Does the Phillips Curve Lie Down as we Age?*

Chadwick Curtis[†]
University of Richmond

Julio Garín[‡]
Claremont McKenna College

Robert Lester[§]
Colby College

This Version: June 20, 2025

Abstract

We study the qualitative consequences of accounting for an unexplored source of heterogeneity in a model with nominal rigidities. Using micro-level data, we present evidence that older individuals are less willing to substitute across varieties of goods. In particular, we estimate the elasticity of substitution for different age groups and find that the youngest cohort (aged 25–34) exhibits a higher elasticity of substitution compared to the oldest group (65+). We incorporate this empirical finding in a Rotemberg model of price adjustment and show that the age distribution affects the slope of the Phillips curve. In a quantitative exercise, we find that this channel accounts for approximately 4.5 percent of the changes observed in its empirical counterpart. Taken together, our results highlight a new channel by which the age distribution of a population could impact both the transmission mechanism and efficacy of monetary policy.

JEL Classification: E21; E40; E52; J11.

Keywords: Consumption; Demographics; Aging; Monetary Policy.

*Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

†E-mail address:ccurtis2@richmond.edu.

‡E-mail address: jgarin@gmail.com.

§E-mail address: rblester@colby.edu.

1 Introduction

In recent years, economic research has advanced our understanding of the effects of household heterogeneity on the transmission of monetary policy. Several prominent papers show that wealth differences capture important channels that are otherwise absent in a representative household environment. A related aspect of this literature emphasizes the role of age. Both areas explore mechanisms in which wealth accumulation - or the lack thereof - and liquid assets play essential roles. In this paper, we show that the age distribution of a population affects monetary policy through a novel channel: differences in the substitution elasticities between the old and the young. We empirically document this behavior and incorporate it into an otherwise standard model with nominal rigidities.

Our contribution has two dimensions. Empirically, we present evidence that the elasticity of substitution is lower for older age groups compared to younger ones. On the theoretical front, we extend a model of monopolistic competition with incomplete nominal price adjustment to include consumer heterogeneity. We then show that incorporating heterogeneity in the elasticity of substitution in the model flattens the Phillips curve. Furthermore, in a quantitative exercise, we evaluate the extent to which this channel accounts for the observed changes in the slope of the curve.

Figure 1 shows the evolution of US expenditure shares across different age groups. There is a notable decline in expenditure shares for the youngest age groups and a significant rise for the older groups, particularly for those ages 65 and above. Based on our empirical findings, the rising share of the oldest age group decreases the average elasticity of substitution (weighted by population expenditure shares). In terms of the model, a lower average elasticity of substitution flattens the Phillips curve. A flatter Phillips curve, in turn, may influence the effectiveness and transmission channel of monetary policy.

Our empirical analysis relies on barcode-level retail purchases from the NielsenIQ Homescan Consumer Panel from 2004–2019. This data captures a large portion of retail purchases which are a significant component of overall expenditures. We aggregate the barcode-level purchases into five age groups ranging between 25 and 65+ and group each age-barcode observation into one of more than 1000 disaggregated retail product modules – detailed categories of similar retail products. We estimate the elasticity of substitution within these product modules by age following methods developed by Feenstra (1994) and Broda and Weinstein (2006) and rely on the application of estimating elasticities using NielsenIQ data in Jaravel (2019).

Our findings show that the youngest cohort (aged 25–34) exhibits a higher mean and median elasticity of substitution compared to the oldest group. Additionally, while the results for the age groups between 35 and 64 are not monotonic, the youngest cohort consistently shows the highest elasticity, whereas the oldest cohort continually exhibits the lowest elasticity among all groups. These findings are particularly stark when we control for income and focus on the top two quartiles of the income distribution, which, on average, represent over 70% of the expenditure share.

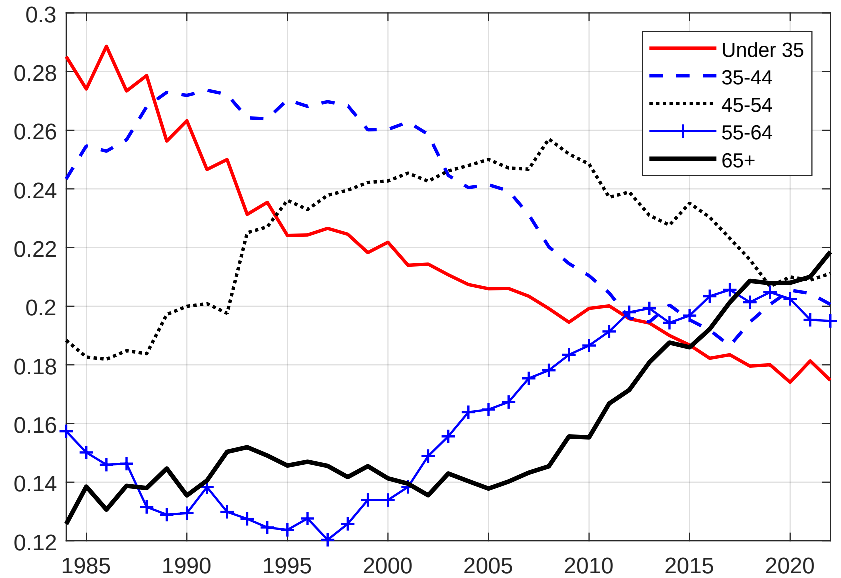

Figure 1: Consumption Expenditure Shares by Age Group

Notes: Data is the share of total consumption by age group of household reference person. Data is from the Bureau of Labor Statistics, Consumer Expenditure Surveys, Demographic Tables.

Based on these empirical results, we study the implications of heterogeneity in the elasticity of substitution by extending the Rotemberg model of price adjustment to include multiple types of consumers. While the slope of the Phillips curve depends on the elasticity of substitution in a representative agent model, the slope in our model is a function of the population-weighted average elasticity. As the population ages, the average elasticity falls, leading to a flatter Phillips curve.

The leading alternative to the Rotemberg model is the Calvo model. To a first-order approximation around a zero-inflation steady state, when production is constant returns to scale the elasticity of substitution does not affect the slope of the Phillips curve in the Calvo model. This may appear to limit the generality of our results. However, we show in Online Appendix B that the elasticity of substitution does affect a firm's decision to update prices in a static menu costs model. Since the Calvo model is a variant of the dynamic menu costs model, in which adjusting prices is costless but only allowed in certain periods, our results extend beyond the Rotemberg framework.

Our theoretical results are important for at least three reasons. First, because the monetary policy transmission mechanism depends on the slope of the Phillips curve, ignoring the age distribution may affect the efficacy with which monetary policy is conducted. Second, the transmission of monetary policy will, in turn, have heterogeneous effects across the population. Finally, our findings relate to empirical research documenting the flattening of the Phillips curve in advanced economies.[1] As these economies are also aging, our model suggests that, all else equal, the Phillips curve will flatten.

Our study introduces a novel dimension by examining how variation in the elasticity of substitution by age affects inflation dynamics, an aspect not explored in existing literature. In doing so, our paper relates to studies that have argued about the stability of structural parameters (e.g. Rubio-Ramirez and Fernández-Villaverde (2007)) since it shows that the overall elasticity of substitution, rather than being stationary, will depend on the demographic composition of the population.

This paper also relates to work highlighting the role of age heterogeneity on monetary policy. Empirically, recent work by Juselius and Takáts (2021), estimating the Phillips curve across countries, finds that demographics affect the level of inflation. Eggertsson et al. (2019) show that through supply and demand factors in the savings market, aging life-cycle savers push down the natural rate of interest, leading to a potentially binding zero lower bound on the nominal interest rate.

Through different channels, other studies have stressed the role of age distribution on the efficacy of monetary policy. For instance, Berg et al. (2021) show that through a wealth effect, interest rate sensitivity on consumption may be age-dependent, while Leahy and Thapar (2022) show that monetary policy could be affected by the age distribution of entrepreneurs. Mangiante (2023) finds composition bundles of older households are weighted more heavily toward products with higher price rigidities, affecting aggregate price adjustments and responsiveness to monetary policy shocks. As in Mangiante (2023), we derive the implications from price sensitivity, however, we focus on the price channel whereby older households themselves are relatively less price-sensitive.

Our work also contributes to the closely related literature that relies on micro-level data which studies the heterogeneous effects that age can have on consumption behavior. Two salient examples are Bornstein (2021) and Aguiar and Hurst (2007). Our results are broadly consistent with Bornstein (2021) who, also relying on NielsenIQ's Consumer Panel data, estimates the degree of persistence of consumption choices across ages and finds that the

[1] See, for instance, Stock and Watson (2020) and Del Negro et al. (2020). On the other hand, Hazell et al. (2022) claim that most of this measured flattening is due to shifting inflation expectations rather than the slope changing.

consumption of younger groups is less persistent than that of older cohorts. Also relying on micro-level data, but through a different channel, Aguiar and Hurst (2007) present evidence that older households spend more time shopping, resulting in lower overall prices for goods.

Finally, our paper connects with the voluminous Heterogeneous Agent New Keynesian - HANK - literature summarized in McKay and Wolf (2023) and Kaplan and Violante (2018). As individuals accumulate wealth over their lifetimes, accounting for the diversity in wealth holdings becomes crucial for understanding an aging population's effect on monetary policy. Nonetheless, even in the absence of alterations in the scale or composition of wealth portfolios, if consumption patterns vary with age, concentrating on the 'wealth channel' may overlook additional economic mechanisms. We discuss the empirical details of these consumption patterns in the next section.

