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: March 27, 2024

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. 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 – namely the 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.

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, and a flatter Phillips curve 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 is 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 consistently exhibits the lowest elasticity among all groups. These findings are particularly stark when we control for income and focus on the

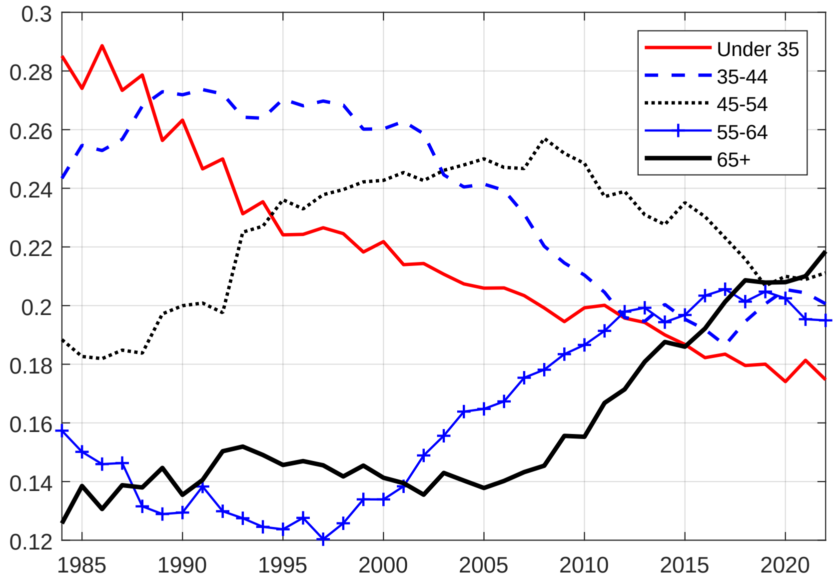

Figure 1: Consumption Expenditure Shares by Age Group

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.

top two quartiles of the income distribution, which, on average, represent over 70% of the expenditure share.

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, 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. Because 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 efficiency 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] Because 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 that highlights the role of age heterogeneity on monetary policy. Empirically, recent work by Juselius and Takáts (2021) documents that demographics affect level inflation in a Phillips curve estimation across countries. 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. 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

[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.

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]

2.1 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

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

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 $\Delta^k ln(p_{mbt}) = \Delta ln(p_{mbt}) - \Delta ln(p_{mkt})$.

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

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

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 do 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]

[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.

[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.

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 young 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.

Figure 2: σ by Age Group

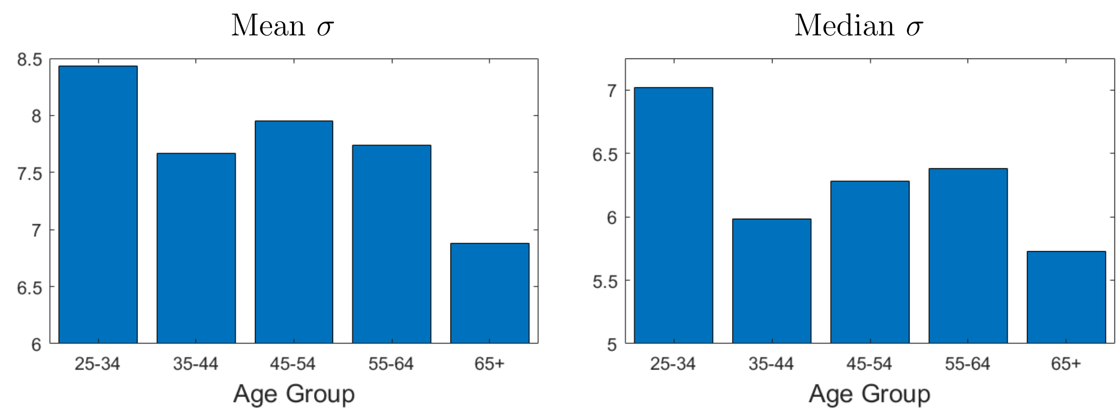

Notes: Mean and median σs are the expenditure weighted values of σ by age over each module. Each product module is required to have at least 20 households purchase in that module (see Online Appendix C for further details). The number of modules with required observations for each age group are: Age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

There is, however, a positive correlation between age and income. A possible concern is that the differences in elasticities by age may be driven by income rather than age itself. To address this, 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

[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.

[6] As in Faber and Fally (2022), we obtain per capita expenditures by year regressing log total expenditures

Table 1: 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
55–65	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.

measure actual income, they do appear to identify relative income levels. In the Online Appendix C, we show per capita expenditures are increasing with (two-year lagged) reported income bins.

