Markups and Marginal Cost over the Firm Life*

Klaus Adam† Tobias Renkin‡ Gabriel Züllig§

June 23, 2025

Abstract

We estimate the dynamics of markups, marginal cost and prices over the life-cycle of Danish manufacturing firms. Markups increase by 8 percentage points over the first 20 years of firms' life. This reflects a substantial decrease of marginal cost which is only partially passed on into prices. These age patterns coincide with product turnover—both introductions and discontinuations—among young firms. We discuss the implications of this finding for a number of macroeconomic theories and trends. For example, we empirically test the hypothesis that declining business dynamism—a reduction in new entrants and thus an increase in the average firm age—led to an increase in the average markup, but find no substantial contribution.

JEL classification: E52, L11, L13

Keywords: markups, firm dynamism, life cycle, business dynamics, inflation

*We are grateful to Timo Boppart, Ryan Kim, Vivien Lewis, Isabelle Mejean, Marta Morazzoni, Filip Rozsypal and Julian Schärer, as well as seminar participants at the 2nd Workshop on Market Power and the Economy at the University of Zurich, the 2nd Workshop on Firm Heterogeneity and Macroeconomics at the University of Bonn, the Kent-BoE-Bristol workshop on Firm Dynamics, Market Structures and Productivity in Macroeconomics at the University of Kent and a Danmarks Nationalbank brownbag seminar for helpful comments and suggestions. Klaus Adam gratefully acknowledges support from the Deutsche Forschungsgemeinschaft through CRC-TR 224 (project C02) and Gabriel Züllig from the Swiss National Science Foundation (grant P2SKP1 194963). The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the authors. They do not necessarily reflect the views of Danmarks Nationalbank or the Swiss National Bank (SNB). All errors and omissions are our own.

†University of Mannheim and CEPR, adam@uni-mannheim.de

‡Danmarks Nationalbank, tobias.renkin@gmail.com

§Swiss National Bank, gabriel.zuellig@snb.ch

1 Introduction

In many developed economies, firms' markups have been rising since the 1980s. Over the same time period, business dynamism, measured by the rate of firm entry and exit, has declined and as a result, the age of the average firm has increased.[1] There are several channels through which a decline in business dynamism can lead to an increase in average markups. Many common pricing models would predict a direct relationship—that younger firms charge lower markups than older firms. This includes models in which firms build up a customer base, or models with strategic complementarities in price-setting (if young firms produce with higher marginal costs than older firms). We focus on this direct link in this paper.[2]

We estimate life-cycle profiles of markups, prices and marginal costs to study the relationship between markups and firm aging in Danish manufacturing industries. We find that markups increase by about 8pp over the first 20 years of an average firm's life. This is a large increase relative to an average markup of about 12% in our sample. Markups increase due to a decline in marginal cost as firms age that is not fully passed-through to prices. In particular, we find that average marginal costs fall by 14% over the first 20 years of a firm's life, while average prices fall only by 7%. Beyond age 20, we find a slight decrease of markups as firms become older.

Over the two decades from 2001 to 2022, the average Danish manufacturing markups has increased by about 10%, while the average firm age has increased from 24 to 41 years. However, despite the clear age pattern in markups, the effect of firm aging on the average markup is small and does not contribute importantly to the observed time trend. This is due to two reasons. First, the aggregate importance of young firms is limited. Second, the small negative effect of age on markups of older firms offsets the larger positive effect of of aging on the markups of younger firms. We find that most of the increase in the average markup is accounted for by time effects that affect all firms equally, rather than age or cohort effects.

We also present evidence on some drivers of the age patterns we observe. We use our production function estimates to show that about 60% of the decrease in average marginal cost we observe over the first 20 years of firms' life is driven by an increase in firms' total factor productivity.

[1] De Loecker et al. (2020) document an increase in markups in the U.S. De Loecker and Eeckhout (2018) provide evidence for increasing markups in numerous countries, including Denmark. De Ridder et al. (2024) provide evidence for French firms addressing several critiques of the production approach to markup measurement. A decline in business dynamism and the entry of new firms in the U.S. has been documented by Decker et al. (2014), Decker et al. (2016) and Pugsley and Șahin (2019). Calvino et al. (2018), Calvino et al. (2020) and Biondi et al. (2025) document declines entry rates and the employment share of entrants in several European countries, including Denmark.

[2] In addition, a decrease in business dynamism could lead to an increase in average markups indirectly if decreased competitive pressure from (potential) entrants entices incumbent firms to increase their markups.

The remainder can be attributed to a combination of scale effects and more efficient input bundles. Moreover, we show that product turnover could contribute to age patterns in markups and marginal cost. Young firms introduce and discontinue products more frequently than older firms, and we show that both introduction of new and discontinuation of existing products is associated with an increase in firm-level markups in our data.

We estimate firms' markups following the production approach of De Loecker and Warzynski (2012). We use firms' physical output (rather than sales) to estimate the output elasticities required for this approach. Our markup estimates are thus not subject to the identification issues raised in Bond et al. (2021) and De Ridder et al. (2024). We use administrative micro-data on firms' balance sheets that contain detailed information on nominal firm-level sales and input expenditures. In addition, we use micro-data underlying the Danish Producer Price Index (PPI) which allows us to construct firm-level price indices and correctly deflate sales to output.

We decompose firm-level markup developments into firm effects (which nest cohort effects), year effects and the age patterns that we are primarily interested in. Since the linear components of those three effects are collinear, we impose two alternative structural restrictions to recover age effects. Both restrictions are commonly used in life-cycle analyses. Our baseline restriction follows Deaton and Paxson (1994) and imposes that firm effects are orthogonal to any linear cohort trend. Alternatively, we impose the restriction that age effects (which we estimate as a polynomial) have a stationary point at a chosen reference age. This assumption has been used by Card et al. (2013, 2016). Most of our results are very similar using either restriction, which we interpret as an informal over-identification test of the restrictions.

Our main contribution is to document age patterns in markups and marginal costs of manufacturing firms. Such patterns are an important input for models of firm dynamics with imperfect competition. To the best of our knowledge, we are the first to show that firms increase their markups over the first years of their life. A larger literature has documented productivity differences between entrants and incumbent firms (Foster et al., 2001, Melitz and Polanec, 2015). This literature usually uses TFP estimates based on revenue data, which conflates differences in physical productivity and markups. The exception is Foster et al. (2008), who use data on physical output and prices of U.S. manufacturing industries from 1977 to 1997 to document the physical TFP and prices of entrants and young firms relative to incumbents. Our approach differs in several aspects from the analysis of Foster et al. First, Foster et al. measure TFP using a Cobb-Douglas production function that is calibrated using cost shares and assumes constant returns to scale, while we estimate a flexible translog production function. Second, Foster et al. study the life cycle of manufacturing plants, while we study the life cycle of manufacturing *firms*. They find that entering plants are more productive and charge similar prices as incumbent plants,[3] suggesting that their markups are higher. In contrast, we find that younger firms are less productive and charge higher prices than older firms. Our results are consistent with van Vlokhoven (2021), who shows that economic profits of public U.S. firms increase with age. Our data covers both public and private firms and thus avoids sample selection issues related to the use of Compustat or similar data, which only covers publicly listed firms. Moreover, we estimate (unit) markups using the widely used production approach of De Loecker and Warzynski (2012), while van Vlokhoven employs an alternative methodology that estimates total profits but doesn't recover markups.

We show that the increasing age profile in markups is driven by a decreasing age profile in marginal costs that is not completely passed-through to prices. This suggests that young firms are less productive than older firms. Given this, the patterns in prices and markups can be explained by strategic complementarity with prices of incumbents that limits the pass-through of variation in cost. Such a setting has been used in Chiavari et al. (2023) or van Vlokhoven (2021). Our findings are not consistent with customer-base models such as in Ignaszak and Sedláček (2023), where an increasing age pattern in markups would be driven by increasing age patterns in prices. The age patterns in markups are also consistent with cross-sectional evidence in Conlon et al. (2023), who show that across U.S. industries, markup growth over the 1980–2019 period is not strongly correlated with price increases and likely driven by productivity growth that is not passed into prices.

Our second contribution is to relate the estimated age pattern to the time trend of average markups (see for example De Loecker et al. (2020) for the US or De Loecker and Eeckhout (2018) for a global comparison). We show that even though young firms charge lower markups than older firms and the firm distribution is shifting toward older firms, the direct effect of aging on the average markup is small. This is driven by two factors. First, there are relatively few young firms in our dataset overall, limiting the impact of their behavior on the aggregate. Second, for firms above age 25 markups slightly decrease with age. As a result, the effect of the aging of young firms is partially offset by a counter-acting effect from the aging of older firms. Overall, our decomposition attributes the bulk of the 2001–2021 increase in Danish markups on time effects that affect all firms equally.

The remainder of the paper is structured as follows. Section 2 introduces the data used to estimate markups and marginal cost, the estimation approach itself is outlined in Section 3. Section 4

[3]We refer to the unweighted regressions in their paper.

identifies age patterns in markups and marginal cost, with Section 5 establishing the robustness of these results. In Section 6 we explore the sources of these age patterns in more detail. Ultimately, we discuss the implications of our results for a range of macroeconomic questions and conclude in Section 7.

2 Data

We combine several firm-level datasets provided by Statistics Denmark to estimate markups at the firm level. We use price data from the Producer Price Index (PPI) survey, accounting data from the accounting statistics (FIRE) and firm demographics from the business register (FIRM).

Price data Our analysis draws on microdata from the Danish Producer Price Index (PPI), which provides detailed price information for manufacturing firms from 2001 to 2022. The dataset comprises approximately 3,000 monthly price quotes from about 600 firms, with products classified using 8-digit Harmonized System (HS) codes. Each firm reports transaction prices in Danish kroner for a stable set of their most representative products. The dataset includes both domestic and export prices. We pool them in our dataset, but consider domestically sold and exported ones as separate products, even if they share the same HS code. One advantage relative to unit value data used for example in De Ridder et al. (2024), is that the survey is designed to allow adjustments for quality changes and product substitutions. In the case of changes to the product, firms report both the price for the updated product as well as a hypothetical price for the same product in the previous period. Another advantage relative to unit value data is that the dataset is strongly balanced, with very few gaps in the price series.

We use this dataset to construct annual firm-specific price deflators which we then combine with sales data to measure real firm output. The deflators are computed as the average of quality-adjusted price changes across all products reported by each firm.[4] On average, the firm-level output price deflators computed from the micro data are consistent with readily available industry-level producer price indices. Details on the construction of firm- and industry-level price indices and their correlations are provided in Appendix A.1.

Accounting data We combine firm-level price deflators with annual data on sales, purchases and inventories from the Danish accounting statistics to construct measures of inputs and output.

[4]The PPI data does not contain weights that indicate the importance of a product within a firm.

The accounting statistics collect headline balance sheet items such as sales and assets from tax data and more detailed items such as intermediate purchases and different kinds of inventories in a large-scale survey of firms. The survey population excludes firms with less than five employees and firms in some sectors such as agriculture and finance, but includes all firms with more than 50 employees. Firms with 20 to 49 employees are included for five years every ten years, firms with 10 to 19 employees for two years every ten years, and firms with 5 to 10 employees are included every 10th year.

We calculate firms' output as annual sales plus the change in the firms' inventory of final goods, deflated with firm-specific deflators constructed from PPI prices. Firms' material input is computed as annual expenditures for intermediates minus the change in intermediate inventories deflated with sector-specific input price indices from the Danish national accounts.[5] Labor input is calculated from the firms' total wage bill deflated with an aggregate wage deflator. We prefer a measure of real labor cost rather than full-time equivalents in the production function, because it provides a quality-adjusted measure of labor input, but our results are robust to using full-time equivalents instead. We measure capital used in production as fixed assets, including both material and immaterial fixed assets, but excluding financial assets.

Firm demographics We use information on firm demographics from the Danish business register. This register contains the date of registration for the universe of firms. We calculate firm age starting with age one in the year of registration. The business register also contains a sector code for each firm, and we will define a sector as a 2-digit NACE code throughout the paper.

Final sample and characteristics Our final sample contains 800 manufacturing firms that we can match between the PPI and accounting statistics. Firms in the sample operate in 15 different 2-digit NACE sectors.[6] Table 1 contains the most important summary statistics for the panel of firms. The median firm in our sample has 100 employees, is observed for 11 (out of 22) years and did not grow over the sample period. The firms in the sample are highly heterogeneous in terms of age. The median firm is 22 years old when we first observe it. As is crucial for our analysis, there is a substantial share of young firms in the sample. For example, we observe around 10% of all firms at least once before they reach age 5 and around 20% of firms before age 10. On the other hand, we observe several firms that are well above 50 years old. The full distribution of age

[5] We refer to Appendix A.1 for details on all factor price deflators.

[6] Because we estimate production function coefficients separately by industry, we only consider sectors with at least 20 firms. The sector distribution is shown in Figure A.2 in the appendix.

Table 1: Firm characteristics

	Mean	p5	p25	p50	p75	p95
Employment (firm average)	217.5	25.7	55.4	100.9	206.2	628.4
Employment growth (first to last obs., %)	9.6	-72.0	-37.0	-8.6	18.0	124.6
Number of years observed	12.2	2	6	11	19	22
Age at first observation	24.7	2	12	22	33	61
Material input share (firm average, %)	51.1	28.0	41.7	50.4	60.3	76.2
Labor input share (firm average, %)	24.8	9.0	18.3	24.5	30.7	39.9

Notes: The panel data is at the annual frequency covering the years 2001–22 with a total of 9,700 observations, but the table shows characteristics at the firm level, with a sample size of 798. Factor input shares are relative to nominal sales in the given year, corrected for the change in inventories of final goods.

across firms is shown in Figure A.3 in the appendix.

On average, the final sample accounts for 39% of total manufacturing employment in Denmark. In the beginning of the sample period, this share is lower at 29%, increases in subsequent years and plateaus at over 40% from 2013 onward. Firms in the sample are thus responsible for an important share of overall manufacturing employment (and output).[7]

However, our sample is not necessarily representative for the typical manufacturing firm, as we require firms to be sampled simultaneously in the PPI and accounting statistics surveys. We define a benchmark sample of firms which consists of all manufacturing firms in the firm register that at one point over the sample period reach at least 50 employees, but is not subject to the constraint of being included in the PPI or accounting statistics. This benchmark population contains 2,000 firms. In Figure A.4 in the appendix, we compare the 800 firms in our sample to the benchmark sample. As one would expect with a size-dependent likelihood to be covered in either the PPI or accounting statistics surveys, our sample is skewed towards larger and older firms. In robustness checks we re-weight firms in our sample to be representative and show that this does not affect our results.

A notable feature of the data illustrated in Figure 1 is that the average age of firms increases strongly over the 22 years covered in the data. The average firm age in the sample is 24 years in 2001 and increases by 17 years to 41 years in 2022 (blue solid line). A similar trend is visible in the

[7] We document that the Danish manufacturing sector is comparable to that in other European countries in two dimensions. First, the firm size distribution is similar to that in other European countries, as we document in Table A.2 using data from Kalemli-Ozcan et al. (2022). Second, Renkin and Züllig (2024) show that Danish producers sell product categories that have a similar demand structure compared to other EU producers. As is standard in the literature, we do not estimate markups for service firms.

median age, which increases by 18 years over the same period. This increase in the average age is representative for manufacturing firms in Denmark. The dashed lines in Figure 1 show the mean and median age in the benchmark sample described above. While firms in the benchmark sample are younger, they age at a similar pace to firms in our merged estimation sample. We show in the Appendix (Figure A.5) that this aging is a result of a decreasing share of startup firms, and that the decline in business dynamism documented in the literature for other countries also occurs in Denmark.

Figure 1: Firm age trends

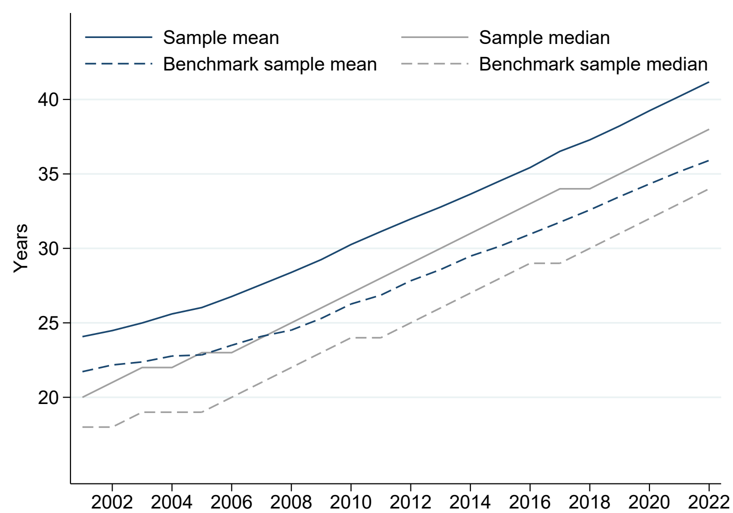

Notes: The sample consists of 800 manufacturing firms for which we have full price and production data, i.e., for which we are able to estimate markups. The benchmark sample is a broader set of approximately 2,000 manufacturing firms.