2 Data

Our data source is the NielsenIQ Homescan Consumer Panel from 2004-2019. This is rotating, nationally representative panel that surveys between 40,000 and 60,000 households each year. The data captures a large fraction of all retail purchases which is a large part of overall household expenditures. Households are asked to scan all purchases of products that have a barcode (a UPC). We refer to each individual product as a *barcode*. Overall, our data captures over 900 million transactions. For each transaction, the price and quantities are recorded which we then match with the demographics of the purchasers. The NielsenIQ data classifies each barcode into larger product departments, which are further subdivided into product groups and product modules. There are 11 departments, and around 120 product groups and 1300 product modules. Departments include categories such as health and beauty, dry grocery, dairy, packaged meat, and general non-grocery merchandise. Product groups are subdivisions of departments that are typically found in close proximity to each other in retail establishments such as office supplies, household cleaners, and frozen pizza. Product modules are the most granular level of aggregation of barcodes we consider. An example of this is the *light beer* product module in the *beer* product group and *alcohol* department.

Our focus is on estimating the elasticity of substitution within product modules. For this, we consider only product modules that are present in all years as some enter and exit in various years. We also omit additional purchases that do not have UPCs, called "magnet" items.[2]

[2]For example, a large category of "magnet" items are fresh produce.

2.1 Estimating the elasticity of substitution by age

For each age group, groups ages 25 – 34, 35 – 44, 45 – 55 and 65+, consider an upper-level utility function

$$\mathbb{U} = u\left(C_1, C_2, ..., C_M\right)$$

where C_m is composite consumption of product module $m \in \{1, 2, ..., M\}$. For each module, C_m is a CES aggregator over barcodes within each module

$$C_m = \left(\sum_{b \in B_m} \left(d_{mb} q_{mb}\right)^{\frac{\sigma_m}{\sigma_m - 1}}\right)^{\frac{\sigma_m - 1}{\sigma_m}} \tag{1}$$

where σ_m is the elasticity of substitution between barcodes b within product module m, q_{mb} is the quantity, and d_{mb} is unobserved quality.

Our estimates of σ_m follow Feenstra (1994) and extensions by Broda and Weinstein (2006). Broda and Weinstein (2010) provide an intuitive explanation of the method. Briefly, there is a supply and demand component to each product. The minimum cost function of Equation (1) for each product module can be represented by

$$\Delta ln(s_{mbt}) = \alpha_{mt} - (\sigma_m - 1)\Delta ln(p_{mbt}) + \epsilon_{mbt} \tag{2}$$

where s_{mbt} is the expenditure share of barcode b at time t in product module m and p_{mbt} is the corresponding per unit price. The change in the unobserved quality (d_{mbt}) is captured in the ϵ_{mbt}. As for the supply side, the inverse supply curve, with ω_m denoting the inverse supply elasticity, can be represented as

$$\Delta ln(p_{mbt}) = \phi_{mt} + \frac{\omega_m}{1 + \omega_m}\Delta ln(s_{mbt}) + \xi_{mbt}. \tag{3}$$

By taking differences with respect to a reference barcode k in Equations (2) and (3), the intercept terms α_{mt} and ϕ_{mt} can be eliminated.[3] We can then rewrite Equations (2) and (3), respectively, as

$$\begin{aligned} \Delta^k ln(s_{mbt}) &= -(\sigma_m - 1)\Delta^k ln(p_{mbt}) + \epsilon^k_{mbt} \\ \Delta^k ln(p_{mbt}) &= \frac{\omega_m}{1 + \omega_m}\Delta^k ln(s_{mbt}) + \xi^k_{mbt} \end{aligned}$$

where k is the reference barcode and $\Delta^k ln(s_{mbt}) = \Delta ln(s_{mbt}) - \Delta ln(s_{mkt})$ and $\Delta^k ln(p_{mbt}) = \Delta ln(p_{mbt}) - \Delta ln(p_{mkt})$.

[3] This is subtracting an expression $\Delta ln(s_{mkt}) = \alpha_{mt} - (\sigma_m - 1)\Delta ln(p_{mkt}) + \epsilon_{mkt}$ for good k on the demand side and an analogous expression for supply.

Assuming the error terms ϵ^k_{mbt} and ξ^k_{mbt} are uncorrelated, the differenced demand and supply expressions can be combined as

$$\begin{aligned}\left(\Delta^k ln(p_{mbt})\right)^2 &= \underbrace{\frac{\omega_m}{(1+\omega_m)(\sigma_m-1)}}_{\theta_{m1}}\left(\Delta^k ln(s_{mbt})\right)^2 \\ &- \underbrace{\frac{1-\omega_m(\sigma_m-2)}{(1+\omega_m)(\sigma_m-1)}}_{\theta_{m2}}\left(\Delta^k ln(p_{mbt})\Delta^k ln(s_{mbt})\right) + u_{mbt}. \qquad (4)\end{aligned}$$

For the estimation, we aggregate quantities and unit prices for each barcode and age group and quarter. We use only continuing barcodes (barcodes available in time t and $t-1$) as to measure per-period changes in quantities and prices. We follow the estimation procedure in Jaravel (2019) who does a similar estimation to ours using the NielsenIQ data, but by income. Equation (4) is estimated using weighted least squares. We then back out σ_m and ω_m from the θ_{m1} and θ_{m2} terms, with the criteria $\sigma_m > 1$ and $\omega_m > 0$. Where these restrictions are not met, we evaluate the objective function for values of $\sigma_m > 1$ to obtain the estimates. For each age group, we follow trade literature that estimates aggregate elasticities as the weighted sector-specific trade elasticities (see e.g., Imbs and Mejean (2017)) by share weighting the product module elasticities (σ_m) by expenditure shares.[4]

2.2 Estimates of the elasticity of substitution by age

Figure 2 plots the elasticity of substitution σ by age group. The left panel depicts the mean values and the median is shown on the right. The patterns show a general decrease in σ by age, although it flattens out near mid-age before descending for the oldest age group. The overall picture for both the mean and median values are similar, but the median estimates are lower. The median estimates of 5.73 for the oldest ages and 7.02 for the youngest are in the range of other empirical work such as Broda and Weinstein (2010) and Hottman et al. (2016). The difference in elasticities across ages reaches a maximum between the oldest and youngest ages at 1.29 for the medians and 1.55 for the means. In comparison, Faber and Fally (2022) estimate a difference in elasticity of substitution of 0.375 between the richest and poorest quintiles.[5] Online Appendix C reports the elasticities by age and year in the sample.

[4]Specifically, for each age we calculate the expenditure share weighted elasticity of each module by year and the overall elasticity σ is the average value across years.

[5]While Faber and Fally (2022) use the NielsenIQ data, their estimation procedure differs from ours and their level of aggregation is at the brand-level across all product departments.

Figure 2: σ by Age Group

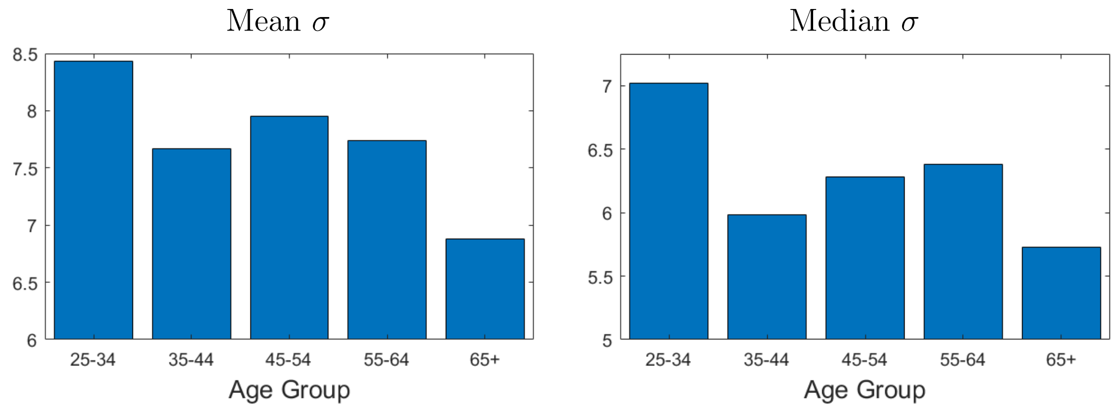

Notes: Mean and median σs are the expenditure weighted values of σ by age over each module. Each product module is required to have at least 20 households purchase in that module (see Online Appendix C for further details). The number of modules with required observations for each age group are: Age 25–34: 378; Age 35–44: 632; Age 45–54: 743; Age 55–64: 768; Ages 65+: 742.