Table 1 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 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 in age 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.

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.

3 Aging and the Phillips Curve

In this section we present a textbook model of monopolistic competition with nominal rigidities, augmented to incorporate the empirical findings of the previous section. We start by considering the behavior of two types of households: young, y, and old, o. Let $C_{j,t}$ be the consumption bundle for a group $j = \{o, y\}$ across good i:

$$C_{y,t} = \left(\int_0^1 c_{i,y,t}^{\frac{\sigma_y - 1}{\sigma_y}} di\right)^{\frac{\sigma_y}{\sigma_y - 1}} \quad \text{and} \quad C_{o,t} = \left(\int_0^1 c_{i,o,t}^{\frac{\sigma_o - 1}{\sigma_o}} di\right)^{\frac{\sigma_o}{\sigma_o - 1}}$$

where σ_j is the elasticity of substitution of age group j. Each household type minimizes expenditures subject to the constraint of achieving an overall level of consumption. The optimization problem for the young is

$$\begin{aligned} &\min \int_0^1 p_{i,t} c_{i,y,t} di \\ &\text{s.t.} \left(\int_0^1 c_{i,y,t}^{\frac{\sigma_y - 1}{\sigma_y}} di\right)^{\frac{\sigma_y}{\sigma_y - 1}} \geq C_{y,t}. \end{aligned}$$

The optimization problem for the old household is analogous. In Online Appendix A, we show that the demand for both young and old consumers are given by

$$\begin{aligned} c_{i,y,t} &= C_{y,t} \left(\frac{p_{i,t}}{P_t}\right)^{-\sigma_y} \\ c_{i,o,t} &= C_{o,t} \left(\frac{p_{i,t}}{P_t}\right)^{-\sigma_o}. \end{aligned}$$

Output of good i is produced with the production function $y_{i,t} = z_t n_{i,t}$, where z_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/z_t and, in equilibrium, supply of good i equals demand for good i:

$$y_{i,t} = c_{i,y,t} + c_{i,o,t} = C_{y,t} \left(\frac{p_{i,t}}{P_t}\right)^{-\sigma_y} + C_{o,t} \left(\frac{p_{i,t}}{P_t}\right)^{-\sigma_o}.$$

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]Because we do not model the ownership structure of the firm, we are silent on origins and specification of λ_t.

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

$$y_{i,t} = c_{i,y,t} + c_{i,o,t} = C_{y,t} \left(\frac{p_{i,t}}{P_t} \right)^{-\sigma_y} + C_{o,t} \left(\frac{p_{i,t}}{P_t} \right)^{-\sigma_o}.$$

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

$$(\sigma_o - 1) C_{o,t} + (\sigma_y - 1) C_{y,t} + \phi Y_t \pi_t (1 + \pi_t) = mc_t (\sigma_o C_{o,t} + \sigma_y C_{y,t}) + \mathbb{E}_t \beta \frac{\lambda_{t+1}}{\lambda_t} \phi \pi_{t+1} (1 + \pi_{t+1}) Y_{t+1}. \quad (5)$$

Note that when $\sigma_y = \sigma_o = \sigma$ 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:

$$\pi_t = \frac{s(\sigma_o - 1) + (1 - s)(\sigma_y - 1)}{\phi} \tilde{mc}_t + \beta \pi^e_{t+1} + H_t \quad (6)$$

where s is the share of consumption by the old, and

$$H_t = \frac{(\sigma_o - 1) C_{o,ss} + (\sigma_y - 1) C_{y,ss}}{\phi Y_{ss}} \times \left[\frac{\sigma_o (C_{o,t} - C_{o,ss}) + \sigma_y (C_{y,t} - C_{y,ss})}{\sigma_o C_{o,ss} + \sigma_y C_{y,ss}} - \frac{(\sigma_o - 1)(C_{o,t} - C_{o,ss}) + (\sigma_y - 1)(C_{y,t} - C_{y,ss})}{(\sigma_o - 1) C_{o,ss} + (\sigma_y - 1) C_{y,ss}} \right].$$

If $\sigma_y = \sigma_o = \sigma$ then $H_t = 0$ and this, once again, collapses to the usual log-linearized Phillips curve under Rotemberg pricing,

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

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

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

Since $\sigma_o < \sigma_y$, a higher consumption share for the old (higher s) flattens the Phillips curve.