3 Estimating firms' markups and marginal cost

We estimate firm markups following the production function approach of De Loecker and Warzynski (2012). Bond et al. (2021) and De Ridder et al. (2024) have recently highlighted the limitations of markup estimation when production functions are estimated using sales. Our estimates are not subject to these limitations, since we use firm-specific prices to construct measures of quantities instead of sales. Moreover, the combination of price data with our markup estimates allows us to recover marginal cost, and to discuss joint dynamics in markups, prices and marginal cost as firms age.

3.1 Estimation approach

The production function approach relies on cost minimization and the assumption of competitive input markets to identify markups, but does not require assumptions on the structure of output markets or demand. Following standard practice, we assume that ex-post observed output Y_{it} of firm i at time t is the combination of planned output Y^*_{it} and a disturbance ϵ_{it} that is realized after production decisions are made:[8]

$$Y_{it} = Y^*_{it}(K_{it}, L_{it}, M_{it}, \Omega_{it}) \exp(\epsilon_{it}) \tag{1}$$

Planned output Y^*_{it} is unobserved by the researcher and is a function of capital K_{it}, labor L_{it}, intermediate material inputs M_{it} and total factor productivity Ω_{it}. We treat materials input M_{it} as a flexible input that can contemporaneously adjust to the current level of productivity Ω_{it}.[9] For a given planned output level Y^*_{it}, the firm chooses material input (and perhaps other flexible inputs) in period t to produce at minimal cost, taking factor prices as given. The first-order condition from this minimization problem together with the definition of the markup as the ratio of the output price over marginal cost ($\mu_{it} = \frac{P_{it}}{\lambda_{it}}$) yields the ratio estimator for the markup:[10]

$$\mu_{it} = \frac{\theta^M_{it}}{\alpha^M_{it}} \exp(-\epsilon_{it}). \tag{2}$$

$\theta^M_{it} = \frac{\partial Y^*_{it}}{\partial M_{it}} \frac{M_{it}}{Y^*_{it}}$ denotes the output elasticity with respect to the material input, $\alpha^M_{it} \equiv \frac{P_{it} Y_{it}}{P^M_t M_{it}}$ the share of expenditures for materials in total revenue, P_{it} the firm's output price, and P^M_t the input price for materials.[11]

Given an estimate for the output elasticity $\hat{\theta}^M_{it}$ and the surprise component $\hat{\epsilon}_{it}$, the estimate for the markup $\hat{\mu}_{it}$ and marginal cost $\hat{\lambda}_{it}$ are given by:

$$\hat{\mu}_{it} = \frac{\hat{\theta}^M_{it}}{\alpha^M_{it} \exp(\hat{\epsilon}_{it})}, \qquad \hat{\lambda}_{it} = \frac{1}{\hat{\mu}_{it}} P_{it}, \tag{3}$$

where α^M_{it} and P_{it} can be directly measured in the data. The second equation in (3) shows how

[8] The disturbance may simply capture measurement error in output.

[9] Other inputs may also be flexible, e.g., labor input, but this is not required for the approach to work. Also, for materials markets the assumption of competitive input markets appears most plausible.

[10] Notice that λ_{it} is also the Lagrange multiplier of the cost minimization problem.

[11] Because the researcher does not observe the expenditure share in terms of revenue anticipated by the firm at the point of the decision but only realized revenues in hindsight, the revenue share has to be corrected for $\exp(\epsilon_{it})$.

we obtain a measure of marginal cost, given the firm-level markup estimate and the observed output price. We now describe how we obtain the estimates $\hat{\theta}_{it}^M$ and $\hat{\epsilon}_{it}$.

Production function estimation We estimate the output elasticity of material inputs using a flexible translog production function (in logarithms):

$$\begin{aligned} y_{it} = {} & \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \beta_{kk} k_{it}^2 + \beta_{ll} l_{it}^2 + \beta_{mm} m_{it}^2 \\ & + \beta_{kl} k_{it} l_{it} + \beta_{km} k_{it} m_{it} + \beta_{lm} l_{it} m_{it} + \beta_{klm} k_{it} l_{it} m_{it} + \omega_{it} + \epsilon_{it}. \end{aligned} \tag{4}$$

All variables in the estimation are real quantities. Output y_{it} is based on revenues deflated with firm-level output price deflators as described in Section 2. The variables on the r.h.s. of Equation (4), i.e. the inputs capital, labor and materials, are deflated as described in Appendix A. ω_{it} is the log of Hicks-neutral productivity. We estimate separate production functions for 16 different 2-digit NACE industries, but suppress industry-level subscripts to keep the notation simple.

Given estimates of the production function coefficients, the output elasticity of material input is:

$$\hat{\theta}_{it}^M = \hat{\beta}_m + \hat{\beta}_{km} k_{it} + \hat{\beta}_{lm} l_{it} + 2\hat{\beta}_{mm} m_{it} + \hat{\beta}_{klm} k_{it} l_{it}. \tag{5}$$

Equation (5) shows that output elasticities vary across firms and over time due to the interaction terms of the translog production function .[12]

We estimate the production function in (4) in two steps following the literature (see e.g. Olley and Pakes (1996), Levinsohn and Petrin (2003), Ackerberg et al. (2015)). The first step removes ϵ_{it} from observed output.[13] The second step addresses the issue of endogeneity. This is necessary because realizations of total factor productivity ω_{it} affect both the firm's choice of material input m_{it} as well as the resulting output y_{it}. Following standard practice, we address this issue using a GMM estimator (Blundell and Bond, 2000, De Loecker, 2011).

Purging ϵ_{it} in the first step The first step uses the fact that the firm observes ω_{it} but not ϵ_{it} when it makes production decisions. If the optimal choice of m_{it} is a monotonic and invertible function in ω_{it} and potentially other information stacked in the vector Ξ_{it}, one can replace $\omega_{it} = m^{-1}(m_{it}, \Xi_{it})$ in equation (4) and thereby consistently estimate ϵ_{it}. Importantly, ϵ_{it} is not a part

[12] With cross-terms equal to zero, i.e. in the Cobb-Douglas case, we would have $\beta_{km} = \beta_{lm} = \beta_{mm} = \beta_{klm} = 0$, and the output elasticity of material inputs would be equal to the constant $\hat{\beta}_{km}$ and equal for all firms.

[13] This is necessary because without additional restrictions, productivity ω_{it} and the shock ϵ_{it} cannot be separately identified.

of Ξ_{it}. The variables that are included among the first-step regressors Ξ_{it} are, as in De Loecker (2011), third-order polynomials of all production function inputs and their interactions as well as time fixed effects to absorb common variation across periods, e.g. factor input trends.

The first-order condition of material inputs in the cost minimization problem contains marginal cost and hence this or prices and the markup itself should be part of Ξ_{it} if they are heterogeneous across firms, as in the case of imperfect competition (Doraszelski and Jaumandreu, 2021, De Ridder et al., 2024). To control for the markup, we include two dimensions along which we conjecture that both markups and prices are heterogeneous. First, we include a firm fixed effect to absorb variety-specific demand that is not accounted for in the data. Second, we include a polynomial in age, since we conjecture that the firm life cycle is an important determinant of markups.[14] The second step uses the fitted values from the first step as the dependent variable, i.e. estimates the elasticity of y^*_{it} with respect to m_{it}.

Addressing endogeneity of ω_{it} and m_{it} in the second step We follow the standard approach in the literature by instrumenting m_{it} with with lagged variable input $m_{i,t-1}$. This is a relevant and valid instrument since values of $m_{i,t-1}$ are predictive of m_{it} but orthogonal to ω_{it}. To ensure that $\mathbb{E}[m_{i,t-1}, m_{it}] \neq 0$, i.e. the instrument is relevant, one can assume that ω_{it} is persistent, i.e. $\omega_{it} = \rho\omega_{i,t-1} + \xi_{it}$.[15]

Crucially for the identification of the production function parameters, our baseline specification uses firm-level deflators from PPI prices, denoted P_{it}, to deflate firm-level sales. De Ridder et al. (2024) show that the moment condition $\mathbb{E}(\hat{\xi}_{it}(\beta)m_{i,t-1}) = 0$ only correctly identifies β if y^*_{it} measures real output. In many empirical settings, however, it is difficult to estimate physical output elasticities because only nominal revenues are observed. For example, De Loecker (2011) deflate nominal revenues with industry-level deflators that are common across firms. Because of the inclusion of time fixed effects in the first step, this essentially estimates the revenue elasticity. In that case, the estimated production function coefficients (and by extension $\hat{\theta}^M_{it}$ and $\hat{\mu}_{it}$) will contain information on the true markup, but be biased by the covariance between the lagged

[14]De Ridder et al. (2024) include prices, or in their case unit values, and market shares instead in the first step instead. They find estimated log markups with their extended controls in the first step have a Pearson correlation of 0.62 with markups estimated without these additional controls for the markup. In our case, it is 0.65. Notice that due the nonparametric nature of the first step estimation, we are not required to know the parameters governing the structural relationship between our additional control and the markup.

[15]This results in the moment conditions $\mathbb{E}(\hat{\xi}_{it}(\beta)X) = 0$, where X is $(k_{it}, l_{it}, m_{i,t-1}, k^2_{it}, l^2_{it}, m^2_{i,t-1}, k_{it}l_{it}, k_{it}m_{i,t-1}, l_{it}m_{i,t-1}, k_{it}l_{it}m_{i,t-1})$. Starting with a set of values for the parameters β and ρ, we compute the implied $\hat{\omega}_{it}$, $\hat{\omega}_{i,t-1}$ and $\hat{\xi}_{it}$ and iterate until said moment conditions hold. The output elasticities implied by the optimized $\hat{\beta}$'s enters in the numerator of the ratio estimator of the markup.

input variable with output prices. The bias might be both positive or negative and wash out in the aggregate. Since we are interested in firm-level markups, however, it is crucial that we measure physical output. We do so by dividing nominal sales by a firm-specific output price deflator based on the PPI. These output prices are observed directly, and thus more precisely measured than unit values used in other datasets. We contrast our results to a version where we estimate the output elasticity using industry-level deflators P_{kt}, i.e., effectively estimate revenue elasticities. In line with De Ridder et al. (2024), we find differences in the level of estimated output and revenue elasticities, but similar aggregate time trends.

3.2 Results of markup estimation

It is useful to discuss some outcomes of our production function estimation and compare the resulting markup estimates to the existing literature. Table B.1 in the appendix presents average material input shares by sector as well as average output elasticities implied by the translog production function coefficients. The mean input share for material inputs over the full sample is 0.51 and the average output elasticity for material inputs is 0.58. The latter compares to an average value of 0.59 for French manufacturing firms in De Ridder et al. (2024) and 0.57 for US firms in Foster et al. (2024). The average estimated output elasticities for labor and capital inputs amount to 0.40 and 0.05, respectively. Following Equation (3), the output elasticity of the flexible input good – material – should be put in relation to its nominal input share to get an estimate of the markup. The resulting average markup over the sample period is 11%.[16]

This is on the lower end of the large range of estimates found in the literature. We attribute this to the fact that our sample is based on both public and private firms as opposed to studies using Compustat data only and that it only covers manufacturing firms, which tend to have lower markups (Marto, 2024). De Ridder et al. (2024) report a baseline log markup estimate of 32% among (larger) manufacturing firms in France, which reduces to 11% when estimated with revenue elasticities. As our output elasticity estimates are in line with the literature, they difference in the reported average markups is due to somewhat higher material input shares among Danish manufacturing firms.

In Figure 2, we show distributions of our markup sample in time and in the cross-section. Panel (a) shows the development of the aggregate (log) markup. We report the unweighted mean, the me-

[16]Henceforth, we follow the convention that we winsorize the resulting markup estimates at the 5th and 95th percentiles and report them in percent above 1, i.e. $100 \cdot \log \hat{\mu}_{it}$. We will show that winsorizing markups has no impact on their trends over the firm life in Section 4.

dian and the employment-weighted mean. All three measures of the aggregate markup increase considerably over time: the mean markup, whether unweighted or employment-weighted, increases by around 10 percentage points (pp) over the sample period, from a level of around 5% in 2002 to around 15% in 2015, after which it moves sideways. The median markup, while consistently below the mean, increases by a similar amount, indicating that the average increase is not driving by just a few superstar firms.

That there is a seminal trend of increasing markups is consistent with De Loecker et al. (2020) and De Loecker and Eeckhout (2018). The latter find an even stronger increase for European firms of roughly 20pp over the same sample period. However, Marto (2024) recently documented that a large part of the increase in the average markup documented in this literature is due to services firms. As our sample only covers manufacturing firms, it is plausible that we find a flatter time profile. Our firms are more comparable to De Ridder et al. (2024), whose sample covers French manufacturing firms from 2010 to 2019. They find that the average markup does not increase significantly over that time frame, mostly because there is a reduction in markups during the Eurozone crisis in 2011-12, with a subsequent recovery. Denmark was much less exposed to this crisis than France was, so it is not surprising that our time profile is more upward-sloping than theirs.

Besides the increasing trend, there are interesting insights on the cyclicality of markups. First, deviations from the increasing time trend appear to be slightly procyclical, as is confirmed in appendix Figure B.3. However, the decrease in markups during recessions is small and brief. Following the global financial crisis, markups appear to even increase disproportionately, especially among large firms, as indicated by the dashed line.[17] Additionally, there is a temporary spike in markups of manufacturing firms of around 4pp in 2021, i.e., during the fast recovery from the Covid recession.

Figure 2(b) shows the cross-sectional distribution of markups. We report the full distribution of markups over the entire sample (grey shaded are) as well as kernel estimates of the markup distribution for the first two years of the sample (blue solid line), two years in the middle (red dashed line) and the two last years of the sample (orange dotted line). The figure illustrates that most of the markup distribution has shifted toward higher markups over the sample period, in line with the parallel increase of the mean and median markup.

When investigating the life cycle patterns of markups over the firm age, we will test our results

[17] Gilchrist et al. (2017) and Renkin and Züllig (2024) show that firms that lose access to external liquidity increase prices beyond credit cost, i.e., increase the markup to generate liquidity internally. Our findings are consistent with this, even though our markup estimation abstracts from financial frictions entirely.

Figure 2: Estimated markups

(a) Time trend

(b) Cross-sectional distribution

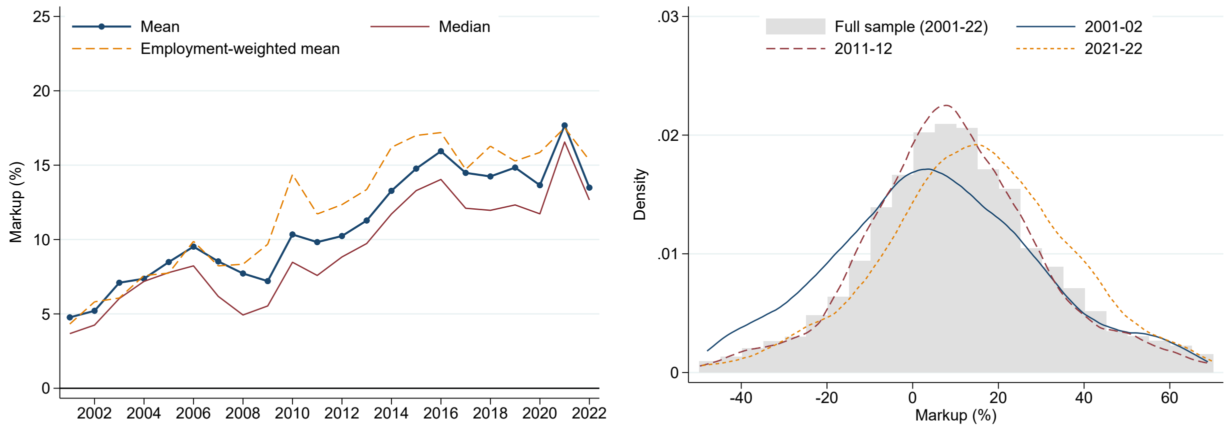

Notes: Aggregated markups estimated using the production function approach outlined in Section 3. Unbalanced sample of a total of 800 manufacturing firms. In panel (a), we first winsorize estimates at the 5th and 95th percentiles, then calculate $100 \cdot \log \hat{\mu}_{it}$ and subsequently depict the (weighted/unweighted) mean (or median).

with a range of alternative estimates for which we vary choices in the production function estimation. In particular, we estimate a version that is not based on firm-level deflators but uses industry-level deflators to calculate output. Furthermore, we vary the definition of inputs (e.g. using employment instead of deflated labor cost as the labor input or material inputs that are not corrected for the change in inventories) and of the first stage (following De Ridder et al. (2024)). In appendix Table B.2 and Figures B.1 and B.2, we show distributions and time trends of these alternative estimates as well as pairwise correlations with our main estimates. These correlations range from 0.66 to 0.91.