To isolate contributions of differences in elasticities across age, we conduct a decomposition exercise capturing differences in elasticities within the same product modules and from differences in the composition in consumption baskets. Here, we take the difference in overall elasticity estimates from age group $a = \{25\text{-}34,\ 35\text{-}44,\ 45\text{-}54,\ 55\text{-}64\}$ relative to age the oldest age group (ages 65+). This difference is approximated as the sum of differences from product module elasticities and differences in expenditure shares in each module:

$$\sigma^a - \sigma^{65+} \approx \underbrace{\sum_{m=1}^{M} \omega_m^{65+}(\sigma_m^a - \sigma_m^{65+})}_{\text{difference from } \sigma} + \underbrace{\sum_{m=1}^{M} \sigma_m^{65+}(\omega_m^a - \omega_m^{65+})}_{\text{difference from weights}}$$

where m is product module, and ω is the expenditure share weights for each module and age. The first terms on the right-hand side captures differences due to elasticities by varying these across modules but holding the expenditure weights constant. This captures differences due to elasticities. The second terms hold the product module elasticities fixed at those for the 65+ age group and vary the expenditure weights for each product module. This captures the composition effect.

Table 1 shows this decomposition. The left columns show the absolute differences in each component and the right columns express the contribution of each component as a share of the total differences. The largest difference in the overall elasticity estimates are between the 65+ and the 25-34 age groups. Here, the absolute differences from both components are the greatest between these groups relative to the others. When focusing on the share of the differences, one-third of the observed differences between the 65+ and the 25-34 age groups

Table 1: Contribution of Differences in Elasticities
Relative to Ages 65+

	Absolute Differences		Share of Difference	
Age	σ_m	weights	σ_m	weights
25–34	0.64	1.36	0.32	0.68
35–44	0.56	0.45	0.55	0.45
45–54	0.55	0.37	0.60	0.40
55–64	0.56	0.39	0.59	0.41
65+	0	0	0	0

are from lower elasticities of the old within the same modules and two-thirds of the difference is due to the composition effect. Moving to the remaining age groups, between 55 and 60 percent of the differences of elasticity estimates between each age group and age 65+ is due to the σ_ms with the remaining shares due to the composition effect.

This shows that differences in elasticities across age groups relative to the old is a combination of the two channels. The old have lower elasticities within the same product modules and their expenditures are more weighted toward products with lower elasticities. Our data includes retail purchases only, but Cravino et al. (2022) find that older age groups purchase more services. Extending on this, Mangiante (2023) shows that the service sectors have higher degrees of price rigidities. Although we cannot address service expenditures in our data, extending our results to include services may unveil additional differences due to the composition effect of expenditure baskets.

Finally, we want to address the positive correlation between age and income and how this may relate to the elasticity estimates. A possible concern is that the differences in elasticities by age may be driven by income rather than age itself. Here, we rerun our elasticity estimates by age but further parse the sample by lower and upper 50 percent of the income distribution (unconditioned by age). One issue with the NielsenIQ Homescan data is that incomes are reported in discrete bins with a two-year lag. Instead of using the reported income groups, we follow Faber and Fally (2022) by estimating expenditures per capita as a proxy for income group.[6] While expenditures may not accurately measure actual income, they do appear to identify relative income levels. In Online Appendix C, we show per capita expenditures are increasing with (two-year lagged) reported income bins.

Table 2 shows the estimated elasticities of substitution by age and income group. We also report the within-age expenditure shares for the higher and lower income levels. For

[6] As in Faber and Fally (2022), we obtain per capita expenditures by year regressing log total expenditures on household size dummies and household-level attributes. We then convert household expenditures into per capita terms by subtracting household size dummy coefficients from household expenditures.

Table 2: σ by Age and Income

	Lower 50% Income			Upper 50% Income		
Age	Mean	Median	Expenditure Share	Mean	Median	Expenditure Share
25–34	7.27	6.95	0.34	8.87	8.97	0.66
35–44	7.16	6.03	0.26	8.20	6.78	0.74
45–54	7.56	6.76	0.20	7.98	6.75	0.80
55–64	7.67	6.58	0.18	8.17	6.54	0.82
65+	6.84	5.81	0.24	6.64	5.85	0.76

Notes: Elasticities are estimated by age and the top and bottom half of estimated income group. Expenditure shares are the average share of total expenditure by each income group within an age group.

lower income levels, mean elasticities are slightly increasing until the 55–64 age group and are lowest for the oldest. The median estimates follow the pooled results, but with a less pronounced difference across ages. At higher income levels, the differences across ages are even starker than the pooled estimates. Broadly, if we consider an aggregate elasticity of substitution within age as an expenditure weighted average between the lower and upper income groups, the final column shows that expenditure shares are substantially higher for the upper income households. This may be why the pooled estimates qualitatively follow the upper income elasticity pattern across age. Moreover, the median elasticity estimates across incomes but within age are similar for ages 45 and above, but below age 45 the elasticities for the lower income group are smaller than the elasticities for the upper income group. Overall, while our main (pooled) elasticities by age are more influenced by the upper income households, the falling pattern of elasticities by age is apparent conditional on being higher income. Consequently, the differences in elasticities by age appear to be from a factor of age unrelated to income. The next section investigates the implications of this heterogeneity for the slope of the Phillips curve.

3 Aging and the Phillips Curve

To understand how our empirical findings from the previous section affect the Phillips curve, we use a model with monopolistic competition and nominal rigidities on the firm side and individuals with different ages on the household side. The details of the full general equilibrium model are included in Online Appendix A.

Let $c_{i,a,t}$ be the consumption of good i at time t for a person of age a. Total consumption of an aged a at time t is a CES composite,

$$c_{a,t} = \left(\int_0^1 c_{i,a,t}^{\frac{\sigma_a-1}{\sigma_a}}\, di\right)^{\frac{\sigma_a}{\sigma_a-1}}$$

where σ_a is the elasticity of substitution of age group a.

Each individual minimizes expenditures subject to the constraint of achieving an overall level of consumption. The optimization problem for group a is

$$\min \int_0^1 P_{i,t} c_{i,a,t} di$$

$$\text{s.t.} \left(\int_0^1 c_{i,a,t}^{\frac{\sigma_a-1}{\sigma_a}}\, di\right)^{\frac{\sigma_a}{\sigma_a-1}} \geq c_{a,t}.$$

In Online Appendix A we show that the demand function for good i can be written as

$$c_{i,a,t} = c_{a,t}\left(\frac{P_{i,t}}{P_{a,t}}\right)^{-\sigma_a}$$

where $P_{a,t}$ is the price index for people of age a. We denote aggregate consumption of age a individuals with upper-case letters. The demand functions can then be written in terms of aggregate variables,

$$C_{i,a,t} = C_{a,t}\left(\frac{P_{i,t}}{P_{a,t}}\right)^{-\sigma_a}.$$

Output of good i is produced with the production function $Y_{i,t} = A_t N_{i,t}$, where A_t is a stationary productivity term that may or may not move over the business cycle and $N_{i,t}$ is labor. Assuming firms pay a real wage of w_t, real marginal cost is given by w_t/A_t and, in equilibrium, supply of good i equals demand for good i:

$$Y_{i,t} = \sum_{a=1}^{N} C_{a,t}\left(\frac{P_{i,t}}{P_{a,t}}\right)^{-\sigma_a}.$$

Following Rotemberg (1982), we assume firms face a convex cost of price adjustment proportional to total output. Future profits are discounted by $\beta\lambda_t$ where $0 < \beta < 1$ is the subjective discount rate and λ_t is some function of the households' marginal utilities of consumption.[7] The dynamic optimization problem for firms entails choosing $P_{i,t}$ and $Y_{i,t}$ to

[7] We do not need to specify the form of λ_t to derive the Phillips curve in this model. For brevity, we do not show it here, but the derivation of the model is fully specified for completeness in Online Appendix A.

maximize

$$\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \lambda_t \left[\frac{P_{i,t} Y_{i,t}}{P_t} - mc_t Y_{i,t} - \frac{\phi}{2} Y_t \left(\frac{P_{i,t}}{P_{i,t-1}} - 1 \right)^2 \right]$$

subject to the constraint

$$Y_{i,t} = \sum_{a=1}^{N} C_{a,t} \left(\frac{P_{i,t}}{P_{a,t}} \right)^{-\sigma_a}.$$

As it is standard in Rotemberg pricing models, there is a symmetric equilibrium in which all firms choose the same price, $P_{i,t} = P_t$. Denoting the period t inflation rate by $\pi_t = P_t / P_{t-1} - 1$, the first order condition for the firm can be expressed as

$$\sum_{a=1}^{N} (\sigma_a - 1) C_{a,t} + \phi Y_t \pi_t (1 + \pi_t) = mc_t \sum_{a=1}^{N} \sigma_a C_{a,t} + \mathbb{E}_t \beta \frac{\lambda_{t+1}}{\lambda_t} \phi \pi_{t+1} (1 + \pi_{t+1}) Y_{t+1}. \tag{5}$$

Note that when σ_a is the same across ages, this collapses to the standard Rotemberg Phillips curve. Sacrificing analytical exactitude, a more intuitive and transparent equation can be obtained by log-linearizing (5), around a zero-inflation steady state:

$$\pi_t = \frac{\sum_{a=1}^{N} (\sigma_a - 1) s_a}{\phi} \tilde{mc}_t + \beta \pi_{t+1}^e + H_t \tag{6}$$

where s_a is the consumption share of the aged a cohort and

$$H_t = \frac{\sum_{a=1}^{N} (\sigma_a - 1) s_a}{\phi} \left(\frac{\sum_{a=1}^{N} \sigma_a (C_{a,t} - C_{a,ss})}{\sum_{a=1}^{N} \sigma_a C_{a,ss}} - \frac{\sum_{a=1}^{N} (\sigma_a - 1)(C_{a,t} - C_{a,ss})}{\sum_{a=1}^{N} (\sigma_a - 1) C_{a,ss}} \right).$$