To gain intuition on 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 $\sigma_y = \sigma_o = \sigma$ gives the standard case, and rearranging

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

where the right-hand side gives the expected benefits minus the costs of time t price adjustments. If the marginal cost increases (i.e. a productivity shock), these costs are scaled by consumption and σ. If σ is high, the demand curve is more sensitive to price adjustments and firms update prices relatively more aggressively. A higher σ results in a larger price change given on the right-hand side – a steeper Phillips curve.

Returning to Equation (5) but reintroducing the consumer heterogeneity

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

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

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.

While we analytically isolate the channel in which heterogeneity in the elasticity of substitution affects the slope of the Phillips curve, we do not provide any quantitative analysis. How these mechanisms play out in a general equilibrium model is an aspect left to future research and our study hints at several pertinent questions. First, to what extent does the aging of the population account for the flattening of the Phillips curve observed empirically? Second, how might the channel explored here affect normative aspects of monetary policy?

References

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.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020): "What's 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.

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.

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. Letting $j = \{y, o\}$

$$\min \int_0^1 p_{i,t} c_{i,j,t} di$$
$$\text{subject to } \left(\int_0^1 c_{i,j,t}^{\frac{\sigma_j - 1}{\sigma_j}} di \right)^{\frac{\sigma_j}{\sigma_j - 1}} \geq C_{j,t}.$$

The first order condition for for good i is

$$p_{i,t} - \lambda_{j,t} c_{i,j,t}^{\frac{-1}{-\sigma_j}} C_{j,t}^{\frac{1}{\sigma_j}} = 0 \Leftrightarrow c_{i,j,t} = C_{j,t} p_{i,t}^{-\sigma_j} \lambda_{j,t}^{\sigma_j}$$

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

$$\left(\int_0^1 c_{i,j,t}^{\frac{\sigma_j - 1}{\sigma_j}} di \right)^{\frac{\sigma_j}{\sigma_j - 1}} = \lambda_{j,t}^{\sigma_j} C_{j,t} \left(\int_0^1 p_{i,t}^{1-\sigma_j} di \right)^{\frac{\sigma_j}{\sigma_j - 1}} = C_{j,t} \Leftrightarrow$$
$$\lambda_{j,t} = \left(\int_0^1 p_{i,t}^{1-\sigma_j} di \right)^{\frac{1}{1-\sigma_j}} = P_{j,t}$$

The demand functions for young and old consumers are given by

$$c_{i,y,t} = C_{y,t} \left(\frac{p_{i,t}}{P_t} \right)^{-\sigma_y}$$
$$c_{i,o,t} = C_{o,t} \left(\frac{p_{i,t}}{P_t} \right)^{-\sigma_o}.$$

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

$$y_{i,t} = c_{i,y,t} + c_{i,o,t} = C_{y,t}\left(\frac{p_{i,t}}{P_t}\right)^{-\sigma_y} + C_{o,t}\left(\frac{p_{i,t}}{P_t}\right)^{-\sigma_o}.$$

Substituting the constraint into the objective and taking the first-order condition with respect to $p_{i,t}$ gives

$$(\sigma_y - 1)C_{y,t}p_{i,t}^{-\sigma_y}P_t^{\sigma_y-1} + (\sigma_o - 1)C_{o,t}p_{i,t}^{-\sigma_o}P_t^{\sigma_o-1} + \phi\frac{Y_t}{p_{i,t-1}}\left(\frac{p_{i,t}}{p_{i,t-1}} - 1\right)$$
$$= mc_t\left(\sigma_o C_{o,t}p_{i,t}^{-\sigma_o-1}P_t^{\sigma_o} + \sigma_y C_{y,t}p_{i,t}^{-\sigma_y-1}P_t^{\sigma_y}\right) + \mathbb{E}_t\beta\frac{\lambda_{t+1}}{\lambda_t}\phi Y_{t+1}\frac{p_{i,t+1}}{p_{i,t}^2}\left(\frac{p_{i,t+1}}{p_{i,t}} - 1\right).$$

The only thing that involves i is $p_{i,t}$ so every firm is going to choose the same price. Note, this verifies out earlier assumption. Let $p_{i,t} = P_t$. Inflation is $\frac{P_t}{P_{t-1}} = 1 + \pi_t$. The (non-linear) Phillips curve is:

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

If $\sigma_y = \sigma_o$ then this is the standard Rotemberg Phillips curve.