We find a consistent increase in markups over the sample period (2001-22) across all specifications of the production function estimation. Across all alternatives, the increase is on average 10pp, ranging from 5pp for the specification with the lowest increase to 15pp for that with the highest. However, the level of estimated markups differs, and markup estimates based on firm-level deflators are more dispersed than estimates based on industry-level deflators which ignore price dispersion between firms (even after winsorizing the estimated markups).

4 Markups over the firm life cyle

We now turn to our main analysis of how markups evolve over the life of the firm. Figure 3 shows averages and interquantile ranges of markups for different age groups, pooling cohorts over the whole 2001-22 sample period. Younger firms charge markups that are 5-10pp lower than those of older firms.

Figure 3: Markups by firm age

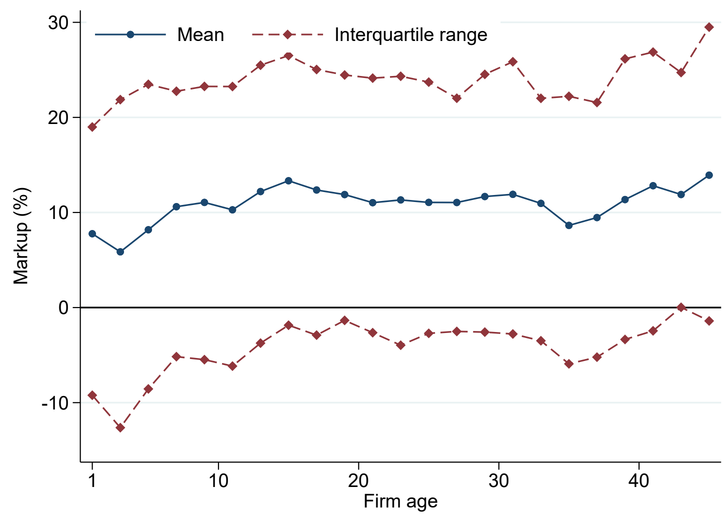

Notes: We estimate markups using the production function estimation approach where firm-level output is deflated using firm-level price deflators. The figure shows the mean and interquartile range of the markup distribution by age bin comprising two years each, but spanning the full sample period.

However, the averages shown in Figure 3 reflect a mixture of age effects, time effects and a firm composition changing over time. To address these issues, we estimate a standard additively separable decomposition of the log markup – or other firm-level outcomes such as prices or marginal costs – into contributions of age effects $f(a)$, sector-time effects τ and firm effects χ (which nest birth-cohort effects):

$$100 \cdot \log \mu_{it} = \beta_0 + f(a_{it}) + \tau_{k(i),t} + \chi_i + \varepsilon_{it} \quad . \tag{6}$$

We estimate age effects $f(a)$ as a 4th-degree polynomial in our main specification. This allows for smoothly evolving age patterns with enough flexibility to potentially show various nonlinear patterns of markups over the life of firms. These age effects are not necessarily all-else-equal causal relationships, but might reflect variation of other variables that evolve with firm age, such as firm size. In robustness checks further below, we control for some of these variables to identify age patterns holding other factors constant.

Our sample covers 22 years, but we observe starting dates for firms founded outside of the sample period and can thus identify age patterns for firms older than 22 years. However, different ranges of the age patterns that we estimate will be identified by different firms aging over this 22-year period.

4.1 Identification and estimation of age patterns

The central problem in estimating decomposition (6) is the collinearity of the linear components of age, time and cohort effects. Age is mechanically a linear combination of a firm's birth cohort and the current year—for any firm-year observation, a firm's age is equal to the current year minus the firm's birth cohort. As a result, any linear age trend $\alpha \cdot a$ is indistinguishable from a time trend $\alpha \cdot t$ that is offset by a negative cohort trend $-\alpha \cdot c$ for entering firms. To estimate unrestricted cohort, time and age effects separately, it would be necessary to observe firms of the same cohort at different ages at the same time. Clearly, that is not possible and we require further restrictions to identify (6). As discussed in more detail in McKenzie (2006), Fosse and Winship (2019) or Schulhofer-Wohl (2018), the collinearity problem applies specifically to the *linear* components of age, time and cohort effects and we need to place a restriction on one component to identify the other two. Nonlinear components can be pinned down from naturally occurring variation.

We employ two different complementary identification strategies that each place one restriction on the linear components. Both are commonly applied in age-time-cohort decompositions in different contexts, notably when estimating life cycle patterns in workers' earnings. The two restrictions yield very similar results on the evolution of markups over firms' lives, but leave some more uncertainty when applied to the dynamics of marginal cost and prices. In addition to the two identifying restrictions below, we always impose the normalization that firm effects sum to zero in the sample, and sector-time effects sum to zero within sectors (i.e. permanent markup differences between sectors are captured in firm effects).

Cohort trend restriction The first restriction is that firm effects in markups—χ_i in Equation (6)—are orthogonal to any linear cohort trend. Formally, we impose that $E(\chi_i N_i c_i) = 0$, where c_i denotes the birth cohort of firm i and N_i the number of years the firm is in the sample. This restriction enables us to identify unrestricted age and time effects, as well as the orthogonal cohort or firm effects. Markups might still systematically vary between cohorts—for example, firms that enter during a recession might have permanently different markups from those that enter during a boom.

An important possibility precluded by not allowing linear cohort trends in firm effects is a constant rate of selection on firm effects. For example, constant differences in exit rates between high- and low-markup firms would result in a linear cohort trend in the markups of surviving firms that we observe over the 2001–22 sample and that linear cohort trend could not be estimated without restrictions. In contrast, differences in exit rates between high- and low-markup firms that vary over firms' life would result in a non-linear pattern in cohort effects. For example, if low-markup firms exhibit a higher exit probability only during the first 10 years of life, the resulting cohort effects would be flat for firms born before 2012 and exhibit an exponential decrease for firms born during the last 10 years of our sample period. Such effects can be captured by our restricted estimates.

The same type of restriction is used in the context of estimating life cycle patterns in labor earnings by Deaton and Paxson (1994) or medical expenses by De Nardi et al. (2010). To estimate the coefficients in Equation (6) under the restriction that $E(c_i N_i \chi_i) = 0$ in practice, we adapt a procedure from Deaton (2019) that we explain in detail in Appendix C.1.

Age polynomial restriction To show that our results are not driven by the particular identification restriction chosen, we use an alternative restriction that constrains the age polynomial to have a stationary point at a chosen reference age, i.e $f'(\bar{a}) = 0$. This can be achieved by a simple variable transformation, because the restriction pins down the linear coefficient of polynomial $f(a)$, α_1, as a function of the higher-order terms. We describe further details on the exact implementation in Appendix C.1.

The age polynomial restriction has been used by Card et al. (2016, 2013) in the context of estimating life cycle patterns in labor earnings. It is in some ways more flexible than the cohort trend restriction. It enables us to identify unrestricted cohort and time effects, while still leaving a wide range of possible age patterns. The restricted age profile can be flat over the full life cycle, increasing or decreasing over the whole range of age except for $\bar{a}$, or feature different kinds of U-shapes or inverse U-shapes.

Our baseline choice of $\bar{a}$ is 25 years. The reference age should be high enough to guarantee that firms have matured and do not systematically grow towards their steady-state size. Based on the age profile of the size of Danish firms, the reference age should thus be at least 15 years (Andersen and Rozsypal, 2021). At the same time, $\bar{a}$ should be low enough to guarantee sufficiently many observations in the region where the flat age profile is imposed, i.e., no higher than 40 years (cf. Figure A.3). Additionally, we have deliberately chosen a reference age which is not equal to a (local) minimum or maximum of the age pattern in markups when estimating them subject to the cohort trend restriction.

Intuitively imposing two alternative restrictions can be seen as an informal over-identification test. If at least one of the restrictions is valid (i.e. not in conflict with the true parameters we would like to estimate), and estimation using both delivers similar results, then we can be confident in the validity of both restrictions.

4.2 Main results

We estimate regression (6) under both restrictions with log markups as the dependent variable, and repeat the same analysis with the same restrictions with log marginal cost and log prices as outcomes. We define marginal cost as the difference between estimated markups and observed price indices, but do not impose restrictions that would require the estimated age polynomials for markups and prices to sum up to the age polynomial for marginal cost. We show the estimated age profiles in Figure 4 and summarize the results by estimating the average changes over the first and second 20-year periods of the firm life in Table 2. Note that Figure 4 and subsequent figures plot the expected markup of firms by age relative to the sample average.

Markups Markups of new entrants are about 6pp below the sample average and significantly lower than those of older firms. Markups then increase and reach a maximum after about 20 years. In total, average markups increase by around 8pp over the first two decades of a firm's life. This is an economically large increase, given an average markup of 11% in the sample, and an increase of the average markup of 10pp over the 2001–2022 period. Once firms reach age 20, markups begin to stagnate or even fall slightly. These patterns are remarkably consistent between our two identifying restrictions, shown in panels (a) and (b).

Marginal cost We repeat the estimation with log marginal cost as the dependent variable. In contrast to markups, marginal cost are a nominal outcome, but factors affecting marginal cost of all firms in a given year—such as changes in input prices—should be captured by time effects. The estimated age patterns for marginal cost are illustrated by orange lines in Figure 4 and shown in the bottom panel of Table 2.

Figure 4: Age trends in markups, prices and marginal cost

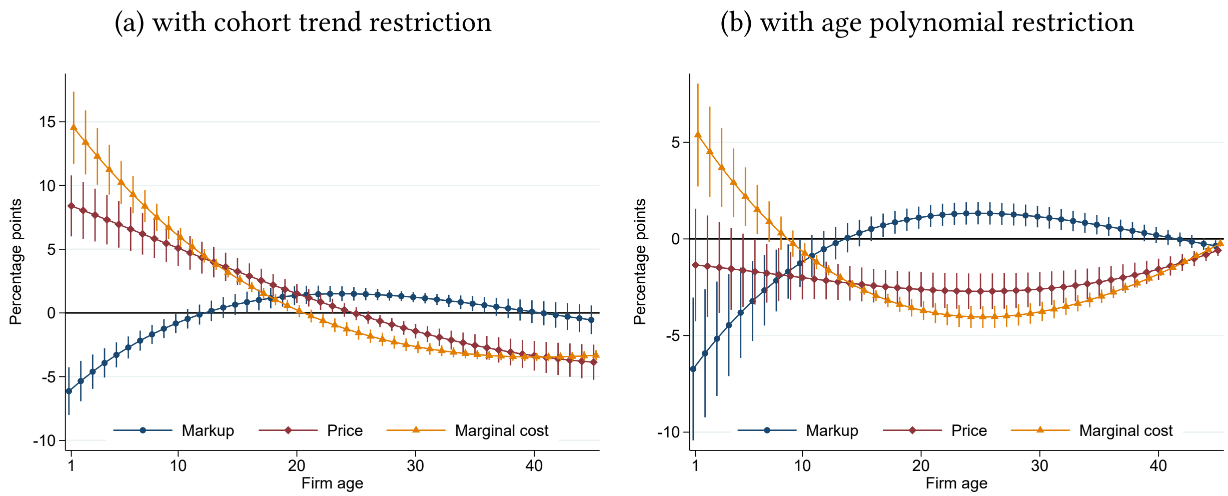

Notes: The figure plots estimated age patterns as 4th-order polynomials in markups, prices and marginal cost relative to the sample mean, based on estimates of Equation (6). As discussed in Section 4.1, two separate identifying restrictions are imposed to disentangle age, time and cohort effects in the panel. The figure includes 95% confidence intervals based on Driscoll-Kraay standard errors.

Marginal cost of entrants are significantly higher than the sample average and depending on the identification restriction, 9 to 14% higher than those of 20-year-old firms. The estimated age patterns are similar with both identifying restrictions over the first 20 years of a firm's life, but slightly flatter with the age polynomial restriction. For older firms, the two restrictions produce somewhat different patterns. If we restrict cohort trends, marginal cost falls by a further 4% between age 20 and age 40 and stagnates after. If we restrict the age polynomial, marginal cost increases by 2% between age 20 and age 40. This leaves some uncertainty over the marginal cost dynamics for firms above age 20.

Prices When we apply the cohort trend restriction to the evolution of quality-adjusted prices, we find that prices are monotonically decreasing in age. Relative to industry-time averages, prices fall by about 7% over the first two decades and another 5% in the subsequent two (see Table 2), by an average annual rate of 0.3%. This is consistent with Adam and Weber (2023), who find decreasing age trends in relative prices for consumer items. When we impose this relative price trend to be flat at the reference age, the decreasing age pattern during the first 20 years of a firm's becomes much flatter (red dots in Figure 4(b)) and there is an increasing age pattern for prices of

Table 2: Changes of markups, prices and marginal cost among young and mature firms

Identification	(1) Cohort trend restriction	(2) Age polynomial restriction
Markups		
Change from age 1 to 20	7.50***	7.83***
	(1.26)	(1.79)
Age 20 to 40	-1.26	-0.91***
	(0.80)	(0.15)
Prices		
Change from age 1 to 20	-6.89***	-1.26
	(1.29)	(1.42)
Age 20 to 40	-4.87***	1.04***
	(0.92)	(0.22)
Marginal cost		
Change from age 1 to 20	-14.38***	-9.10***
	(1.63)	(1.59)
Age 20 to 40	-3.61***	1.95***
	(0.31)	(0.19)
Observations	9,700	9,700
Firms	798	798

Notes: Estimates of Equation (6) under two different restrictions to identify the age polynomial $f(a)$. We report the implied percentage point change of markups from age 1 and 20, i.e. $100 \cdot \log \hat{\mu}(a = 20) - 100 \cdot \log \hat{\mu}(a = 1)$ and from 20 to 40 years, respectively. Marginal cost are defined as the log difference between the firm price deflator and the estimated markup. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

firms above age 20.[18]

Taken together, our main results show that young firms have steeply increasing markups, and declining marginal costs that are only partially passed through to prices. We show in Section 5 that this result is robust to a range of alternative choices in our estimation. Markups stop increasing around age 20 and slightly decline for older firms. The results for marginal cost and prices of older firms are less clear. Under both restrictions marginal cost and price dynamics become more similar for older firms, consistent with relatively stable markups.

[18]Note that the case that the restriction of a flat age profile at age 25 holds for prices is theoretically much weaker than it is for markups and we therefore put more weight on the results using the cohort trend restriction in panel (a).

Cohort trends Our first identification restriction requires absence of a cohort trend in firm effects. In Figure D.1 in the appendix, we test for the presence of this (linear) cohort trend when we impose the age polynomial restriction, where the cohort trends remain unrestricted. The result is that the restriction is violated at the 95% significance level, indicating that our two identification strategies are to some extent complementary.

4.3 The contribution of firm aging to the increase of the average markup

The average firm age increased by around 15 years over the course of the 22-year sample period, while the average markup has increased by around 10pp. In this section, we ask whether and to what extent the former trend directly contributed to the latter.

Our estimates of identified age, time(-sector) and firm (which nests cohort) effects now allow us to decompose changes of the average markup over time into these three components. Figure 5 decomposes the fitted values of our two models (with different identification restrictions) into contributions by, i.) changes in the age composition of firms (blue), ii.) entry and exit of firms with heterogeneous firm effects (red) and iii.) time effects that affect all active firms equally (orange). The latter is estimated separately for each industry and then averaged across industries using the number of observations in each year.

Despite the clear age patterns presented before, the role of firm aging for average markups is close to zero. We investigate this more closely by further decomposing the age effect into contributions from firms below the age of 20 whose aging should— according to our estimations—have an unambiguously positive effect on the average markup. In appendix Figure D.2, we show that their contribution is around 1pp, still falling short of the increase of the average markup of 10pp by an order of magnitude. On the other hand, the further aging of the firms older than 20 years contributes negatively to the average markup, as some of them have decreasing markups. The net age effect of the two age groups is then very close to zero.

We find that the fixed firm effects contribute to a continuously increasing markup, in total about 2pp over the sample period. This indicates that over time, high-markup firms replace low-markup firms.