If σ_a does not vary across ages then $H_t = 0$ and this, once again, collapses to the usual log-linearized Phillips curve under Rotemberg pricing,

$$\pi_t = \frac{\sigma - 1}{\phi} \tilde{mc}_t + \beta \pi_{t+1}^e.$$

Returning to our Phillips curve expressed in Equation (6), the slope with respect to marginal cost is

$$\frac{\sum_{a=1}^{N} (\sigma_a - 1) s_a}{\phi}. \tag{7}$$

With heterogeneity in σ_a, changes in consumption shares (s_a) will also change the slope of the Phillips curve. To clearly see this effect of aging on the slope of the Phillips curve, consider a case of two age groups, $a = \{y, o\}$, where y is the young and o is the old with respective elasticities σ_y and σ_o. We can rewrite the slope of the Phillips curve in equation (7) as

$$\frac{s_o(\sigma_o - 1) + (1 - s_o)(\sigma_y - 1)}{\phi}$$

where s_o is the consumption share of the old. If $\sigma_o < \sigma_y$, then the Phillips curve becomes flatter as the older generation consumes a higher fraction of output. This corresponds to our general empirical findings.

To understand the channel in which the elasticity of substitution affects the slope of the Phillips curve, we return to the non-linear Phillips curve given in Equation (5). Ignoring heterogeneity by setting $\sigma_a = \sigma$ gives the standard case, and rearranging

$$(\sigma - 1)C_t - mc_t C_t = \mathbb{E}_t \beta \frac{\lambda_{t+1}}{\lambda_t} \phi \pi_{t+1}(1 + \pi_{t+1}) Y_{t+1} - \phi Y_t \pi_t (1 + \pi_t)$$

where the right-hand side gives the expected benefits minus the costs of time t price adjustments.

A lower σ implies that products are less substitutable. As a result, firms, which now have more market power due to the lower elasticity, are less sensitive to changes in marginal cost. These firms can afford to delay price adjustments, leading to larger but less frequent price changes. Consequently, small changes in marginal costs are less likely to trigger immediate price adjustments, as a firm can either absorb the cost or wait for a later time to pass it on, without losing demand. This, in turn, causes inflation to respond less sharply to changes in marginal cost, resulting in a flatter Phillips Curve. In other words, compared to the standard case, the presence of a cohort with a lower elasticity implies that marginal costs have to move relatively more to exert the same inflationary pressure.

Returning to Equation (5) but reintroducing the consumer heterogeneity and keeping two groups, y and o,

$$(\sigma_o - 1)C_{o,t} + (\sigma_y - 1)C_{y,t} - mc_t(\sigma_o C_{o,t} + \sigma_y C_{y,t}) = \mathbb{E}_t \beta \frac{\lambda_{t+1}}{\lambda_t} \phi \pi_{t+1}(1 + \pi_{t+1}) Y_{t+1} - \phi Y_t \pi_t (1 + \pi_t),$$

the intuition described above holds. The right-hand side is the same as before, but now the marginal cost is scaled by the consumptions of different cohorts weighted by their elasticities. In the data, relative consumption is growing for the old relative to the young and the old correspondingly have a lower elasticity. Overall, this leads to a lower sensitivity of price adjustments and hence a flatter Phillips curve.

Before turning to the quantitative implications of the model, it is instructive to compare our modeling choices to some common alternatives. The most common alternative to the

Rotemberg model is the Calvo model of staggered price adjustment. With Calvo pricing and constant returns to scale in production, to a first-order approximation around a zero-inflation steady state the slope of the Phillips curve is independent of the elasticity of substitution. Whereas firms in the Rotemberg model internalize the effects of their pricing power in deciding the path of price adjustment, the probability of price adjustment is independent of pricing power in the Calvo model. In fact, in the Calvo model, a production function with decreasing returns to scale induces marginal cost dispersion across firms and a higher elasticity flattens the Phillips curve.[8] Because there is no marginal cost (or price) dispersion in the Rotemberg model, this effect is absent.

Finally, although less commonly used than either the Rotemberg or Calvo approaches, models with fixed costs of price adjustment, such as Golosov and Lucas Jr (2007), add important microfoundations to price-setting decisions. Accordingly, we analyze a one-period menu cost model in Online Appendix B. While the elasticity of substitution has an ambiguous effect on a firm's decision to update in general, we show that a plausible calibration implies that the willingness of a firm to update its prices is increasing in the elasticity of substitution. Thus, the conclusion in the fixed-cost model is similar to our findings in the environment with Rotemberg pricing. Adjudicating the proper choice of model is beyond the scope of this paper. Nevertheless, the conclusion that a lower share-weighted elasticity flattens the Phillips curve is, although not general, also not confined to a special case.

3.1 Quantitative Implications

To what extent does demographic change flatten the Phillips curve in the model? Our goal here is not to conduct a full-scale quantitative evaluation. However, as one reviewer pointed out, it is helpful to provide some insight into the significance of this demographic channel. We agree and, in this section, we use the estimated elasticities and the consumption data by age to parameterize the model to match the estimated slope of the Phillips curve from the empirical literature. We then feed the observed change in consumption shares by age through the model to isolate the observed change in the slope of the Phillips curve and compare this to findings in the literature.

In the theoretical framework we considered N age groups. In this exercise, we rewrite the slope of the log-linearized Phillips curve with respect to marginal cost in Equation (7) for

[8]See Galí (2015), Chapter 3.3. On the other hand, decreasing returns to scale does not affect the slope of the Phillips curve under Rotemberg pricing.

five groups in order to follow our empirical analysis as closely as possible:

$$\frac{\sum_{a=1}^{5} s_a(\sigma_a - 1)}{\phi} \tag{8}$$

where s_a is the consumption share of group a and $\sum_{a=1}^{5} s_a = 1$. Mapping this to the data, we consider the age groups to align in order with ages $a = \{<35, 35-44, 45-54, 55-64, 65+\}$.

We next take Equation 8 to the data to assess the contribution of aging on the change in its slope. We parameterize ϕ to match the observed slope of the Phillips curve whose estimated values range from 0.05 to 0.06. We select the midpoint (0.055) from Gagliardone et al. (2023) who estimate the Phillips curve using marginal costs from Danish firm-level data from 1999–2019. We do this by bringing our empirical estimates of σ by age (data in the left panel in Figure 2) and consumption shares by age (data in Figure 1) to Equation (8). Using 2022 data to parameterize the Phillips curve, we obtain $\phi = 122$ to match $\frac{\sum_{a=1}^{5} s_{2022,a}(\sigma_a - 1)}{122} = 0.055$, where $s_{2022,a}$ is the consumption share of age group a in 2022.

To compare how consumption age shares affects the slope of the Phillips curve, we next feed in the consumption shares by age in 1984, $s_{1984,a}$, into the expression: $\frac{\sum_{a=1}^{5} s_{1984,a}(\sigma_a - 1)}{122} = 0.056$. This change represents just over a two percent decrease in the slope of the Phillips curve from 1984-2022 ($\frac{0.055-0.056}{0.056} * 100\% = -2.3\%$). As mentioned in the introduction, there is debate about whether–and to what extent,–the slope of the Phillips curve has changed. Furlanetto and Lepetit (2024) survey some of the recent literature and report that the median estimates indicate the Phillips curve slope has been cut in half, although there is large degree of uncertainty around these estimates. Taking this as a benchmark to compare the changes in the slope in our model, the demographic channel we propose accounts for around 4.5 percent of the observed changes in the Phillips curve (i.e. $2.3/50 = 4.5$, so a 2.3 percent reduction in our model accounts for approximately 4.5 percent of the 50 percent reduction in the data). It is reasonable to argue that this effect is not large quantitatively, but is still one of the factors that may have caused the slope of the Phillips curve to flatten.

4 Conclusion

We add to the expansive literature investigating the consequences of heterogeneity in New Keynesian models. We do so by departing from previous papers that have focused on the effects of differences in wealth and access to liquidity, and present evidence older individuals have a lower elasticity of substitution. Motivated by an increase in the share of older households in the US population, we incorporate this finding into a model of nominal rigidities. We find that as the average elasticity of substitution decreases, so does the slope of the Phillips curve. This finding, in turn, has implications for the transmission channel of monetary policy.

We analytically isolate the channel in which heterogeneity in the elasticity of substitution affects the slope of the Phillips curve and provide a fairly straightforward quantitative analysis of its relative importance. An important question remains unexplored: how might the channel we examine affect normative aspects of monetary policy?

References

Aguiar, M. and E. Hurst (2007): "Life-cycle prices and production," American Economic Review, 97, 1533–1559.

Berg, K. A., C. C. Curtis, S. Lugauer, and N. C. Mark (2021): "Demographics and monetary policy shocks," Journal of Money, Credit and Banking, 53, 1229–1266.

Bornstein, G. (2021): "Entry and profits in an aging economy: The role of consumer inertia," Review of Economic Studies, forthcoming.

Broda, C. and D. E. Weinstein (2006): "Globalization and the Gains from Variety," The Quarterly journal of economics, 121, 541–585.

