris, P., and D. J. Wilson (2004): "Quantifying Embodied Technological Change," *Review of Economic Dynamics*, 7(1), 1–26.

Schmitt-Groh, 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.

Schmitt-Grohe, S., and M. Uribe (2012): "On quality bias and inflation targets," *Journal of Monetary Economics*, 59(4), 393–400.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97(3), 586–606.

Suarez, C., and J. De Jorge (2012): "Growth and evolution of technical efficiency among entrant versus incumbent firms: the Spanish retail sector," *The Service Industries Journal*, 32(14), 2265–2285.

Wolman, A. L. (2011): "The Optimal Rate of Inflation with Trending Relative Prices," *Journal of Money, Credit and Banking*, 43, 355–384.

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

Yorukoglu, M. (1998): "The Information Technology Productivity Paradox," *Review of Economic Dynamics*, 1(2), 551–592.

Yun, T. (2005): "Optimal Monetary Policy with Relative Price Distortions," *American Economic Review*, 95(1), 89–109.