The largest share of the markup increase is explained by time-sector fixed effects. We present them at the sector level in Figure D.3 in the appendix. All of the 15 sectors have a clearly visible increasing time trend under both identification restrictions. Therefore, the aggregate increase is not driven by a changing sector composition. In fact, the contribution from reallocation from

Figure 5: Decomposition of the average markup increase

(a) with cohort trend restriction (b) with age polynomial restriction

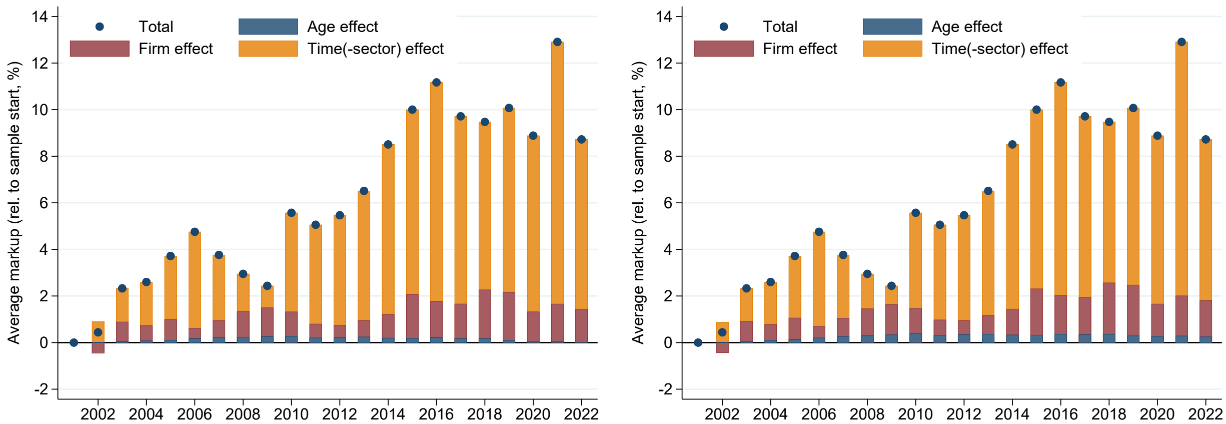

Notes: Having identified the age, firm and time effects in Equation (6) under two different identification restrictions, we decompose the change in the average markup into these three components by taking fitted values of the components in our two estimated models and aggregate them over all firms by year.

low- to high-markup sectors is virtually zero in our sample. Another interesting finding is that the time effects are more strongly pro-cyclical than the aggregate markup itself—from 2007 to 2009, it decreased by 2pp. This is partially offset by a composition effect, as firms with low markups were more likely to exit over the Great Recession, which drove up the average markup.

5 Robustness checks

Our key findings are that marginal cost decline over the first 20 years of firms' existence, and that this cost advantage does not (fully) pass through to price changes, but instead markups rise. In this section, we show that these age patterns are robust not only to the identification restriction needed to identify age patterns, as already shown in Section 4.2, but also to numerous alternative choices in the estimation of markups, specification choices in the estimation of age patterns, as well as several sample restrictions.

Estimation of production functions and markups Our analysis depends on well-estimated firm-level markups, the estimation of which requires a multitude of choices (described in Section 3).

A critical choice is whether or not to deflate firm-level sales with firm- or industry-level prices to measure firm output and estimate output elasticities. The latter is standard in the literature following De Loecker and Warzynski (2012), but a number of papers, e.g. Bond et al. (2021), have highlighted the importance of deflating with firm-level prices for the identification of the *level* of markups. De Ridder et al. (2024) have documented that observing physical output is much less important to track markup *changes*, and since all our specifications include a firm fixed effect, we would not expect sector-level deflators to have a large impact on our results. The results in columns (1) and (2) of Table 3 confirm this. While qualitatively similar, the magnitudes of the increasing (decreasing) age pattern of markups (marginal cost) is about 2pp weaker over the first 20 years.

Table 3: Robustness of age pattern in terms of markup estimation

	(1)	(2)	(3)	(4)
	Markup estimates with industry price deflator		Markup estimates as De Ridder et al. (2024)	
Identification:				
Cohort trend restriction	✓		✓	
Age polynomial restriction		✓		✓
Markups				
Change from age 1 to 20	5.71***	6.01***	4.01***	5.89***
	(1.15)	(1.08)	(0.86)	(1.05)
Age 20 to 40	-1.05***	-0.73***	-2.58***	-0.61***
	(0.25)	(0.12)	(0.49)	(0.10)
Marginal cost				
Change from age 1 to 20	-12.38***	-7.04***	-10.90***	-7.15***
	(2.17)	(2.09)	(1.89)	(2.17)
Age 20 to 40	-3.87***	1.75***	-2.29*	1.65***
	(1.11)	(0.29)	(1.33)	(0.27)
Observations	9,673	9,673	9,700	9,700
Firms	798	798	798	798

Notes: Percentage point and percent changes of markups and marginal cost with variations in markup estimation. Columns (1) and (2) use sector-level deflators to compute firm-level output instead of our PPI-based firm-level deflators. Columns (3) and (4) use additional control variables in the first stage of the production function estimation, namely the firm's market share and log price. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The second robustness check concerns the first stage of the production function estimation, which purges measurement error and shocks that occur after production decisions are made—ϵ in the notation of Section 3—from output. Doraszelski and Jaumandreu (2021) show that under certain conditions, the markup itself should be part of the control variables to help identify this term. Because the markup is unobserved in the first stage, De Ridder et al. (2024) propose to include a firms' log price as well as its market share as proxies. Our main results are robust to adding these two variables among the controls in the first stage of the production function estimation (see columns (3) and (4) of Table 3).

We also implement a series of robustness checks related to choices in the production function estimation. In our main specification, we use deflated payroll expenditures as a measure of labor input. Alternatively, we can use full-time equivalent employment, which is directly measured as a real variable, but the disadvantage that it is not adjusted for the quality of labor input. Additionally, our main specification uses an inventory correction for material inputs (by the change of intermediate inputs) and outputs (by the change of final-goods inventories) to reflect production in a year, rather than sales, as is convention in accounting statistics.[19] While more correct, the role of inventories is often disregarded in the markup estimation literature due to data limitations. In Table D.1 in the appendix, we show that neither the inventory adjustment nor our choice of labor input have a substantial impact on the estimated age patterns.

We also address outliers in our sample of estimated markups. These are subject to measurement error and show a large variance, particularly when estimated with firm price deflators (see Figure B.2). Therefore, we have winsorized left-hand side variables at the 5th and 95th percentiles in our main specification. In Table D.2 we include a version with unwinsorized markup and marginal cost estimates, which leads to somewhat larger standard errors and a version where we winsorize more strongly, namely at the 10th and 90th percentiles. Both of them confirm the increasing age pattern of markups and decreasing marginal cost trend for young firms.

Estimation of age patterns Age correlates with both size and the financial position of firms. To assess whether our age patterns are driven by these factors, we include additional control variables, for example polynomials of employment and real output as we use it in the production function estimation, the log market share of the firm within its 2-digit sector, the leverage ratio as well as the ratios of cash to sales and cash to total assets. Table D.3 shows that the markup increase does not diminish with these control variables, i.e., the markup increase among young firms is not driven by firms increasing in size. Interestingly, the decrease in marginal cost becomes somewhat less pronounced when controlling for firm characteristics, indicating that decreasing marginal cost might to some extent root in increasing returns to scale of young firms.

[19] We have shown in Figure B.2 that this can change both the time trend and the cyclical properties of estimated markups.

When we identify age patterns using the age polynomial restriction, we define a specific firm age at which the slope of the polynomial is restricted to be zero–25 years in our main specification. In Table D.4 we report the results for alternative reference ages between 15 and 35 years, respectively. The result that markups rise significantly and marginal cost fall is robust to the choice of the reference age, but the slopes increase (in absolute terms) when the reference age is increased. In Table D.5 in the appendix we show that our results do not depend on the order of the polynomial we choose to approximate $f(a_{it})$.

Sample As we discussed in Section 2, our sample of firms is not representative for the population of Danish manufacturing–the firms in our sample are larger and older than the typical manufacturing firm. Our sample might also contain more fast-growing young firms that cross sampling thresholds quicker than others. To alleviate concerns that this drives our results, we repeat the estimation with sampling weights that adjust the weight of age-size cells to the correct weight in the population of Danish manufacturing firms (see Figure A.4). The results, shown in Table 4, show patterns among young firms that are at least as strong as in the unweighted sample.

The slope of our age profiles is estimated from piecewise variation in age, i.e., from firms we observe in at least two subsequent years. Because firms have different ages at the start of the sample, we can estimate an age pattern for a long age range, even though we observe each firm for a maximum of 22 years. However, each year-on-year change is conditional on the firm surviving at least up to that point. For example, firms which exit between age 10 and age 11 will not be used to estimate the markup change between these two age points. If exit is dependent on (unobserved) markup changes, this might bias our age profile. Unfortunately, the sample size does not allow to estimate age profiles conditional on exit. What we can do, however, is to show estimates based on a sample of firms for which we know they eventually reach 20 years of age in columns (3) and (4) of Table 4. The markup increase among young firms is stronger. This is indicative of some degree of firm selection on markup profiles among young firms. Interestingly, the estimated pattern for marginal cost is not sensitive to subsetting our sample to only surviving firms.

We also make sure that our results are not specific to particular large sectors by excluding the two largest ones—food production and machinery production—from our sample. As shown in the last two columns of Table 4, excluding these two sectors increases the magnitudes of the age patterns in markups over the first 20 years by around 3pp, while leaving the decrease in marginal cost largely unchanged.

Table 4: Robustness of age pattern to sample

	(1)	(2)	(3)	(4)	(5)	(6)
	Sampling weights		Conditional on reaching age 20		Excluding largest sectors	
Identification:						
Cohort trend restr.	✓		✓		✓	
Age polynomial restr.		✓		✓		✓
Markups						
Age 1 to 20	8.66***	8.60***	11.75***	11.03***	10.38***	10.87***
	(1.32)	(1.20)	(1.38)	(1.87)	(1.52)	(2.02)
Age 20 to 40	-1.00**	-1.06***	-0.61	-1.37***	-1.82**	-1.31***
	(0.47)	(0.19)	(0.86)	(0.15)	(0.86)	(0.18)
Marginal cost						
Age 1 to 20	-15.19***	-9.52***	-14.06***	-7.73***	-14.57***	-8.53***
	(1.12)	(1.02)	(2.26)	(1.86)	(1.77)	(1.30)
Age 20 to 40	-3.93**	2.04***	-4.89***	1.77***	-4.30***	2.05***
	(0.40)	(0.16)	(0.33)	(0.25)	(0.50)	(0.20)
Observations	9,700	9,700	8,656	8,656	6,086	6,086
Firms	798	798	645	645	501	501

Notes: Percentage point and percent changes of markups and marginal cost with variations in the sample. Columns (1) and (2) use firm-level weights to adjust the joint distribution of firm size and age to that of the overall population of manufacturing firms in Denmark, in particular to scale down variation used from large firms, which are over-represented. (3) and (4) conditions the sample on firms for which we know that they become at least 20 years old, (5) and (6) exclude the food and machinery manufacturing sectors. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 Sources of age patterns

Next, we explore two possible sources for the age patterns documented so far. The first one asks whether they are driven by productivity or scale effects; the second one takes a more granular look at the product level and documents that marginal cost reductions by young firms coincide with both the introduction and the discontinuation of new products.

6.1 Production function origins

Minimizing the cost function given output leads to a marginal cost function that negatively depends on TFP. In the case of Cobb-Douglas production with constant returns to scale, $\log mc_{it} \propto$

$-\omega_{it}$. In our case of the translog production function,

$$\log mc_{it} \propto -\omega_{it} + h(k_{it}, l_{it}, m_{it}, \boldsymbol{\beta}), \tag{7}$$

marginal cost depend linearly on the negative of TFP as well as a nonlinear function in the factor inputs and production function parameters. The latter component exists due to non-constant returns to scale stemming from the non-homotheticity of the production function. In the following, we decompose the decrease in marginal cost into these two components.

Given our estimated production functions, we back out TFP from Equation (4):

$$\hat{\omega}_{it} = y_{it} - \hat{\varepsilon}_{it} - \hat{f}(k_{it}, l_{it}, m_{it}). \tag{8}$$

For these estimates of TFP, we then repeat the age polynomial estimation under the cohort trend restriction described in Section 4.1.[20] The green line in Figure 6 reports the estimated age pattern inverted–because an increase in TFP decreases marginal cost–and normalized to 0 at age 1.

Figure 6: Age pattern in marginal cost and TFP

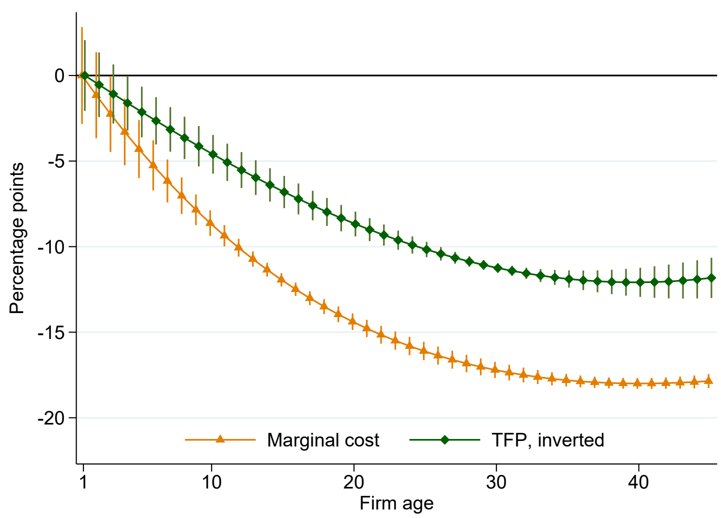

Notes: The orange lines shows estimates of marginal cost over the firm life, as already shown in Figure 4(a); the green line repeats the age polynomial estimation with estimates of TFP from the translog production function. 95% confidence intervals based on Driscoll-Kraay standard errors.

Firm-level TFP strongly and continuously decreases over the firm life and flattens out only after more than 30 years. This leads to a direct decrease in marginal cost of about 12%.

[20] We do not use our second identifying age polynomial restriction because there is no theoretical justification for a flattening of TFP relative to the market average at any reference age.

The fact that marginal cost decreases somewhat stronger than TFP increases suggests that there is a role for non-constant returns to scale. If these are increasing in scale, the decrease in marginal cost unexplained by TFP, i.e., the difference between the orange and green lines, could be driven by young firms increasing in size. Beyond size, it could also be that young firms start using a more efficient bundle of inputs, as they learn about the nonlinearity of their production function.

6.2 Product turnover

Finally, we explore whether life cycle patterns of markups and marginal cost are associated with changes of firms' portfolio of product produced, i.e., with the introduction of new products and the discontinuation of previously produced ones.

We link our firm-level data with survey data on product-level sales of Danish manufacturing firms. In this survey (which is the basis of the Danish PRODCOM statistics), manufacturing firms with more than 10 employees report their quarterly sales at the level of 8-digit CN codes. The advantage of this dataset is that it covers all products of all firms, rather than a selection of products for a selection of firms as in the PPI, allowing us to get a full picture of the products produced in each firm and year. The disadvantage of the data is that it contains more measurement error, both in the number of products–particularly due to changes in the reported CN codes–and in the reported sales and quantities (and units).[21] Nevertheless, we can measure the approximate number of products, or CN8 codes, produced by the 800 firms in our sample of markups. Another drawback is that we only have the product-level data up to 2019. For these years, we can match 88% of our panel of firm markups to an observation in the PRODCOM statistics.

The mean and median numbers of reported CN8 codes per firm are 5.9 and 2, indicating a highly skewed distribution of products among a relatively small number of multiproduct firms. Around a quarter of all firm identifier never report more than one CN8 codes.

We define two new dummy variables, namely d^N for new product introductions, equal to one whenever positive sales of a new 8-digit CN code at the firm level are reported. This is the case for 16% of firm-year observations.[22] Similarly, we define d^D for product discontinuations, i.e., whenever a CN code is reported by that firm for the last time and is no longer observed in the subsequent year, which is the case for 18% of all observations.

Interestingly, product turnover—both introductions and discontinuations—is substantially higher

[21]This is why we do not use the unit values we can compute in this data as the firm-level price deflators for the computation of markups.

[22]We exclude new 8-digit CN codes that might appear due to changes in the CN classification.

among young firms. Figure 7(a) shows that among firms less than 20 years old, around 20% of firms introduce a new product in a given year and between 20 to 25% discontinue at least one product. As a consequence, the number of products that the typical firm produces decreases in age. For firms aged 20 and older, the number of new and discontinued products is about 15% a year. Panel (b) confirms higher product turnover among young firms at the product level: Among very young firms, around 7% of all observed product codes are new products, whereas among firms above 20 years of age, this share decreases to around 4%.

The shapes of these age patterns roughly coincide with the firm life-cycle patterns of markups and marginal cost estimated in Section 4.2. This suggests that a higher rate of product rotation could be a source of markup increases for young firms. Young firms could experiment with different products, and keep those which they are either effective at producing, or where demand conditions allow them to charge high markups. In both cases, we would expect average firm-level markups to increase on average after a product has been either discontinued or introduced.

To investigate this hypothesis further, we run local projections of the form

$$100 \cdot (\log x_{i,t+h} - \log x_{i,t-1}) = \beta_h^E d_{it}^E + \tau_{k(i),t+h} + u_{i,t+h} \quad , \tag{9}$$

where x denotes either the firm-level markup, price, or marginal cost and $E = \{N, D\}$ is either a new product introduction or discontinuation event in year t. The coefficients β_h^N and β_h^D for $h = \{-3, -2, ..., 3\}$ trace out the behavior of firm-level variables around these events. Estimates from these regressions are purely descriptive and should not be interpreted as causal, as both the introduction and discontinuation of products are endogenous firm choices. Aggregate trends are absorbed by the sector-time fixed effects τ.