——— (2010): "Product creation and destruction: Evidence and price implications," American Economic Review, 100, 691–723.

Cravino, J., A. Levchenko, and M. Rojas (2022): "Population aging and structural transformation," American Economic Journal: Macroeconomics, 14, 479–498.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020): "Whats up with the Phillips Curve?" National Bureau of Economic Research Working Paper.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2019): "A model of secular stagnation: Theory and quantitative evaluation," American Economic Journal: Macroeconomics, 11, 1–48.

Faber, B. and T. Fally (2022): "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," The Review of Economic Studies, 89, 1420–1459.

Feenstra, R. C. (1994): "New product varieties and the measurement of international prices," The American Economic Review, 157–177.

Furlanetto, F. and A. Lepetit (2024): "The Slope of the Phillips Curve," Board of Governors of the Federal Reserce System Finance and Economics Discussion Series, 2024-043.

Gagliardone, L., M. Gertler, S. Lenzu, and J. Tielens (2023): "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications," National Bureau of Economic Research, 31382.

Galí, J. (2015): Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications, Princeton University Press.

Golosov, M. and R. E. Lucas Jr (2007): "Menu costs and Phillips curves," Journal of Political Economy, 115, 171–199.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022): "The slope of the Phillips Curve: evidence from US states," The Quarterly Journal of Economics, 137, 1299–1344.

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016): "Quantifying the sources of firm heterogeneity," The Quarterly Journal of Economics, 131, 1291–1364.

Imbs, J. and I. Mejean (2017): "Trade elasticities," Review of International Economics, 25, 383–402.

Jaravel, X. (2019): "The unequal gains from product innovations: Evidence from the us retail sector," The Quarterly Journal of Economics, 134, 715–783.

Juselius, M. and E. Takáts (2021): "Inflation and demography through time," Journal of Economic Dynamics and Control, 128, 104136.

Kaplan, G. and G. L. Violante (2018): "Monetary policy according to HANK," Journal of Economic Perspectives, 32, 167–194.

Leahy, J. V. and A. Thapar (2022): "Age structure and the impact of monetary policy," American Economic Journal: Macroeconomics, 14, 136–173.

Mangiante, G. (2023): "Demographic Trends and the Transmission of Monetary Policy," Working Paper: SSRN 4149657.

McKay, A. and C. Wolf (2023): "Monetary Polic and Inequality," Journal of Economic Perspectives, 37, 121–144.

Rotemberg, J. J. (1982): "Monopolistic price adjustment and aggregate output," The Review of Economic Studies, 49, 517–531.

Rubio-Ramirez, J. F. and J. Fernández-Villaverde (2007): "How structural are structural parameters?" National Bureau of Economic Research Macroeconomics Annual, 22.

Stock, J. H. and M. W. Watson (2020): "Slack and cyclically sensitive inflation," Journal of Money, Credit and Banking, 52, 393–428.

Online Appendix

Does the Phillips Curve Lie Down as we Age?

Chadwick Curtis, Julio Garín, and Robert Lester

A Derivation of Main Model

For completeness, we start with the foundations of the framework presented in the main text. Recall that each type of person minimizes expenditure subject to the constraint of achieving an overall level of consumption.

$$\min \int_0^1 P_{i,t} c_{i,a,t} di$$
$$\text{subject to } \left(\int_0^1 c_{i,a,t}^{\frac{\sigma_a - 1}{\sigma_a}} di \right)^{\frac{\sigma_a}{\sigma_a - 1}} \geq c_{a,t}.$$

$c_{i,a,t}$ is the consumption of good i of a person of age a at time t. Lower-case letters denote per-capita variables. The first order condition for for good i is

$$P_{i,t} - \lambda_{a,t} c_{i,a,t}^{\frac{-1}{\sigma_a}} c_{a,t}^{\frac{1}{\sigma_a}} = 0 \Leftrightarrow c_{i,a,t} = c_{a,t} P_{i,t}^{-\sigma_a} \lambda_{a,t}^{\sigma_a}$$

where $\lambda_{a,t}$ is the multiplier on the constraint. Substituting the FOC into the constraint (and evaluating it at equality) gives

$$\left(\int_0^1 c_{i,a,t}^{\frac{\sigma_a - 1}{\sigma_a}} di \right)^{\frac{\sigma_a}{\sigma_a - 1}} = \lambda_{a,t}^{\sigma_a} c_{a,t} \left(\int_0^1 P_{i,t}^{1-\sigma_a} di \right)^{\frac{\sigma_a}{\sigma_a - 1}} = c_{a,t} \Leftrightarrow$$
$$\lambda_{a,t} = \left(\int_0^1 P_{i,t}^{1-\sigma_a} di \right)^{\frac{1}{1-\sigma_a}} = P_{a,t}$$

The demand function for a person of age a is

$$c_{i,a,t} = c_{a,t} \left(\frac{P_{i,t}}{P_{a,t}} \right)^{-\sigma_a}.$$

Note that we can write this demand function in aggregate variables as well. Denoting

$C_{i,a,t}$ as the aggregate consumption of good i at time t for all people of age a gives

$$C_{i,a,t} = C_{a,t}\left(\frac{P_{i,t}}{P_{a,t}}\right)^{-\sigma_a}.$$

Assume that people of every age group are part of a large household. The household planner puts equal weight on their utility. Let the number of people of age a be given by ν_a. The intertemporal utility function is given by

$$U = \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \left[\sum_{a=1}^{N} \nu_a \frac{\left(c_{a,t}(1-n_{a,t})^{\theta_a}\right)^{1-\gamma}}{1-\gamma}\right]$$

where γ is the coefficient of relative risk aversion, θ_a governs the marginal utility of leisure for a person of age a.

Letting the general price level be P_t, the household's budget constraint in real terms is given by

$$\sum_{a=1}^{N} \nu_a c_{a,t} + B_t = \frac{P_{t-1}}{P_t}(1+i_{t-1})B_{t-1} + w_t \sum_{a=1}^{N} \nu_a n_{a,t} + D_t.$$

B_t is the aggregate quantity of bonds held by the household. Bonds purchased in period t pay a nominal interest rate of i_t. D_t are dividends earned by the firm and rebated to the household.

The first order conditions are given by

$$\begin{aligned} c_{a,t}:\ & c_{a,t}^{-\gamma}(1-n_{a,t})^{\theta_a(1-\gamma)} = \lambda_t \\ n_{a,t}:\ & \theta_a c_{a,t}^{1-\gamma}(1-n_{a,t})^{\theta_a-1-\theta_a\gamma} = \lambda_t w_t \\ B_t:\ & \lambda_t = \beta \mathbb{E}_t \lambda_{t+1}\frac{1+i_t}{1+\pi_{t+1}}. \end{aligned}$$

Denote aggregate variables by big letters. Accordingly,

$$\begin{aligned} C_{a,t} &= \nu_a c_{a,t} \\ N_{a,t} &= \nu_a n_{a,t}. \end{aligned}$$

A.1 Firm Side

Firms, indexed by i, produce a differentiated good, $y_{i,t}$. The production function for good i is

$$Y_{i,t} = A_t \sum_{a=1}^{N} N_{a,i,t}$$

The marginal cost of producing another unit of $Y_{i,t}$ is

$$mc_t = \frac{w_t}{A_t}.$$

The output produced by firm i gets sold to consumers of different ages:

$$Y_{i,t} = \sum_{a=1}^{N} C_{i,a,t} = \sum_{a=1}^{N} C_{a,t} \left(\frac{P_{i,t}}{P_{a,t}} \right)^{-\sigma_a}.$$

The dynamic optimization problem for firms entails choosing $P_{i,t}$ and $Y_{i,t}$ to maximize

$$\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \lambda_t \left[\frac{P_{i,t} Y_{i,t}}{P_t} - mc_t Y_{i,t} - \frac{\phi}{2} Y_t \left(\frac{P_{i,t}}{P_{i,t-1}} - 1 \right)^2 \right]$$

subject to the constraint

$$Y_{i,t} = \sum_{a=1}^{N} C_{a,t} \left(\frac{P_{i,t}}{P_{a,t}} \right)^{-\sigma_a}.$$

Let $\mu_{i,t}$ denote the Lagrangian multiplier on the constraint. The first order conditions are

$$Y_{i,t}: \quad mc_t - \frac{P_{i,t}}{P_t} = \mu_{i,t}$$

$$\begin{aligned} P_{i,t}: \quad & \frac{Y_{i,t}}{P_t} - \phi \frac{Y_t}{P_{i,t-1}} \left(\frac{P_{i,t}}{P_{i,t-1}} - 1 \right) \\ & + \mu_{i,t} \sum_{a=1}^{N} \sigma_a C_{a,t} P_{a,t}^{\sigma_a} P_{i,t}^{-\sigma_a - 1} \\ + \beta \mathbb{E}_t & \left[\frac{\lambda_{t+1}}{\lambda_t} \phi Y_{t+1} \frac{P_{i,t+1}}{P_{i,t}^2} \left(\frac{P_{i,t+1}}{P_{i,t}} - 1 \right) \right] = 0. \end{aligned}$$

Next, combine these two equations to eliminate the multiplier.