Panels (c) and (d) of Figure 7 depict estimates of β^N (left) and β^D (right) with patterns around events of product turnover. When a new product is introduced, firm-level markups gradually rise by around 1pp. Firm-level price deflators first increase by approximately the same amount, reflecting the fact that marginal cost do not fall on impact. Three years after, however, marginal cost fall, such that a 1.5pp increase in the markup leads to a (quality-adjusted) price increase of 1pp.

When a product is discontinued, on the other hand, marginal cost fall and markups increase immediately. The size of both effects is around 1pp in year $t + 1$, which is the first full calendar year in which the discontinued product is no longer observed. Prices, however, do not change in a statistically significant way.

Figure 7: Product turnover and markup and marginal cost changes

(a) Firms with new/discontinued products

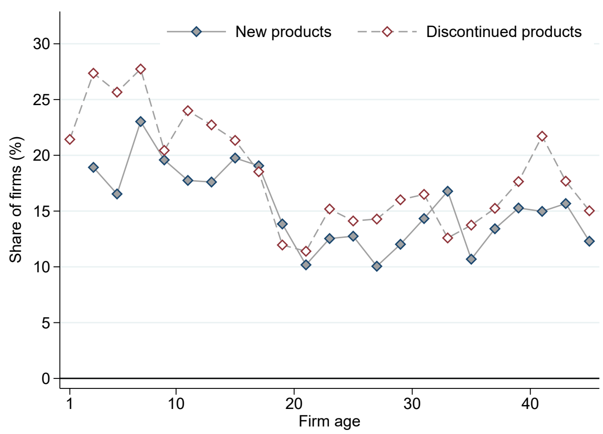

(b) New/discontinued products

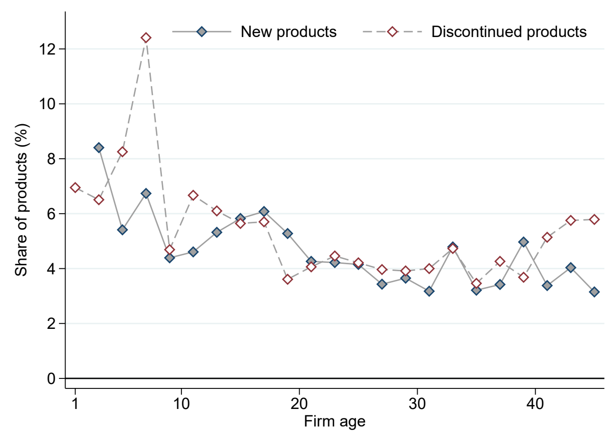

(c) μ and mc around new product introductions

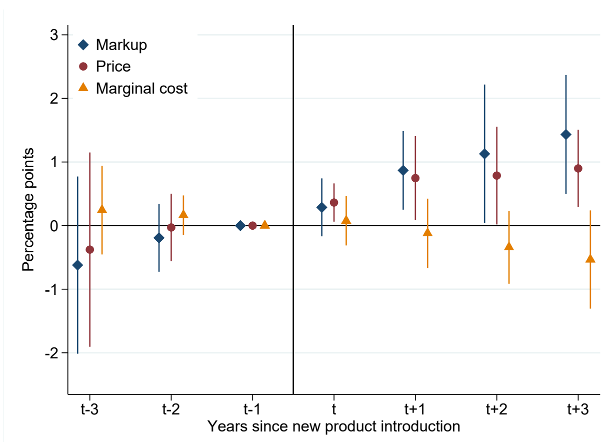

(d) μ and mc around discontinuation of products

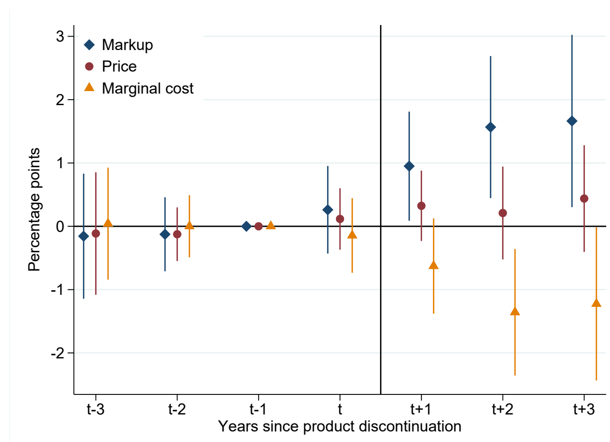

Notes: (a) shows the share of our sample firms with at least one CN8 code reported for the first time by that firm in a given year. The x-axis combines two years of age at a time. The dummy is not defined for firms of age 1, as all their products are new. Discontinued products shows the share of firms which report at least one CN8 code for the last time. Panel (b) reports the same statistics, but as a share of all products instead of as a share of firms. The lower two panels show local projection estimates (9) for the estimated markups (and prices and marginal cost) on a dummy for whether or not the firm has introduced or discontinued at least one product in a given year. 95% confidence intervals based on Driscoll-Kraay standard errors.

Together, these results suggest that firms use both new product introductions and product discontinuations to upgrade their markup in a statistically and economically significant way.[23] They also suggest that higher rates of product introduction and discontinuations among younger firms are important factors driving the increase in markups (and the fall in marginal cost) over the first two decades of firms' lives.

7 Implications and conclusion

We think that our empirical results have a number of important implications for macroeconomic research.

Implications for aggregate markup trends A first important take-away from our results follows directly from Section 4.3 and is that firm aging has at best a small direct effect on the aggregate markup. Even though markups increase substantially over the first 20 years of a firms' life, young firms have a small weight in the aggregate. In both our structural decompositions, the largest part of the increase in markup observed over the sample period comes from time effects that increase markups for all firms equally. While this still leaves many possible sources of increasing aggregate markups, it rules out several explanations that rely on a changing composition of firms across sectors, age or cohorts. First and foremost, it rules out that the evident aging of the average firm that follows from declining business dynamism was a significant driver of the aggregate markup. What it does not rule out is indirect effects of declining business dynamism, for example a lower number of competitors, which allows all remaining firms to charge higher markups (see e.g. Akcigit and Ates (2021)). Second, it excludes the possibility that the average markup increase is driven by "superstar firms", which in our decomposition would be attributed to the firm fixed effects.

Implications for models of firm dynamics Many benchmark models of firm dynamics rely on borrowing constraints to generate realistic patterns of firm dynamics. Young firms are typically constrained, so profits and markups are their primary determinant of growth, as they allow them to grow out of their borrowing constraints. If markups of young firms are systematically

[23] Notice that we estimate changes in markups around product turnover only. Firms might additionally increase markups on existing products. Decreasing product-level markups can also be consistent with a firm-level markup increasing in firm age, as long they can charge substantially higher markups on new products compared to existing ones and have a sufficiently high rate of product introduction. However, the data do not allow us to estimate product-level markups to test this.

low, as is the pattern we document in this paper, it would slow down the rate at which young firms grow to their optimal size.

In most firm dynamics models, markups are constant, and such a channel is absent. Ignaszak and Sedláček (2023), Roldan-Blanco and Gilbukh (2023) and Chiavari (2024), among others, build models in which firms can invest in their customer base. This yields markup age patterns consistent with our results, but in these models, low markups of young firms are driven by low prices, rather than high cost as in our results. The facts we document are rather consistent with a setup that features strategic complementarity in price-setting and young firms that are less productive than older firms. While this is beyond the scope of this paper, it would be interesting to integrate such a setup in a firm dynamics models to quantify the impact of age dynamics in markups on the firm size distribution.

Implications for models of creative destruction The link between product turnover, declining marginal cost and rising markups has implications for Schumpeterian models of innovation. In these models (see e.g. Klette and Kortum (2004), Akcigit and Kerr (2018)), innovation happens when a firm improves an existing product, introduces a new one (both internal innovation) or when a new firm adds a new product to the market to compete with existing ones (creative destruction). Garcia-Macia et al. (2019) argue that incumbents are responsible for an outsized share of innovation, and within incumbent firms, improving existing products is more important than introducing new ones. The latter is challenged by empirical evidence in Argente et al. (2024), which shows that new products are the main source of sales growth at the firm level and that sales within existing products decline over time. Our finding of persistent declines in firm-level marginal cost supports the view that internal innovation is indeed important. However, we show that within firm-level developments, changes in the product portfolio *are* an important contributor. When firms decide whether to introduce a new product, they face a trade-off between being able to better compete for market share and cannibalizing their own existing products (Argente et al., 2024). Our finding is that after events of product introductions actually observed in the data, i.e. considered optimal by the firm, the firm-level price and markup both increase persistently, indicating that new products do not necessarily cannibalize the existing product portfolio.

Implications for optimal inflation Our empirical findings even have implication for monetary policy. Adam and Weber (2019) show that in an economy with heterogeneous firms whose marginal cost decrease in age and who face nominal rigidities, the optimal inflation target is positive. The intuition is as follows: In an efficient economy, the allocation of production is such that relative prices reflect the relative marginal cost of firms. If marginal cost of firms are decreasing over a firm's life, this can happen in two ways. First, if the aggregate price level is constant (i.e., there is no aggregate inflation), the nominal price of each firm relative to the aggregate must decline over time. New firms charge a higher price than the average firm. Over time, they lower it, but the new higher prices of entrants keeps the aggregate price level unchanged. The second case is that the nominal price of a firm is constant over time (and thus in its age), but in that case, the aggregate price level increases. Under sticky prices, the second solution is the more optimal one, as no nominal prices need not be changed. This finding stands in sharp contrast to a canonical literature which finds a zero-inflation steady state to be optimal.

To estimate the optimal inflation target implied by their theory, Adam et al. (2022) and Adam and Weber (2023) estimate (linear) age trends of relative prices in CPI micro data for the UK, Germany, France and Italy. However, these relative price trends only act as a proxy for trends in (unobserved) relative marginal cost. Our consistent result is that markups increase in age, at least for large parts of the firm age distribution. This means that relative price age trends overestimate trends in relative marginal cost.[24] As a consequence, estimating the optimal inflation target based on relative price trends alone likely underestimates the optimal inflation target.

Conclusion Using price and production data, we have estimated markups for manufacturing firms in Denmark. We have documented the evolution of markups, prices and marginal cost over the life of these and found strongly decreasing age patterns in marginal cost over the first two decades, which is partially reflected in an increasing markup. These findings are important for several fields studying either firm behavior or macroeconomic outcomes and policy.

References

Ackerberg, Daniel, Kevin Caves, and Garth Frazer, "Identification Properties of Recent Production Function Estimators," *Econometrica*, 2015, *86* (6), 2411–2451.

[24]While the dependent variables in our regressions with markups, prices and marginal cost enter in absolute (log) terms, the inclusion of sector-time fixed effect makes the interpretation equivalent to that of markups, prices and marginal cost relative to an aggregate trend, just like in the papers cited above.

Adam, Klaus and Henning Weber, "Optimal trend inflation," *American Economic Review*, 2019, *109* (2), 702–737.

— **and** — , "Estimating the optimal inflation target from trends in relative prices," *American Economic Journal: Macroeconomics*, 2023, *15* (3), 1–42.

— **, Erwan Gautier, Sergio Santoro, and Henning Weber**, "The Case for a Positive Euro Area Inflation Target: Evidence from France, Germany and Italy," *Journal of Monetary Economics*, 2022, *132*, 140–153.

Akcigit, Ufuk and Sina T. Ates, "Ten facts on declining business dynamism and lessons from endogenous growth theory," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 257–298.

— **and William R. Kerr**, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

Andersen, Svend Greniman and Filip Rozsypal, "What can old firms tell us about the effect of age on firm size?," *Danmarks Nationalbank Working Paper*, 2021.

Argente, David, Munseob Lee, and Sara Moreira, "The Life Cycle of Products: Evidence and Implications," *Journal of Political Economy*, 2024, *132* (2), 337–390.

Biondi, Filippo, Sergio Inferrera, Matthias Mertens, and Javier Miranda, "Declining job reallocation in Europe: The role of shocks, market power, and technology," 2025.

Blundell, Richard and Stephen Bond, "GMM Estimation with Persistent Panel Data: an Application to Production Functions," *Econometric Reviews*, 2000, *19* (3), 321–340.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch, "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data," *Journal of Monetary Economics*, 2021, *121*, 1–14.

Calvino, Flavio, Chiara Criscuolo, and Carlo Menon, "A cross-country analysis of start-up employment dynamics," *Industrial and Corporate Change*, August 2018, *27* (4), 677–698.

— , — **, and Rudy Verlhac**, "Declining business dynamism: Structural and policy determinants," OECD Science, Technology and Industry Policy Papers 94 November 2020. Series: OECD Science, Technology and Industry Policy Papers Volume: 94.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, May 2016, *131* (2), 633–686.

_ , Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, August 2013, *128* (3), 967–1015.

Chiavari, Andrea, "Customer Accumulation, Returns to Scale, and Secular Trends," *mimeo*, 2024.

_ , Marta Morazzoni, and Danila Smirvov, "Heterogeneous Markups Cyclicality and Monetary Policy," 2023.

Conlon, Christopher, Nathan H. Miller, Tsolmon Otgon, and Yi Yao, "Rising Markups, Rising Prices?," *AEA Papers and Proceedings*, May 2023, *113*, 279–283.

De Loecker, Jan, "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity," *Econometrica*, 2011, *79* (5), 1407–1451.

_ and Frederic Warzynski, "Markups and firm-level export status," *American Economic Review*, 2012, *102*, 2437–2471.

_ and Jan Eeckhout, "Global Market Power," *NBER Working Paper*, 2018, *24768.*

_ , _ , and Gabriel Unger, "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

De Ridder, Maarten, Basile Grassi, and Giovanni Morzenti, "The Hitchhiker's Guide to Markup Estimation," *mimeo*, 2024.

Deaton, Angus, *The Analysis of Household Surveys (Reissue Edition with a New Preface): A Microeconometric Approach to Development Policy*, Washington, DC: World Bank, January 2019.

_ and Christina Paxson, "Saving, Growth, and Aging in Taiwan," in "Studies in the Economics of Aging," University of Chicago Press, January 1994, p. w4330.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda, "Where has all the skewness gone? The decline in high-growth (young) firms in the U.S.," *European Economic Review*, July 2016, *86*, 4–23.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The role of entrepreneurship in US job creation and economic dynamism," *Journal of Economic Perspectives*, 2014, *28* (3), 3–24.

Dedola, Luca, Mark Strøm Kristoffersen, and Gabriel Züllig, "Price Synchronization and Cost Pass-through in Multiproduct Firms: Evidence from Danish Producer Prices," *mimeo*, 2019.

Doraszelski, Ulrich and Jordi Jaumandreu, "Reexamining the De Loecker & Warzynski (2012) method for estimating markups," *CEPR Discussion Paper*, 2021, *DP16027.*

Fosse, Ethan and Christopher Winship, "Bounding Analyses of Age-Period-Cohort Effects," *Demography*, October 2019, *56* (5), 1975–2004.

Foster, Lucia, John C. Haltiwanger, and C. J. Krizan, "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in "New Developments in Productivity Analysis," University of Chicago Press, January 2001, pp. 303–372.

_ , John Haltiwanger, and Chad Syverson, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, March 2008, *98* (1), 394–425.

_ , _ , and Cody Tuttle, "Rising markups or changing technology?," *NBER Working Paper*, 2024, *30491.*

Garcia-Macia, Daniel, Hsieh Chang-Tai, and Peter J. Klenow, "How Destructive Is Innovation?," *Econometrica*, 2019, *87* (5), 1507–1541.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek, "Inflation Dynamics during the Financial Crisis," *American Economic Review*, March 2017, *107* (3), 785–823.

Ignaszak, Marek and Petr Sedláček, "Customer Acquisition, Business Dynamism and Aggregate Growth," *mimeo*, 2023.

Kalemli-Ozcan, Sebnem, Bent Sorensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yesiltas, "How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications," *NBER Working Paper*, 2022, *21558.*

Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Levinsohn, James and Amil Petrin, "Estimating production functions using inputs to control for unobservables," *Review of Economic Dynamics*, 2003, *70* (2), 317–341.

Marto, Ricardo, "Structural Change and the Rise in Markups," *mimeo*, 2024.

McKenzie, David J., "Disentangling Age, Cohort and Time Effects in the Additive Model*," *Oxford Bulletin of Economics and Statistics*, 2006, *68* (4), 473–495.

Melitz, Marc J. and Sašo Polanec, "Dynamic Olley-Pakes productivity decomposition with entry and exit," *RAND Journal of Economics*, 2015, *46* (2), 362–375.

Nakamura, Emi and Jón Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

Nardi, Mariacristina De, Eric French, and John B. Jones, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, February 2010, *118* (1), 39–75.

Olley, Steve G. and Ariel Pakes, "The dynamics of productivity in the telecommunications equipment industry," *Econometrica*, 1996, *64* (6), 1263–1297.