$$\begin{aligned}\frac{Y_{i,t}}{P_t} - \phi\frac{Y_t}{P_{i,t-1}}\left(\frac{P_{i,t}}{P_{i,t-1}} - 1\right)\\ +\left(mc_t - \frac{P_{i,t}}{P_t}\right)\sum_{a=1}^{N}\sigma_a C_{a,t}P_{a,t}^{\sigma_a}P_{i,t}^{-\sigma_a-1}\\ +\beta\mathbb{E}_t\left[\frac{\lambda_{t+1}}{\lambda_t}\phi Y_{t+1}\frac{P_{i,t+1}}{P_{i,t}^2}\left(\frac{P_{i,t+1}}{P_{i,t}} - 1\right)\right] = 0.\end{aligned}$$

Focusing on the symmetric equilibrium in which firms choose the same price and quantity implies that the price indexes across different age groups are the same, $P_{j,t} = P_t$). Let $P_{i,t} = P_t$. Inflation is $\frac{P_t}{P_{t-1}} = 1 + \pi_t$. The (non-linear) Phillips curve is:

$$\sum_{a=1}^{N}(\sigma_a - 1)C_{a,t} + \phi Y_t\pi_t(1+\pi_t) = mc_t\sum_{a=1}^{N}\sigma_a C_{a,t} + \mathbb{E}_t\beta\frac{\lambda_{t+1}}{\lambda_t}\phi\pi_{t+1}(1+\pi_{t+1})Y_{t+1}.$$

A.2 Aggregation

Because each firm chooses the same price, they also produce the same level of output. The aggregate production function is therefore given by

$$Y_t = A_t\sum_{a=1}^{N}N_{a,t}.$$

Dividends are given by

$$D_t = Y_t - w_t\sum_{a=1}^{N}N_{a,t} - \frac{\phi}{2}Y_t\pi_t^2.$$

Substituting this into the household budget constraint and assuming that bonds are in zero net supply gives

$$\begin{aligned}\sum_{a=1}^{N}C_{a,t} + B_t = \frac{P_{t-1}}{P_t}(1+i_{t-1})B_{t-1} + w_t\sum_{a=1}^{N}N_{a,t} + D_t \Leftrightarrow\\ \sum_{a=1}^{N}C_{a,t} = w_t\sum_{a=1}^{N}N_{a,t} + Y_t - w_t\sum_{a=1}^{N}N_{a,t} - \frac{\phi}{2}Y_t\pi_t^2 \Leftrightarrow\\ \sum_{a=1}^{N}C_{a,t} + \frac{\phi}{2}Y_t\pi_t^2 = Y_t.\end{aligned}$$

To close the model, assume that the central bank follows the following rule in setting the nominal interest rate.

$$i_t = i_{ss} + \varphi\pi_t$$

with the parameter $\varphi > 1$ satisfying the Taylor principle.

A.3 Summarizing Equilibrium Conditions

The endogenous variables are: $c_{a,t}, n_{a,t}, C_{a,t}, N_{a,t}, \lambda_t, mc_t, w_t, Y_t, \pi_t, i_t$.

$$c_{a,t}^{-\gamma}(1-n_{a,t})^{\theta_a(1-\gamma)} = \lambda_t \tag{9}$$

$$\theta_a c_{a,t}^{1-\gamma}(1-n_{a,t})^{\theta_a-1-\theta_y\gamma} = \lambda_t w_t \tag{10}$$

$$\lambda_t = \beta E_t \lambda_{t+1}\frac{1+i_t}{1+\pi_{t+1}} \tag{11}$$

$$C_{a,t} = \nu_a c_{a,t} \tag{12}$$

$$N_{a,t} = \nu_a n_{a,t} \tag{13}$$

$$mc_t = \frac{w_t}{A_t} \tag{14}$$

$$\sum_{a=1}^{N}(\sigma_a-1)C_{a,t} + \phi Y_t\pi_t(1+\pi_t) = mc_t\sum_{a=1}^{N}\sigma_a C_{a,t} + \mathbb{E}_t\beta\frac{\lambda_{t+1}}{\lambda_t}\phi\pi_{t+1}(1+\pi_{t+1})Y_{t+1}. \tag{15}$$

$$Y_t = A_t\sum_{a=1}^{N}N_{a,t} \tag{16}$$

$$\sum_{a=1}^{N}C_{a,t} + \frac{\phi}{2}Y_t\pi_t^2 = Y_t \tag{17}$$

$$i_t = i_{ss} + \varphi\pi_t \tag{18}$$

A.4 Rotemberg Pricing: Log-linear approximation

As mentioned, we assume steady-state inflation rate is 0. A couple notes. First, the resource constraint is

$$Y_t = \sum_{a=1}^{N}C_{a,t} + \frac{\phi}{2}\pi_t^2 Y_t.$$

In steady state,

$$Y_{ss} = \sum_{a=1}^{N}C_{a,ss}.$$

Let s_a be the share of output consumed by people of age a in steady state. So

$$C_{a,ss} = s_a Y_{ss}.$$

Second, evaluating the Phillips curve in steady state shows that steady-state marginal cost is

$$mc_{ss} = \frac{\sum_{a=1}^{N}(\sigma_a-1)C_{a,ss}}{\sum_{a=1}^{N}\sigma_a C_{a,ss}}.$$

This can be simplified to

$$mc_{ss} = \frac{\sum_{a=1}^{N}(\sigma_a - 1)s_a}{\sum_{a=1}^{N}\sigma_a s_a}.$$

Noting that $\sum_{a=1}^{N} s_a = 1$ and letting $\sum_{a=1}^{N} s_a\sigma_a = \bar{\sigma}$ be the weighted average of the elasticities, we can write steady-state marginal cost as

$$mc_{ss} = \frac{\bar{\sigma} - 1}{\bar{\sigma}}.$$

We then take the natural log of both sides of the Phillips curve. This gives

$$\ln\left[\sum_{a=1}^{N}(\sigma_a - 1)C_{a,t} + \phi Y_t\pi_t(1+\pi_t)\right] = \ln\left[mc_t\sum_{a=1}^{N}\sigma_a C_{a,t} + \mathbb{E}_t\beta\frac{\lambda_{t+1}}{\lambda_t}\phi\pi_{t+1}(1+\pi_{t+1})Y_{t+1}\right].$$

Now, taking a first-order Taylor series expansion around the steady state gives

$$\frac{\sum_{a=1}^{N}(\sigma_a - 1)(C_{a,t} - C_{a,ss}) + \phi Y_{ss}\pi_t}{\sum_{a=1}^{N}(\sigma_a - 1)C_{a,ss}} = \frac{(mc_t - mc_{ss})\sum_{a=1}^{N}\sigma_a C_{a,ss}}{mc_{ss}\sum_{a=1}^{N}\sigma_a C_{a,ss}} + \frac{mc_{ss}\sum_{a=1}^{N}\sigma_a(C_{a,t} - C_{a,ss})}{mc_{ss}\sum_{a=1}^{N}\sigma_a C_{a,ss}} + \frac{\beta\phi Y_{ss}\pi_{t+1}^e}{mc_{ss}\sum_{a=1}^{N}\sigma_a C_{a,ss}}.$$

Note that the multipliers completely disappear because we are approximating around an inflation rate of 0. Then, we can write the log-linearized Phillips curve as

$$\frac{\phi\pi_t}{\sum_{a=1}^{N}(\sigma_a - 1)s_a} + A_t = \tilde{mc}_t + B_t + \frac{\beta\phi\pi_{t+1}^e}{mc_{ss}\sum_{a=1}^{N}\sigma_a s_a}$$

where $\tilde{mc}_t$ is marginal cost's percent deviation from steady state.[9] Using the steady-state condition for marginal cost, we can write this as

$$\frac{\phi\pi_t}{\sum_{a=1}^{N}(\sigma_a - 1)s_a} + A_t = \tilde{mc}_t + B_t + \frac{\beta\phi\pi_{t+1}^e}{\sum_{a=1}^{N}(\sigma_a - 1)s_a}.$$

Doing some cross multiplication yields

$$\pi_t = \frac{\sum_{a=1}^{N}(\sigma_a - 1)s_a}{\phi}\tilde{mc}_t + \beta\pi_{t+1}^e + H_t$$

[9] A_t and B_t are given by $A_t = \frac{\sum_{a=1}^{N}(\sigma_a - 1)(C_{a,t} - C_{a,ss})}{\sum_{a=1}^{N}(\sigma_a - 1)C_{a,ss}}$ and $B_t = \frac{\sum_{a=1}^{N}\sigma_a(C_{a,t} - C_{a,ss})}{\sum_{a=1}^{N}\sigma_a C_{a,ss}}$.

where $H_t = \frac{\sum_{a=1}^{N}(\sigma_a - 1)s_a}{\phi}(B_t - A_t)$.

Finally, we can write the Phillips curve as

$$\pi_t = \frac{\bar{\sigma} - 1}{\phi}\tilde{mc}_t + \beta\pi^e_{t+1} + H_t.$$

If σ_a is constant across ages, then $H_t = 0$, and this collapses to the usual log-linearized Phillips curve under Rotemberg pricing:

$$\pi_t = \frac{\sigma - 1}{\phi}\tilde{mc}_t + \beta\pi^e_{t+1}.$$

Returning to our Phillips curve, the slope with respect to marginal cost is

$$\frac{\bar{\sigma} - 1}{\phi}. \tag{19}$$

To the extent that older people have a lower σ_a, an older population will flatten the Phillips curve.

B Menu Cost Pricing

With Rotemberg pricing, a lower elasticity of substitution flattens the Phillips curve. This result is derived in a model in which firms face a quadratic cost of adjustment. With quadratic costs, it is optimal for firms to adjust their price every period, but stop short of the optimal (frictionless) reset price.

The leading alternative to Rotemberg is the Calvo model in which firms get to reset their prices with some exogenous probability. However, under Calvo pricing, the elasticity of substitution parameter does not show up directly in the slope of the Phillips curve. In what follows, we argue that the probability of price adjustment itself likely depends on the value of σ. Specifically, we assume that firms pay a fixed cost, i.e. a menu cost, to update prices. For most parameterizations we find that, for a given menu cost, firms will be more likely to change prices when demand is relatively elastic. As demand becomes more inelastic, firms change prices less frequently which means that prices will be "stickier" *à la* Calvo. So the probability of price adjustment, which is commonly given as a structural parameter in these models, will actually be affected by the age distribution – also consistent with the message in Rubio-Ramirez and Fernández-Villaverde (2007). All else equal, an older population maps to a lower probability of price adjustment which, in the Calvo framework, flattens the Phillips curve.

Consider a monopolist who faces a CES demand curve given by $Q = AP^{-\sigma}$ where $\sigma > 1$ is the price elasticity of demand and A represents market size. Assume a linear cost function,

$$C(Q) = \phi Q.$$

A constant returns to scale production function maps into a linear cost function, so this is without loss of generality. From a firm's perspective, this demand function is equivalent up to a constant of the monopolistic competition demand curve (because the firm takes the aggregate price level and aggregate income as exogenous).

Profit maximization implies an optimal price of

$$P^* = \frac{\sigma}{\sigma - 1}\phi$$

which is the typical markup over marginal cost.

Suppose that the firm comes into the period with a preset price of $\bar{P}$, i.e. the price on the menu that was presumably selected in an earlier period. The firm's profit under $\bar{P}$ is

$$\Pi_{\text{fixed}} = A\bar{P}^{1-\sigma} - \phi A\bar{P}^{-\sigma}.$$

If marginal cost changes between periods (i.e. a productivity shock), then $\bar{P}$ is no longer optimal. Suppose the firm faces a menu cost, f, of adjusting prices. If it pays the adjustment cost, the firm will adjust prices all the way to P^*. In the case that the firm adjusts to the optimal price, profit is given by

$$\Pi^*_{\text{var}} = AP^{*^{1-\sigma}} - \phi AP^{*^{-\sigma}} - f.$$

It follows that the firm will adjust its price if an only if

$$\Pi^*_{\text{var}} > \Pi_{\text{fixed}} \Leftrightarrow AP^{*^{1-\sigma}} - \phi AP^{*^{-\sigma}} - (A\bar{P}^{1-\sigma} - \phi A\bar{P}^{-\sigma}) > f.$$

Define $\hat{f}$ as the menu cost where the firm is indifferent to adjusting prices. Formally,

$$\hat{f} = AP^{*^{1-\sigma}} - \phi AP^{*^{-\sigma}} - (A\bar{P}^{1-\sigma} - \phi A\bar{P}^{-\sigma}).$$

If $\hat{f}$ is increasing in σ then a firm facing a more price-sensitive demand curve will require a larger f to keep its prices the same. That is, for a given f, the firm will be more likely to change the price the larger the elasticity (σ). Intuitively, the more price sensitive is a firm's demand curve the more willing they will be to pay the fixed cost and update prices. The derivative is

$$\frac{\partial \hat{f}}{\partial \sigma} = \frac{\partial \Pi^*_{\text{var}}}{\partial \sigma} - \frac{\partial \Pi_{\text{fixed}}}{\partial \sigma}.$$

Applying the envelope theorem to the first derivative on the right-hand side and simplifying results in

$$\frac{\partial \hat{f}}{\partial \sigma} = AP^{*^{-\sigma}} \ln P^* \left[-P^* + \phi\right] - A\bar{P}^{-\sigma} \ln \bar{P} \left[-\bar{P} + \phi\right].$$

The sign of this derivative is ambiguous. Intuitively, as long as $\bar{P} > \phi$, profits under the optimal price and the fixed price are both decreasing in σ and it is not obvious which profit function decreases faster. Assuming $\bar{P}$ is five percent below the optimal price, Figure 3 shows how $\hat{f}$ depends on σ and ϕ.[10]

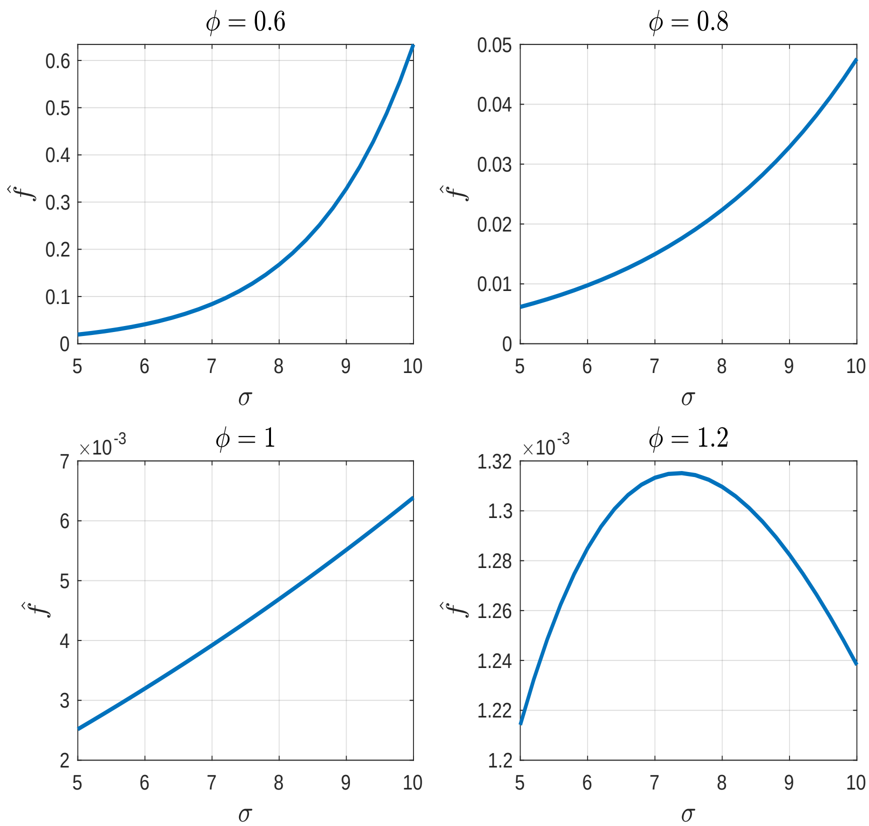

Figure 3: $\bar{P} = 0.95P^*$

Likewise, the Figure 4 shows plot when the $\bar{P}$ is five percent too high.

[10]Because A does not affect the sign of the derivative, we assume $A = 1$ in all of the exercises in this Appendix.

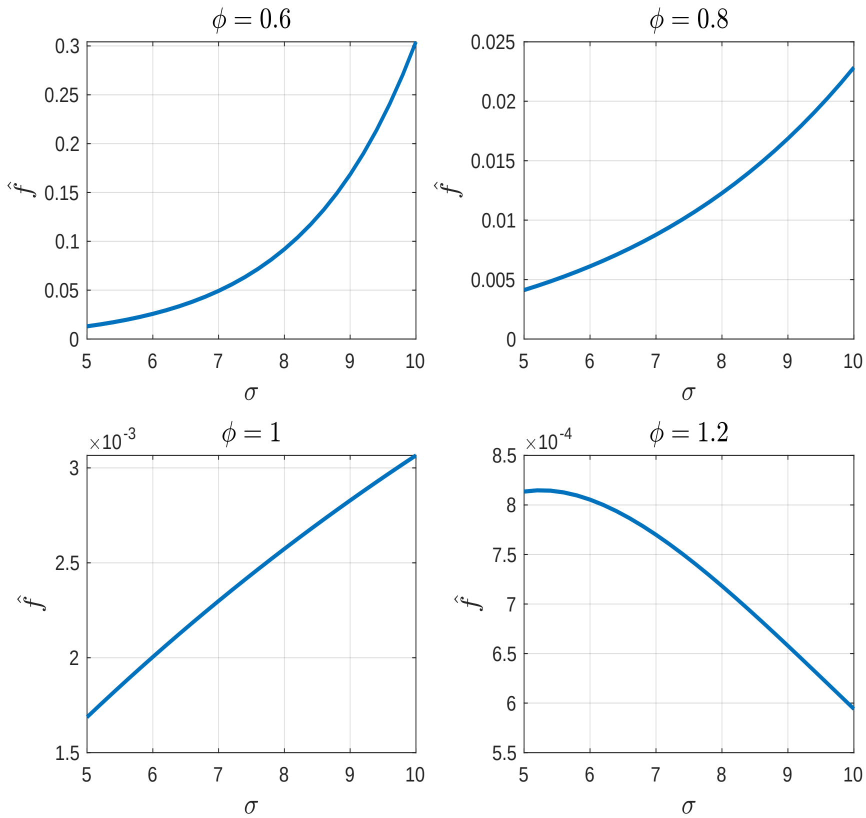

Figure 4: $\bar{P} = 1.05P^*$

We can discipline ϕ to be in a quantitatively relevant range to see the role of σ. When prices are flexible, the profit-to-output ratio reduces to $\phi/(\sigma-1)$. In the data, profit share of GDP is between 5 and 10 percent. The plot below shows how $\hat{f}$ changes when $\phi = 0.5$ over the range of $6 < \sigma < 11$, the empirically relevant range of σ. The profit share in this case is between 5 and 10 percent. In the empirically relevant range for ϕ, Figure 5 shows that $\hat{f}$ increases with σ for several different values of $\bar{P}$ which is consistent with a flatter Phillips curve.