Pugsley, Benjamin Wild and Ayşegül Şahin, "Grown-up Business Cycles," *The Review of Financial Studies*, March 2019, *32* (3), 1102–1147.

Renkin, Tobias and Gabriel Züllig, "Credit supply shocks and prices: Evidence from Danish firms," *American Economic Journal: Macroeconomics*, 2024, *16* (2), 1–28.

Roldan-Blanco, Pau and Sonia Gilbukh, "Firm Dynamics and Pricing Under Customer Capital Accumulation," *mimeo*, 2023.

Schulhofer-Wohl, Sam, "The Age-Time-Cohort Problem and the Identification of Structural Parameters in Life-Cycle Models," *Quantitative Economics*, 2018, *9* (2), 643–658.

van Vlokhoven, Has, "The rise in profits and fall in firm entry: A tale of the life cycle of profits," *mimeo*, 2021.

Appendices

A Details on the data

A.1 Firm- and industry-level output and input deflators

A.1.1 Firm-level output deflators from PPI data

In the PPI survey, producers report each month the transaction price quote of a representative and time-consistent set of their products, including temporary sales.[25] We refer to $\tilde{P}_{jim}$ as the reported nominal transaction price for good j of firm i in month m. Over the sample we use (2001–22), the raw PPI data covers 1,220 firms and 8,512 products in total. This covers both the domestic and export waves of the PPI survey. If a good is sold both domestically and internationally, we treat them as separate products.

One advantage of this data is that price of individual products is tracked for a relatively long time, i.e., we have few gaps in the individual price series. The average length of a product spell we observe is 122, i.e., longer than 10 years out of the 22 years of the sample period.

Another advantage–for example relative to unit value data–is that in the case that any feature of the product such as size or quality is changed, firms are also asked to report the hypothetical price of the exact same product in the previous month. We refer to quality-adjusted prices as P_{jim}. This allows us to compute quality-adjusted firm level inflation rates at monthly frequency:

$$\Delta p_{jim} = \log \tilde{P}_{jim} - \log P_{ji,m-1}. \tag{A.1}$$

One drawback of the data is that we do not observe within-firm quantities or product weights. The mean firm reports prices for 7 products in total over the sample period. As firms are asked to report prices for a “representative” sample of products, we assign uniform weights across products within a firm when computing firm-level inflation rates

$$\Delta p_{im} = \frac{1}{J} \sum_{j} \Delta p_{jim}. \tag{A.2}$$

[25]When applying sales filter “B” of Nakamura and Steinsson (2008) to the raw data, 0.3% of price quotes are identified as sales (see Dedola et al. (2019)). Temporary sales are not a prominent feature in the Danish PPI.

Subsequently, we generate firm-level price deflators as

$$P_{im} = \prod_{1}^{m} P_{i,m-1}(1 + \Delta p_{im}) \tag{A.3}$$

and then average over all monthly observations in a calendar year. Finally, we index the resulting price levels at the annual frequency P_{it} to a base year (2015). For firms not reporting prices in 2015 but in other years, we let their deflators in the first year of observation take the average value of the deflators of all other firms. However, as all our regressions include firm fixed effects, the *level* of prices of each firm will not be important.

A.1.2 Industry-level output deflators

Much of the literature on production function and markup estimation deflates nominal sales using industry-level price deflators, P_{kt}. We do this as well in robustness checks and in order to verify (the mean of) our firm-level prices.

Sector-level PPIs are published by Denmark Statistics in different datasets, two of which are relevant for us: PRIS4215 and PRIS4015. The advantage of the former is that it contains consistent time series for the entire time span, starting in 2000. However, it has the disadvantage that it groups together certain industries even at the two-digit NACE levels. For example, basic metals (NACE 24) and fabricated metal products (NACE 25) are collected in one index (CH). In the second dataset–PRIS4015–all series are published at the 2-digit NACE level, but start only in 2005. Wherever a series is available at the 2-digit level, we use data from PRIS4215. If not, we use the 2-digit series starting in 2005 from PRIS4015 and link it to the slightly more aggregated series from PRIS4215 in the years prior to 2005. For two sectors, namely food manufacturing (NACE 10) and manufacturing on other non-metallic mineral products (NACE 23), the PRIS4215 dataset allows us to match price series at a slightly finer level of granularity. The complete list of NACE codes we match to price series from any of the PPI datasets is contained in Table A.1.

A.1.3 Verification of firm-level output deflators

As the estimation approach discussed in Section 3 hinges on a reliable measure of real output, we show that our firm-level output deflators on average accurately reflect industry price dynamics which are published by Statistics Denmark.

We first show that annual inflation rates of our own firm-level output deflators (Δp_{it}) and the

Table A.1: Sources for industry-level output deflators

NACE	Industry	Time	PPI dataset	Series
10.1	Food: Meat	2000-	PRIS4215	10001 Meat
10.2	Food: Fish	2000-	PRIS4215	10002 Fish
10.3	Food: Fruit and vegetables	2000-	PRIS4215	10005 Other food products
10.4	Food: Oils and fats	2000-	PRIS4215	10005 Other food products
10.5	Food: Dairy products	2000-	PRIS4215	10003 Dairy products
10.6	Food: Grain mill, starch prod.	2000	PRIS4215	10004 Grain mill a. bakery
10.7	Food: Bakery	2000-	PRIS4215	10004 Grain mill a. bakery
10.8	Food: Other	2000-	PRIS4215	10005 Other food products
10.9	Food: Prepared animal feeds	2000-	PRIS4215	10005 Other food products
13	Textiles	-2005	PRIS4215	CB Textiles and leather
		2005-	PRIS4015	13 Textiles
16	Wood	-2005	PRIS4215	CC Wood, paper a. print.
		2005-	PRIS4015	16 Wood
17	Paper	-2005	PRIS4215	CC Wood, paper a. print.
		2005-	PRIS4015	17 Paper
20	Chemicals	2000-	PRIS4215	CE Chemicals
22	Rubber and plastic	2000-	PRIS4215	22000 Rubber and plastic
23.1	NMM:* Glass	2000-	PRIS4215	23001 Glass and ceramic
23.2	NMM: Refractoryproducts	2000-	PRIS4215	23001 Glass and ceramic
23.3	NMM: Clay building materials	2000-	PRIS4215	23001 Glass and ceramic
23.4	NMM: Porcelain, ceramic prod.	2000-	PRIS4215	23001 Glass and ceramic
23.5	NMM: Cement, lime and plaster	2000-	PRIS4215	23002 Concrete and bricks
23.6	NMM: Articles of concrete	2000-	PRIS4215	23002 Concrete and bricks
23.7	NMM: Stone	2000-	PRIS4215	23002 Concrete and bricks
23.9	NMM: Abrasive and other	2000-	PRIS4215	23002 Concrete and bricks
24	Basic metals	-2005	PRIS4215	CH Basic and fabr. metals
		2005-	PRIS4015	24 Basic metals
25	Fabricated metal products	-2005	PRIS4215	CH Basic and fabr. metals
		2005-	PRIS4015	25 Fabr. metals
26	Computers and electronics	2000-	PRIS4215	CI: Electronic components
27	Electrical equipment	2000-	PRIS4215	CJ Electrical equipment
28	Machinery and equipment	2000-	PRIS4215	CK Machinery
29	Motor vehicles	2000-	PRIS4215	29000 Motor vehicles
31	Furniture	-2005	PRIS4215	CM Furniture and other
		2005-	PRIS4015	31 Furniture
32	Other manufacturing	-2005	PRIS4215	CM Furniture and other
		2005-	PRIS4015	32 Other manuf.

Notes: To construct P_{kt}, we use industry-level PPIs published by Denmark Statistics in two datasets (PRIS4215 and PRIS4015), with different time and industry coverage. See Section A.1.2 for details.

sector-level equivalent which are publicly available (Δp_{kt}) have a high degree of correlation. In principle, the latter are a weighted average of the former, but our main challenge is that we do not observe any weights of the micro price data used to construct the aggregate PPI and sector-level subindices. Figure A.1(a) shows a binned scatter plot and the fit of a linear regression (blue solid line). The estimated coefficient is 0.81 (with a standard error of 0.03). Our PPI micro data thus accurately reflects price changes as measured with industry price deflators, on average.

Figure A.1: Firm- vs. industry-level deflators

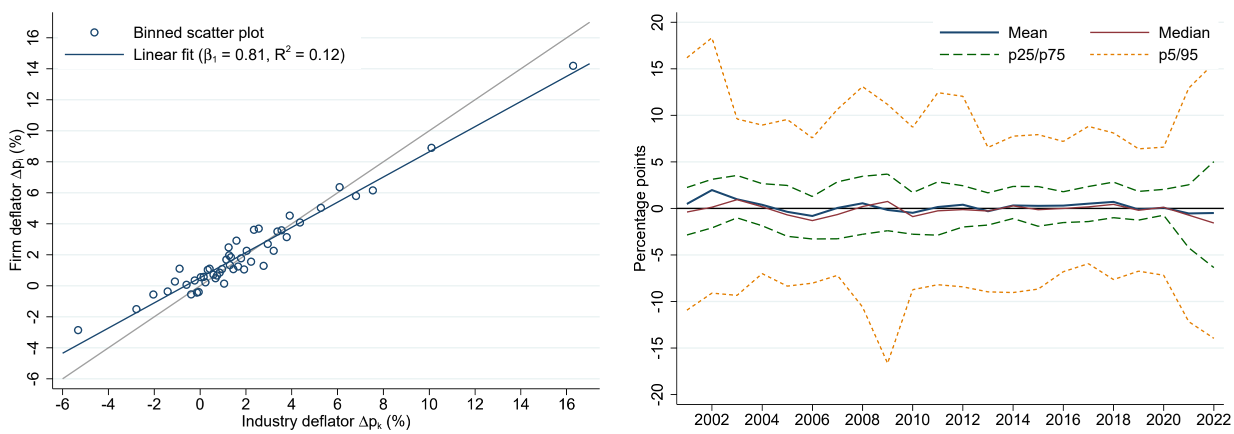

(a) Δp_{it} vs. Δp_{kt}

(b) Distribution of $\Delta p_{it} - \Delta p_{kt}$ by year

Notes: Relationship between annual inflation rates at the firm level (using the PPI micro data, as explained in Section A.1.1) and the level of the associated industry (using publicly available PPI indices, as explained in Section A.1.2).

While industry-price deflators are informative, there is a large degree of price dispersion, even within sector-years. To show this, we compute, the percentage point difference between the changes of firm and industry deflators. In Figure A.1(b), we show moments of this distribution for each year. There are two conclusions that support the validity of our subsequent analysis: First, there is no systematic bias that would lead to diverting time trends between the average firm and industry deflators (and thus in markups estimated with either firm- or industry-level price deflators). Second, it becomes clear that there are large deviations of firm-level price changes relative to the weighted industry average. As is shown by the dashed lines, around half the firms in most years increase or decrease their prices more than 2pp more than the average industry price. The fact that output deflators behave heterogeneously across firms highlights the importance of using firm-level output deflators in our production function estimation.

A.1.4 Factor input deflators

As is standard in the literature, we deflate all factor inputs before they enter the production function. In the following, we discuss the construction of the deflators for each input factor in turn. The price series used for this are uniform within industries, assuming perfect competition in input markets.

Material We compute material input prices at the sector level from annual input-output tables at the 2-digit NACE level. Let $P^M_{klt} M_{klt}$ be the nominal input values of sector k purchased from sector l. Input-output tables are published in current values, as above, and in previous-year prices, i.e., $P^M_{kl,t-1} M_{klt}$. We compute sector-level input inflation rates as the weighted average of the log difference of the two, whereas the weights are the lagged nominal input shares:

$$\Delta p^M_{kt} = \sum_l \left(\frac{P^M_{kl,t-1} M_{kl,t-1}}{\sum_l P^M_{kl,t-1} M_{kl,t-1}} \right) \left(\log P^M_{klt} M_{klt} - \log P^M_{kl,t-1} M_{klt} \right). \tag{A.4}$$

Finally, we generate a material input price deflator as the accumulated product of Δp^M_{kt}.

Labor Our production function features real labor cost as labor input to better reflect quality differences in labor than with simple full-time equivalents (which we use in robustness checks). To deflate nominal labor cost, we sum over labor expenses reported in the revenue statement over all firms (not just the 800 firms in our sample) and divide by the sum of full-time equivalents of these firms.

Capital Denmark Statistics publishes time series on the capital stock (dataset NAHK) on total fixed assets (as well as subcomponents such as buildings, structures, equipment and intellectual property products. We use the ratio of nominal and real total fixed assets to deflate firms' capital stock.

A.2 Firm sample

Figure A.2: Industry composition

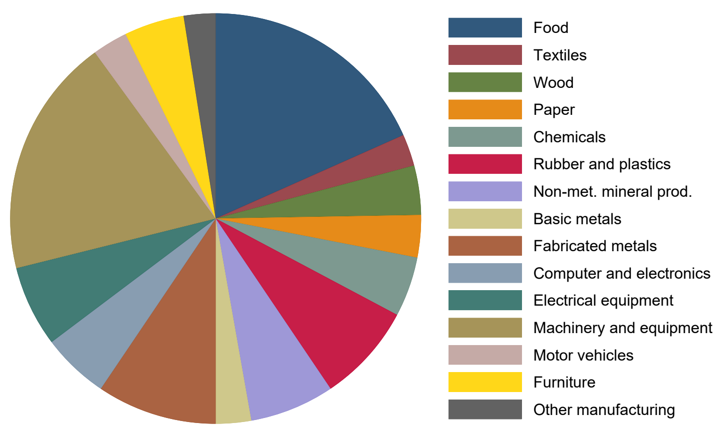

Notes: Distribution of 2-digit NACE industry codes in the sample of matched firms and price quotes.

Figure A.3: Age distribution

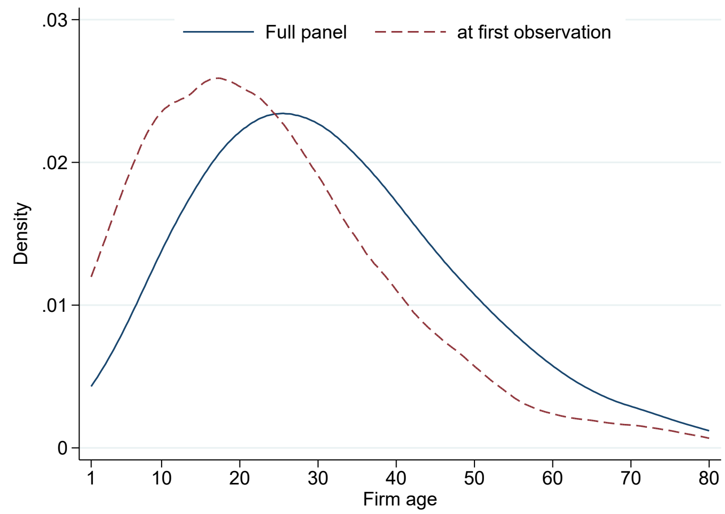

Notes: Distribution of firm age in the data, displayed with an Epanechnikov kernel density with a bandwidth of 5 years. While the blue line uses the full panel, the dashed red line uses only the first observation of the 800 firms, i.e. the observation with their youngest age.

Table A.2: Firm size distribution in Denmark and the rest of Europe

	Gross output			Employment		
Size bin by employment:	Denmark, benchmark	Europe, Orbis	Europe, Eurostat	Denmark, benchmark	Europe, Orbis	Europe, Eurostat
1-19	0.12	0.07	0.10	0.14	0.10	0.22
20-249	0.34	0.30	0.29	0.41	0.40	0.37
250+	0.54	0.63	0.60	0.46	0.50	0.41

Notes: Kalemli-Ozcan et al. (2022) compute the distribution of gross output and employment of manufacturing firms in the Oribs Global databaset across three size bins, comparing a total of 20 countries: Austria, Belgium, Czechia, Estonia, Finland, France, Germany, Greece, Hungary, Italy, Latvia, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden and the United Kingdom. The data only covers the year 2006. We compute a weighted average across all twenty countries using manufacturing gross value added and manufacturing employment as weights. Kalemli-Ozcan et al. (2022) compare their data to official statistics they take from Eurostat. To characterize the Danish manufacturing sector, we compute the same statistics using only the 2006 register (of the full sample of manufacturing firms, not our selected regression sample). The table documents that the size distribution of Danish manufacturing firms is similar to that in other European economies in terms of the distribution of employment. In terms of gross output, large firms (making up 54% of gross output) are somewhat underrepresented, but this is especially compared to large economies such as Germany, France and the UK (70, 72 and 74%, respectively) and more similar to other small European economies.

Figure A.4: Joint age and size distribution of sample firms relative to benchmark sample

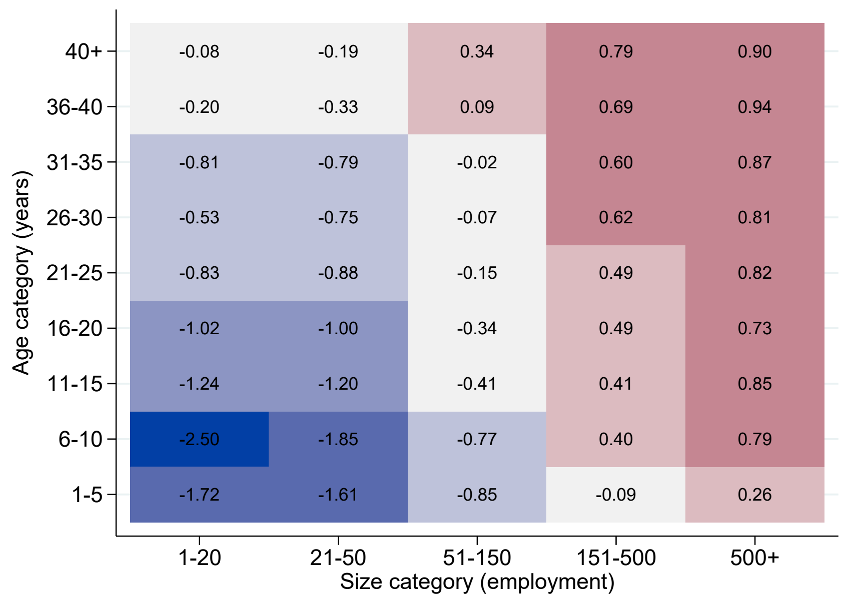

Notes: We first compute the share of firm-year observations for each size-age cell in our sample and compute log differences to the equivalent share in the benchmark sample. The benchmark sample consists of all manufacturing firms in Denmark that have 50 employees at least once. Positive values (shaded red) indicate that firms of the respective size-age combination are overrepresented in our sample.

Figure A.5: Declining business dynamism and the increase of average firm age

(a) Entry and exit

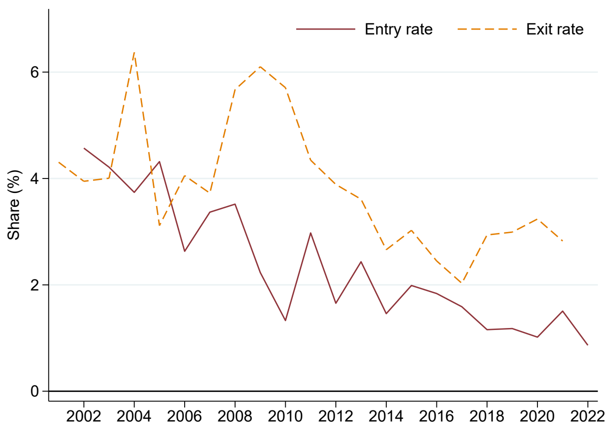

(b) Contribution to average age

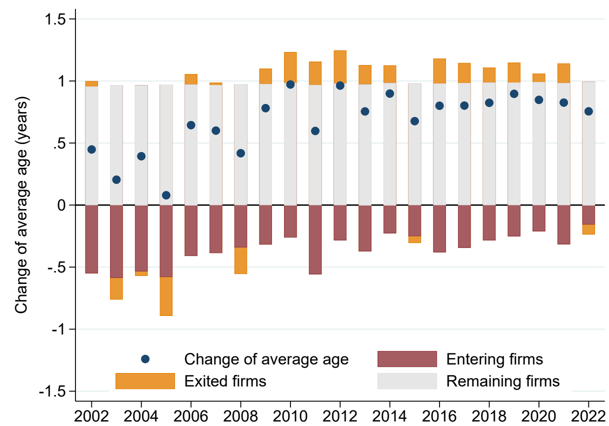

Notes: In panel (a), we show the share of firms that enter the benchmark sample in a given year (red) and leave the benchmark sample in the subsequent (yellow), both as a share of all firms. We use the broader benchmark sample to compute these entry and exit rates because it is larger and our main sample is constrained by sampling in the PPI and accounting statistics surveys, i.e., not by firms actually entering and exiting the market, but to some extent by firms entering and exiting the survey. The benchmark sample uses all firm observations of manufacturing firms with at least 10 employees and reaching at least 50 employees at least once over the course of the sample period. Entry and exit are thus defined for the years where the firm crosses the 10-employee threshold, i.e., it does not by definition have age 1. In panel (b) we show how entry and exit of firms affect the average age of firms in the benchmark sample. The blue bullets show the change of the average age across all firms, i.e., first differences of the dashed blue line in Figure 1. The contribution by entering firms, depicted in dark red, is calculated as $\phi_t^{\text{enter}}(\bar{a}_t^{\text{enter}} - \bar{a}_{t-1})$, where ϕ^{enter} is the share of new firms. If the average age of new firms – which is usually 1 – is below the average age of continuing firms, i.e., if $\bar{a}_t^{\text{enter}} - \bar{a}_{t-1} < 0$, they contribute negatively. The contribution by exiting firms (depicted in orange bars) is -$\phi_{t-1}^{\text{exit}}(\bar{a}_{t-1}^{\text{exit}} - \bar{a}_{t-1}^{\text{remain}})$. If the average exiting firm is younger than the average remaining firm, the contribution by exiters on average age is positive. Remaining firms are depicted in light grey bars. Their contribution is close to one as almost all firms remain and become one year older.

B Details on estimated production functions and markups

Table B.1: Input shares and output elasticities by industry

Industry	N	Firms	α^M	α^{M*}	$\hat{\theta}^M$	$\hat{\theta}^L$	$\hat{\theta}^K$	$100 \cdot \log \hat{\mu}$
Food	1,710	146	0.62	0.62	0.71	0.28	0.06	14.44
Textiles	251	21	0.56	0.56	0.61	0.17	0.10	6.95
Wood	366	31	0.52	0.52	0.59	0.44	0.03	12.48
Paper	314	26	0.50	0.50	0.67	0.27	0.02	30.70
Chemicals	497	38	0.52	0.53	0.54	0.39	0.12	-2.33
Rubber and plastics	795	61	0.48	0.49	0.52	0.48	0.05	3.88
Non-met. mineral prod.	702	54	0.40	0.41	0.46	0.48	0.06	10.51
Basic metals	250	23	0.61	0.61	0.57	0.27	0.09	-6.00
Fabricated metals	810	74	0.47	0.48	0.55	0.43	0.09	16.41
Computer and electronics	531	43	0.43	0.44	0.47	0.57	-0.00	10.22
Electrical equipment	576	50	0.51	0.52	0.50	0.47	0.00	-5.97
Machinery and equipment	1,904	151	0.50	0.50	0.52	0.48	0.03	4.34
Motor vehicles	288	22	0.52	0.54	0.60	0.39	0.05	11.85
Furniture	419	37	0.46	0.46	0.78	0.15	0.11	51.82
Other manufacturing	287	21	0.42	0.43	0.61	0.46	0.03	36.91
All	9,700	798	0.51	0.51	0.58	0.40	0.05	11.10

Notes: Before computing means, all variables are winsorized at the 5th and 95th percentile by industry and year. *Our baseline measure of the material input share is corrected for the change in intermediate goods inventories since the end of the previous year. In the fourth column, we show the raw input share not corrected for changes in inventories in order to show that this does not systematically affect the level of the nominal input share.

Table B.2: Alternative markup estimates

	Overall	Emp. wt.	Mean 2001-02	2011-12	2021-22
Main markup estimate	11.10	11.93	4.99	10.03	15.61
Industry output deflator	5.79	6.41	-0.09	4.66	1.38
De Ridder et al. first stage	2.82	3.04	-4.02	1.54	7.81
Employment as labor input	4.98	10.47	-5.77	3.92	9.97
Uncorrected material input	12.41	13.79	7.59	11.30	14.27
Main, not winsorized	11.24	12.82	2.85	10.46	16.27
	Median	**SD**	**Skew**	**Kurt**	**Corr w/ main**
Main markup estimate	9.80	23.81	0.23	2.80	1.00
Industry output deflator	7.27	20.37	-0.57	3.37	0.68
De Ridder et al. first stage	4.52	23.61	-0.49	3.14	0.66
Employment as labor input	7.47	40.59	-0.21	2.64	0.68
Uncorrected material input	10.40	25.26	0.44	3.08	0.84
Main, not winsorized	10.05	30.82	-1.24	18.73	0.91

Notes: We report moments of the distribution of markups, which we estimate as described in Section 3, then winsorize at the 5th and 95th percentile (unless mentioned otherwise) and express in percent over marginal cost, i.e., $100 \cdot \log \hat{\mu}_{it}$. Deviations relative to the main markup estimate are as follows: The second row uses industry output deflators instead of firm-specific deflator series. The third row adds two two variables to the first stage in the production function estimation, namely the market share of the firm and the log price, to control for the markup itself. The fourth row uses log employment in full-time equivalents as an input into the production function instead of deflated labor cost; the fifth control neither material inputs nor outputs for the change in intermediate and final goods inventories.

Figure B.1: Pairwise correlations between main markup estimates and alternatives

(a) Output deflator and first stage

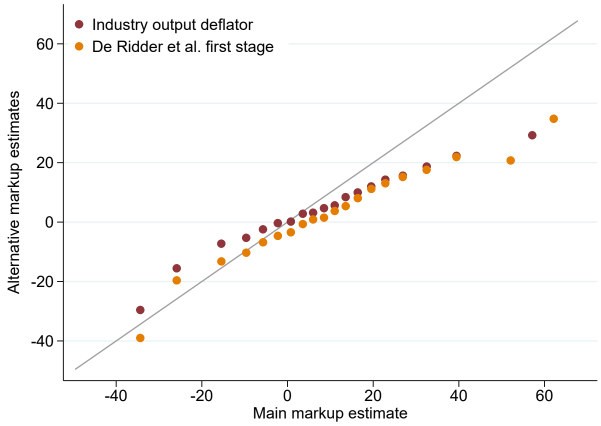

(b) Production function inputs

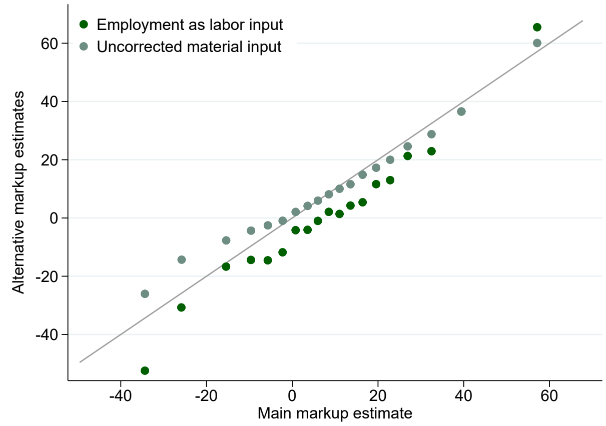

Notes: Binned scatter plots of the main markup estimates (on the x-axis) against two alternative estimates at a time. Alternatives are described in Section 5.

Figure B.2: Time trend of average markup with alternative markup estimates

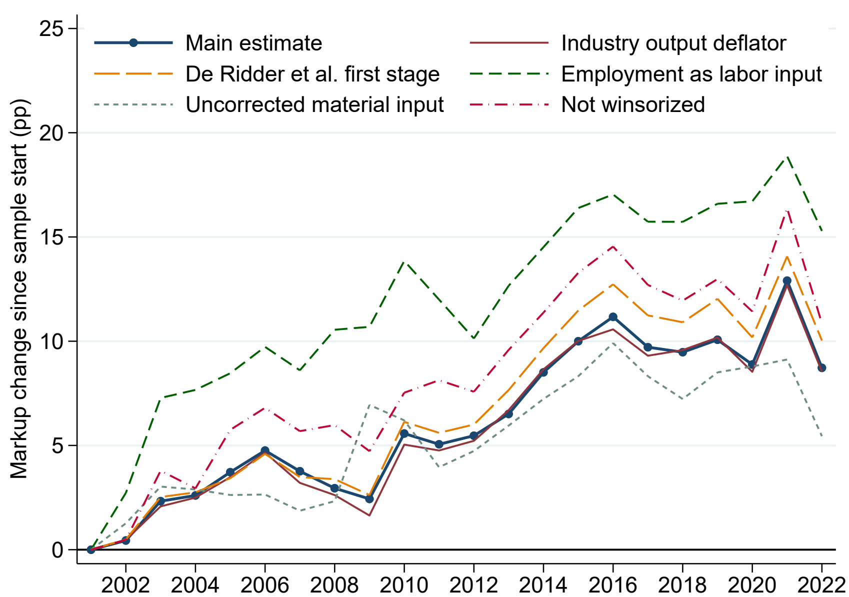

Notes: We report the unweighted average of the distribution of markups, which we estimate as described in Section 3, then winsorize at the 5th and 95th percentile (unless mentioned otherwise) and express in percent over marginal cost, i.e., $100 \cdot \log \hat{\mu}_{it}$. Besides our main estimates shown in blue dots, we compare the indexed average to a number of alternatives, the specifications to which are described in Section 5. Series are indexed to 0 by subtracting the 2001 values.

Figure B.3: Business cycle cyclicality of average markup

(a) Main estimate

(b) Alternative estimates

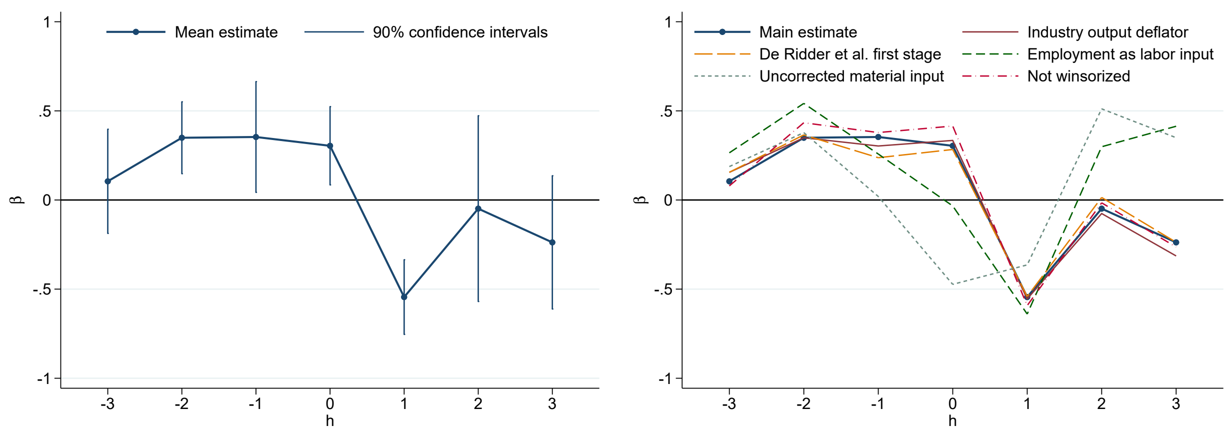

Notes: We report estimates $\hat{\beta}_{1,h}$ of the following regression $100 \cdot \log \bar{\mu}_{t+h} = \beta_{0,h} + \beta_{1,h}\hat{y}_t + \gamma \text{year}_t + u_t$, the dynamic covariances of the output gap $\hat{y}$ (which we hp-filter from annual GDP with a smoothing parameter of 6.25) after accounting for a linear time trend γ. We use Newey-West standard errors to compute confidence intervals. Positive coefficients indicate procyclical markups. With our main markup estimates, we find that the average markup is procyclical and slightly leading the business cycle, as indicated by positive values for $h = \{-2, -1, 0\}$. In panel (b), we confirm this for most of our alternative estimates of the mean markup. The only alternative for which markups are countercyclical is when material input and output are not adjusted for the change in inventories. Markups are *more procyclical* when proper inventory adjustments are made.

C Details on identification restrictions

C.1 Cohort trend restriction

As outlined in section 4.1, the first restriction we impose is that $E(c_i N_i \chi_i) = 0$, where c_i is a firm's birth cohort, and N_i is the number of observations for that firm in the sample. We implement this constraint following Deaton (2019). We derive an auxiliary regression that identifies the parameters in the regression:

$$y_{it} = f(a_{it}) + \sum_t \sum_k \tau_{kt} m_{it} + \sum_i \chi_i d_{it}, \tag{B.1}$$

This is a restatement of decomposition (6) with notation making the included dummy variables explicit. m_{it} are sector-year dummies, d_{it} are firm dummies and χ_i are the associated coefficients (fixed effects). We impose the normalization that

$$\sum_i N_i \chi_i = 0, \tag{B.2}$$

i.e. the firm fixed effects sum to zero. Moreover, we impose the restriction that

$$\sum_i c_i N_i \chi_i = 0, \tag{B.3}$$

i.e. the firm effects are orthogonal to any linear cohort trend.