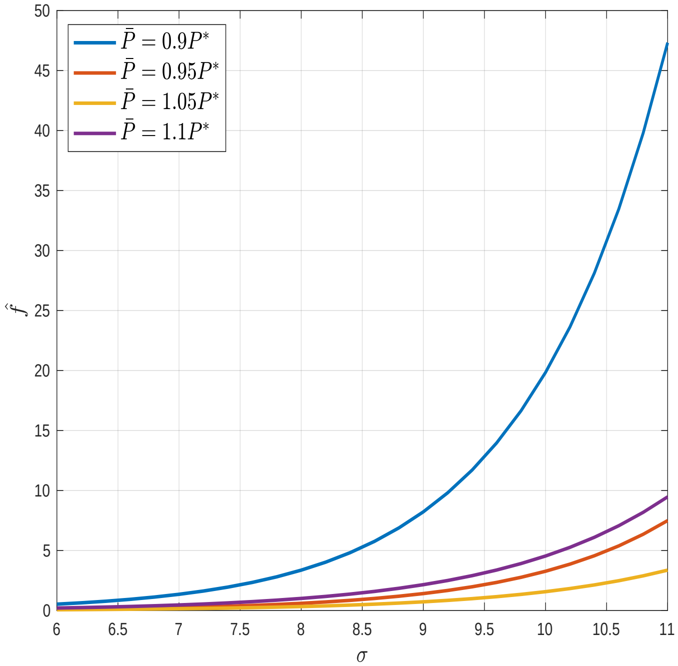

Figure 5: $\hat{f}$ as a Function of σ

C Summary of Data

Table 3 summarizes the NielsenIQ Homescan data sample in the main analysis. Table 4 shows the number of estimated product module σ_ms by age group and by age-income category. In our analysis, we require that at least 20 households purchase a product in each group to be included in the estimation. At younger and lower income households, there are fewer observations that match this criteria and thus there are fewer elasticities estimated within each product module.

Table 3: Summary of Data

Ave. Households per year	57,355
Number of Observations (summed at age-period-barcode)	6,402,134
Number of Product Modules	1,117
Ave. Total Expenditures per year (projection factor weighted)	$314,392,448,666

Table 4: Number of estimated Product Module σ_ms by Group

	by Age	Income and Age lower 50%	Income and Age upper 50%
25-34	378	187	220
35-44	632	383	529
45-54	743	459	667
55-64	768	471	688
65+	742	481	655

C.1 Income groups

One issue with the Homescan data is that incomes are reported in discrete bins at a two-year lag. To estimate relative income, we follow Faber and Fally (2022) to estimate expenditures per capita as a proxy for income group. For this, we obtain per capita expenditures by year regressing log total expenditures on household size dummies and household-level attributes. We then adjust household expenditures by netting our household size dummy coefficients to get per capita expenditure estimates. While expenditures may not accurately measure actual income, they do appear to identify relative income levels. Table 5 regresses the log adjusted per capita expenditures we constructed on income dummies and household size controls for 2004, 2009, 2014, and 2019 (every 5 years in our sample). Note that with only few exceptions, the coefficients on the income dummies are monotonically increasing even though the income bin is reported from 2 years prior. This pattern suggests that the expenditures are monotonically increasing with reported income levels. Thus, estimated expenditures appear to be appropriate for capturing relative income.

Table 5: Per Capita Consumption Estimate and Income Bin

	2004	2009	2014	2019
\$5000-7999	-0.0170	-0.0104	-0.0102	-0.0604
	(0.0355)	(0.0351)	(0.0354)	(0.0367)
\$8000-9999	0.0156	-0.00577	-0.0405	-0.0782*
	(0.0366)	(0.0352)	(0.0315)	(0.0337)
\$10,000-11,999	0.0230	-0.0255	0.0121	-0.0251
	(0.0342)	(0.0323)	(0.0292)	(0.0305)
\$12,000-14,999	0.0598	0.0274	0.00461	-0.0434
	(0.0318)	(0.0291)	(0.0265)	(0.0271)
\$15,000-19,999	0.0971**	0.0421	0.0246	0.00580
	(0.0304)	(0.0276)	(0.0245)	(0.0249)
\$20,000-24,999	0.119***	0.0555*	0.0705**	0.0183
	(0.0299)	(0.0267)	(0.0235)	(0.0236)
\$25,000-29,999	0.149***	0.0837**	0.0913***	0.0524*
	(0.0301)	(0.0266)	(0.0234)	(0.0235)
\$30,000-34,999	0.171***	0.125***	0.109***	0.0711**
	(0.0299)	(0.0264)	(0.0233)	(0.0232)
\$35,000-39,999	0.181***	0.134***	0.146***	0.0991***
	(0.0301)	(0.0266)	(0.0234)	(0.0234)
\$40,000-44,999	0.196***	0.143***	0.147***	0.0753**
	(0.0301)	(0.0266)	(0.0236)	(0.0233)
\$45,000-49,999	0.235***	0.165***	0.165***	0.104***
	(0.0303)	(0.0266)	(0.0233)	(0.0232)
\$50,000-59,999	0.237***	0.193***	0.169***	0.115***
	(0.0296)	(0.0258)	(0.0226)	(0.0223)
\$60,000-69,999	0.248***	0.203***	0.189***	0.131***
	(0.0298)	(0.0261)	(0.0229)	(0.0226)
\$70,000-99,999	0.290***	0.221***	0.199***	0.143***
	(0.0294)	(0.0254)	(0.0221)	(0.0217)
\$100,000 +	0.329***		0.212***	0.177***
	(0.0299)		(0.0223)	(0.0217)
\$100,000 - 124,999		0.249***		
		(0.0262)		
\$125,000 - 149,999		0.224***		
		(0.0289)		
\$150,000-199,999		0.237***		
		(0.0300)		
\$200,000+		0.212***		
		(0.0329)		
Household size controls	Yes	Yes	Yes	Yes
N	39,577	60,502	61,554	61,480

Notes: Standard errors in parentheses. * denotes 10%, ** 5%, and *** 1% significance.

C.2 Elasticity of substitution estimates by year and age

Table 6 gives the estimates of the elasticity of substitution within product modules by each age group and each year in our sample.

Table 6: σ estimates by year and age

	Age				
	25-34	35-44	45-54	55-64	65+
2004	6.380	6.101	6.562	6.246	5.642
	(0.088)	(0.056)	(0.061)	(0.081)	(0.046)
2005	6.388	6.108	6.564	6.244	5.656
	(0.092)	(0.055)	(0.060)	(0.084)	(0.046)
2006	6.351	6.066	6.566	6.230	5.665
	(0.094)	(0.053)	(0.060)	(0.081)	(0.045)
2007	9.720	8.554	8.886	8.562	7.153
	(0.149)	(0.113)	(0.106)	(0.114)	(0.090)
2008	10.745	9.231	9.312	9.124	7.769
	(0.156)	(0.121)	(0.111)	(0.120)	(0.100)
2009	10.740	8.916	9.122	8.917	7.735
	(0.157)	(0.118)	(0.108)	(0.118)	(0.099)
2010	11.155	9.306	9.464	9.238	8.135
	(0.161)	(0.122)	(0.110)	(0.119)	(0.104)
2011	11.509	9.721	9.657	9.551	8.611
	(0.162)	(0.127)	(0.112)	(0.121)	(0.110)
2012	8.029	7.257	7.547	7.276	6.460
	(0.131)	(0.089)	(0.084)	(0.095)	(0.070)
2013	7.730	7.168	7.468	7.263	6.458
	(0.132)	(0.088)	(0.082)	(0.093)	(0.069)
2014	7.584	7.166	7.492	7.344	6.573
	(0.127)	(0.087)	(0.083)	(0.093)	(0.072)
2015	7.614	7.173	7.476	7.378	6.702
	(0.125)	(0.087)	(0.083)	(0.093)	(0.075)
2016	7.459	7.176	7.484	7.399	6.716
	(0.118)	(0.085)	(0.083)	(0.092)	(0.075)
2017	7.486	7.211	7.495	7.469	6.817
	(0.119)	(0.086)	(0.083)	(0.093)	(0.076)
2018	7.584	7.284	7.551	7.490	6.799
	(0.122)	(0.088)	(0.084)	(0.094)	(0.076)
2019	7.697	7.322	7.549	7.497	6.792
	(0.128)	(0.089)	(0.084)	(0.094)	(0.075)

Notes: Standard errors in parentheses.