We can express restriction (B.2) in terms of the fixed effect of two (arbitrary) reference firms 1 and 2 from two different cohorts.

$$\chi_1 = -\frac{N_2}{N_1}\chi_2 - \frac{1}{N_1}\left(\sum_{i \neq 1,2} N_i \chi_i\right) \tag{B.4}$$

$$\chi_2 = -\frac{N_1}{N_2}\chi_1 - \frac{1}{N_2}\left(\sum_{i \neq 1,2} N_i \chi_i\right) \tag{B.5}$$

We do the same for restriction (B.3):

$$\chi_1 = -\frac{c_2}{c_1}\frac{N_2}{N_1}\chi_2 - \frac{1}{c_1 N_1}\left(\sum_{i\neq 1,2} c_i N_i \chi_i\right) \tag{B.6}$$

$$\chi_2 = -\frac{c_1}{c_2}\frac{N_1}{N_2}\chi_1 - \frac{1}{c_2 N_2}\left(\sum_{i\neq 1,2} c_i N_i \chi_i\right) \tag{B.7}$$

To simplify notation, we require that $N_1 = N_2$ and use two reference firms that are present for the whole sample period. We can then use (B.7) and (B.4) to express χ_1 and (B.6) and (B.5) to express χ_2 as a function of the other fixed effects and the imposed restrictions:

$$\chi_1 = \sum_{i\neq 1,2} \chi_i \left(\frac{c_i - c_2}{c_2 - c_1}\right)\frac{N_i}{N_1} \tag{B.8}$$

$$\chi_2 = -\sum_{i\neq 1,2} \chi_i \left(\frac{c_i - c_1}{c_2 - c_1}\right)\frac{N_i}{N_2} \tag{B.9}$$

Finally, (B.8) and (B.9) can be plugged into (B.1) to obtain

$$y_{it} = f(a_{it}) + \sum_t \sum_k \tau_{kt} m_{it} + \sum_{i\neq 1,2} \chi_i \left(d_i + d_1\left(\frac{c_i - c_2}{c_2 - c_1}\right)\frac{N_i}{N_1} - d_2\left(\frac{c_i - c_1}{c_2 - c_1}\right)\frac{N_i}{N_2}\right). \tag{B.10}$$

Consequently, we can run the following regression with properly transformed firm dummies $\tilde{d}_i$ to obtain estimates of $f(a)$, τ_{kt} and χ_i with the constraint imposed:

$$y_{it} = f(a_{it}) + \sum_t \sum_k \tau_{kt} m_{it} + \sum_{i\neq 1,2} \chi_i \tilde{d}_{it}, \tag{B.11}$$

with

$$\tilde{d}_{it} = d_{it} + d_{1t}\left(\frac{c_i - c_2}{c_2 - c_1}\right)\frac{N_i}{N_1} - d_{2t}\left(\frac{c_i - c_1}{c_2 - c_1}\right)\frac{N_i}{N_2} \tag{B.12}$$

The estimates for the final two fixed effects χ_1 and χ_2 follow from (B.8) and (B.9).

C.2 Age polynomial restriction

We impose that the age polynomial fulfills $f'(\bar{a}) = 0$ for some age $\bar{a}$. This pins down the linear coefficient of $f(a)$ as a function of the higher order coefficients. For fourth degree polynomial, the derivative is given by:

$$f'(a) = \alpha_1 + 2\alpha_2 a + 3\alpha_3 a^2 + 4\alpha_4 a^3 \tag{B.13}$$

and with the constraint imposed:

$$\alpha_1 = -2\alpha_2\bar{a} - 3\alpha_3\bar{a}^2 - 4\alpha_4\bar{a}^3 \tag{B.14}$$

Plugging back into the estimation equation:

$$y_{it} = \alpha_2(a_{it}^2 - 2\bar{a}) + \alpha_3(a_{it}^3 - 3\bar{a}^2) + \alpha_4(a_{it}^4 - 4\bar{a}^3) + \tau_{s(i),t} + \chi_i \tag{B.15}$$

We can thus identify α_2, α_3 and α_4 from a regression with properly transformed age terms $\tilde{a}_{it}^2 = (a_{it}^2 - 2\bar{a})$, $\tilde{a}_{it}^3 = (a_{it}^3 - 3\bar{a}^2)$ and $\tilde{a}_{it}^4 = (a_{it}^4 - 4\bar{a}^3)$. The linear coefficient α_1 follows from the constraint (B.14).

D Supplementary Results and Robustness Checks

Figure D.1: Cohort trends in estimated firm effects

(a) with cohort trend restriction

(b) with age polynomial restriction

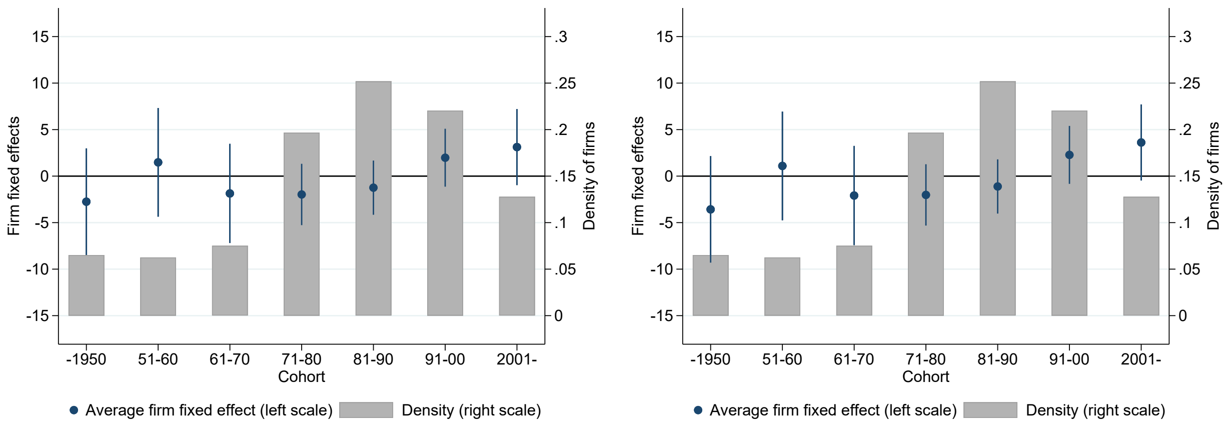

Notes: In Section 4.2 we have identified age patterns of markups under two separate restrictions, namely the restriction that the linear cohort trend is equal to zero, and that the age polynomial has a slope of 0 at a certain age. The figure plots the identified firm fixed effects χ_i by firm birth cohort grouped by decade (and a larger bin for cohorts before 1950 to guarantee a large enough number of observations) for both identification restrictions. In the case of the cohort trend restriction, a linear regression of χ_i on c_i yields a precisely estimated 0, by construction. In the case of the age polynomial restriction, the coefficient is 0.018 with a p-value of 0.084, giving some evidence that under this restriction, later-born cohorts have higher markups. The fact that the restriction placed in panel (a) does not hold under the restriction placed in panel (b) indicates that our two identification restrictions are, at least to some extent, complementary and do not mutually depend on one another.

Figure D.2: Contribution of age composition to average markup: Young vs. old firms

(a) with cohort trend restriction

(b) with age polynomial restriction

Notes: Having identified the age effects in Equation (6) under two different identification restrictions, we take fitted values of the age component for each firm and subsequently aggregate them over young firms (age 1-20) and older firms (20 years and older) by year.

Figure D.3: Contribution of time-sector effects to average markup

(a) Cohort trend restriction, effects

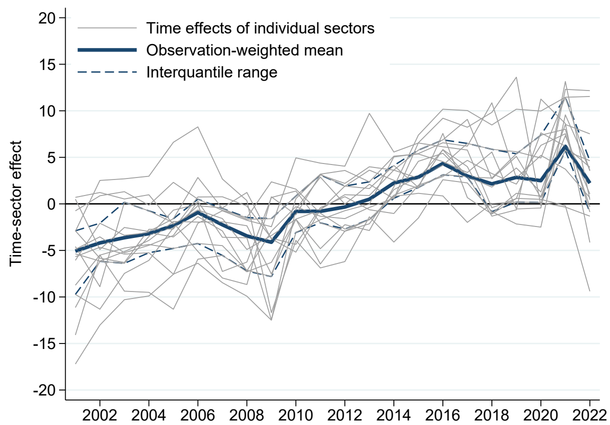

(b) Age polynomial restriction, effects

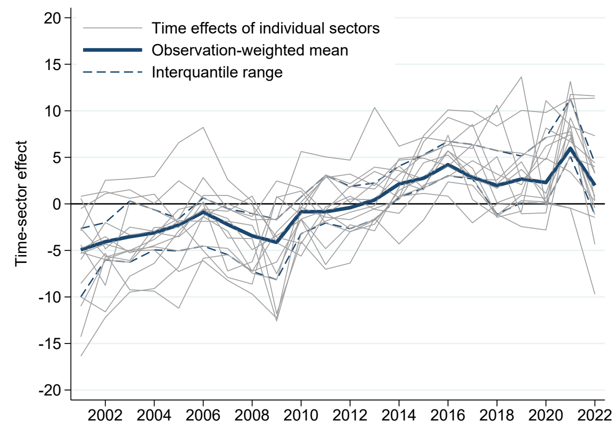

(c) Cohort trend restriction, contribution

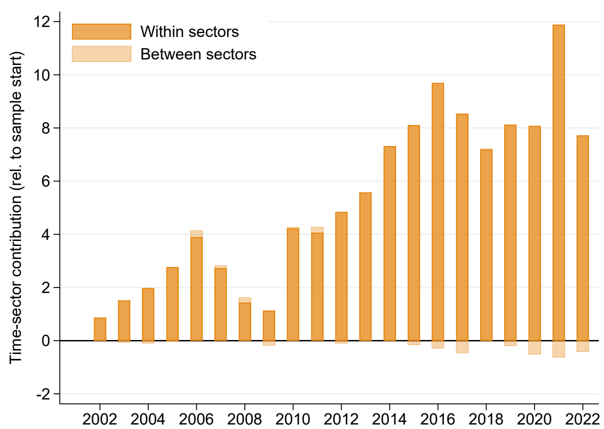

(d) Age polynomial restriction, contribution

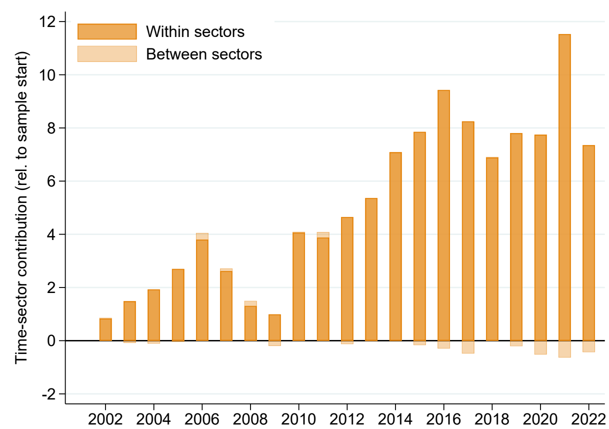

Notes: We have identified sector-time effects from age and firm effects in Equation (6) under two different identification restrictions. In the top two panels, we show each of the sector-time fixed effects, along with the mean and interquantile range across the 15 sectors. To construct panels at the bottom, we decompose the sample-average sector-time fixed effect into within-sector and between-sector components relative to the sample start, following $\bar{\tau}_t = \sum_k s_{k,01}(\tau_{k,t} - \tau_{k,01}) + \sum_k (s_{k,t} - s_{k,01})\tau_{k,t}$, where s is the share of firms in a specific sector k.

Table D.1: Robustness of age patterns to input variables in production function estimation

	(1)	(2)	(3)	(4)
	Employment instead of deflated labor cost		No inventory correction	
Identification:				
Cohort trend restriction	✓		✓	
Age polynomial restriction		✓		✓
Markups				
Change from age 1 to 20	4.18***	9.72***	6.42***	7.01***
	(1.50)	(1.92)	(0.81)	(1.46)
Age 20 to 40	-4.10***	-0.71***	-1.39	-0.77***
	(0.94)	(0.21)	(0.97)	(0.16)
Marginal cost				
Change from age 1 to 20	-10.54***	-10.99***	-13.27***	-8.24***
	(1.71)	(2.10)	(1.33)	(1.50)
Age 20 to 40	-0.52	1.76***	-3.49***	1.81***
	(0.73)	(0.18)	(0.43)	(0.15)
Observations	9,588	9,588	9,624	9,624
Firms	793	793	795	795

Notes: Percentage point and percent changes of markups and marginal cost with variations in markup estimation. Columns (1) and (2) use (log) employment in full-time equivalents instead of deflated labor cost as labor input l_{it} and columns (3) and (4) use material input m_{it} (output y_{it}) without adjusting them for the change of intermediate (final) goods inventories. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.2: Robustness of age patterns to outliers

	(1)	(2)	(3)	(4)
	No winsorizing		Winsorizing at 10th/90th percentiles	
Identification:				
Cohort trend restriction	✓		✓	
Age polynomial restriction		✓		✓
Markups				
Change from age 1 to 20	6.16***	9.58**	6.57***	6.11***
	(1.89)	(2.85)	(0.94)	(1.43)
Age 20 to 40	-4.24***	-0.64***	-0.39	-0.87***
	(1.36)	(0.20)	(0.83)	(0.14)
Marginal cost				
Change from age 1 to 20	-12.33***	-10.55***	-11.89***	-6.58***
	(2.10)	(2.38)	(1.38)	(1.43)
Age 20 to 40	-0.31	1.55***	-3.95***	1.64***
	(0.76)	(0.24)	(0.33)	(0.18)
Observations	9,608	9,608	9,700	9,700
Firms	795	795	798	798

Notes: Percentage point and percent changes of markups and marginal cost with variations in the percentiles at which we winsorize the estimated markups. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.3: Robustness of age patterns to control variables

	(1)	(2)	(3)	(4)
	Control for firm size		Control for additional firm characteristics	
Identification:				
Cohort trend restriction	✓		✓	
Age polynomial restriction		✓		✓
Markups				
Change from age 1 to 20	10.47***	9.51***	9.85***	8.75***
	(0.91)	(1.39)	(0.87)	(1.33)
Age 20 to 40	-0.05	-1.06***	0.17	-0.98***
	(0.70)	(0.11)	(0.70)	(0.10)
Marginal cost				
Change from age 1 to 20	-11.50***	-5.79***	-11.70***	-5.97***
	(1.59)	(1.54)	(1.56)	(1.54)
Age 20 to 40	-4.60***	1.41***	-4.59***	1.44***
	(0.34)	(0.19)	(0.37)	(0.18)
Observations	9,699	9,699	9,697	9,697
Firms	798	798	798	798

Notes: Percentage point and percent changes of markups and marginal cost when estimating age patterns with additional control variables. Columns (1) and (2) include a 4th-order polynomial of log employment and log real output as we use it in the production function estimation. Columns (3) and (4) additionally include the log market share in the 2-digit NACE sector in Denmark, the leverage ratio and well as the ratios of cash to sales and cash to total assets. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.4: Robustness of age patterns to reference age in age polynomial restriction

	(1)	(2)	(3)	(4)
	$\bar{a} = 15$	$\bar{a} = 20$	$\bar{a} = 30$	$\bar{a} = 35$
Identification:				
Cohort trend restriction				
Age polynomial restriction	✓	✓	✓	✓
Markup				
Change from age 1 to 20	3.47***	6.00***	9.05***	9.73***
	(0.88)	(1.44)	(1.97)	(1.99)
Age 20 to 40	-5.50***	-2.83***	0.37***	1.09***
	(1.06)	(0.49)	(0.11)	(0.18)
Marginal cost				
Change from age 1 to 20	-3.53***	-6.56***	-11.20***	-12.93***
	(70)	(1.21)	(1.83)	(1.97)
Age 20 to 40	7.81***	4.62***	-0.27***	-2.08***
	(1.12)	(0.58)	(0.8)	(0.23)
Observations	9,700	9,700	9,700	9,700
Firms	798	798	798	798

Notes: Percentage point and percent changes of markups and marginal cost under the age polynomial restriction to identify age patterns, with variations in the reference age $\bar{a}$. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table D.5: Robustness of age patterns to the order of the age polynomial

	(1)	(2)	(3)	(4)
Identification:	3rd-order polynomial		5th-order polynomial	
Cohort trend restriction	✓		✓	
Age polynomial restriction		✓		✓
Markups				
Change from age 1 to 20	3.83***	2.53***	8.26***	9.14***
	(0.85)	(0.63)	(1.77)	(3.07)
Age 20 to 40	0.65***	-0.72***	-1.64	-0.72*
	(0.20)	(0.15)	(1.20)	(0.42)
Marginal cost				
Change from age 1 to 20	-12.14***	-5.86***	-16.56***	-12.81***
	(1.12)	(0.73)	(1.10)	(1.41)
Age 20 to 40	-4.78***	1.84***	-2.54***	1.41***
	(0.37)	(0.19)	(0.61)	(0.34)
Observations	9,700	9,700	9,700	9,700
Firms	798	798	798	798

Notes: Percentage point and percent changes of markups and marginal cost with variations in the order of the polynomial to approximate $f(a_{it})$. See Table 2 for details. Driscoll-Kraay standard errors in brackets; significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.