A.1 Rotemberg Pricing: Log-linear approximation

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

$$Y_t = C_{y,t} + C_{o,t} + \frac{\phi}{2}\pi_t^2 Y_t.$$

In steady state,

$$Y_{ss} = C_{o,ss} + C_{y,ss}.$$

Let s be the share of output consumed by the old people in steady state. So

$$C_{o,ss} = sY_{ss} \tag{7}$$
$$C_{y,ss} = (1 - s)Y_{ss}. \tag{8}$$

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

$$mc_{ss} = \frac{(\sigma_o - 1)C_{o,ss} + (\sigma_y - 1)C_{y,ss}}{\sigma_o C_{o,ss} + \sigma_y C_{y,ss}}.$$

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

$$\begin{aligned}&\ln\left[(\sigma_o-1)C_{o,t}+(\sigma_y-1)C_{y,t}+\phi Y_t\pi_t(1+\pi_t)\right]\\ &=\ln\left[mc_t(\sigma_o C_{o,t}+\sigma_y C_{y,t})+\mathbb{E}_t\beta\frac{\lambda_{t+1}}{\lambda_t}\phi\pi_{t+1}(1+\pi_{t+1})Y_{t+1}\right].\end{aligned}$$

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

$$\begin{aligned}\frac{(\sigma_o-1)(C_{o,t}-C_{o,ss})+(\sigma_y-1)(C_{y,t}-C_{y,ss})+\phi Y_{ss}\pi_t}{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}=&\frac{(mc_t-mc_{ss})(\sigma_o C_{o,ss}+\sigma_y C_{y,ss})}{mc_{ss}\sigma_o C_{o,ss}+mc_{ss}\sigma_y C_{y,ss}}\\ &+\frac{mc_{ss}\sigma_o(C_{o,t}-C_{o,ss})+mc_{ss}\sigma_y(C_{y,t}-C_{y,ss})}{mc_{ss}\sigma_o C_{o,ss}+mc_{ss}\sigma_y C_{y,ss}}\\ &+\frac{\beta\phi Y_{ss}\pi^e_{t+1}}{mc_{ss}\sigma_o C_{o,ss}+mc_{ss}\sigma_y C_{y,ss}}.\end{aligned}$$

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

$$\frac{\phi Y_{ss}\pi_t}{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}+A_t=\tilde{mc}_t+B_t+\frac{\beta\phi Y_{ss}\pi^e_{t+1}}{mc_{ss}\sigma_o C_{o,ss}+mc_{ss}\sigma_y C_{y,ss}}$$

where $\tilde{mc}_t$ is marginal cost's percent deviation from steady state.[8] Using the steady-state condition for marginal cost, we can write this as

$$\frac{\phi Y_{ss}\pi_t}{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}+A_t=\tilde{mc}_t+B_t+\frac{\beta\phi Y_{ss}\pi^e_{t+1}}{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}.$$

Doing some cross multiplication yields

$$\pi_t=\frac{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}{\phi Y_{ss}}\tilde{mc}_t+\beta\pi^e_{t+1}+H_t$$

where $H_t=\frac{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}{\phi Y_{ss}}(B_t-A_t)$. Recall the convention introduced that defined $C_{o,ss}=sY_{ss}$ and $C_{y,ss}=(1-s)Y_{ss}$.

Writing the Phillips curve this way gives

$$\begin{aligned}\pi_t&=\frac{(\sigma_o-1)sY_{ss}+(\sigma_y-1)(1-s)Y_{ss}}{\phi Y_{ss}}\tilde{mc}_t+\beta\pi^e_{t+1}+H_t\Leftrightarrow\\ \pi_t&=\frac{s(\sigma_o-1)+(1-s)(\sigma_y-1)}{\phi}\tilde{mc}_t+\beta\pi^e_{t+1}+H_t.\end{aligned}$$

[8] A and B are given by $A_t=\frac{(\sigma_o-1)(C_{o,t}-C_{o,ss})+(\sigma_y-1)(C_{y,t}-C_{y,ss})}{(\sigma_o-1)C_{o,ss}+(\sigma_y-1)C_{y,ss}}$ and $B_t=\frac{\sigma_o(C_{o,t}-C_{o,ss})+\sigma_y(C_{y,t}-C_{y,ss})}{\sigma_o C_{o,ss}+\sigma_y C_{y,ss}}$.

If $\sigma_y = \sigma_o = \sigma$ and the old and young consumptions are identical, $H_t = 0$, 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

$$\frac{s(\sigma_o - 1) + (1-s)(\sigma_y - 1)}{\phi}. \tag{9}$$

Assuming $\sigma_o < \sigma_y$ as we find in the data, a higher share of old people flattens 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.[9]

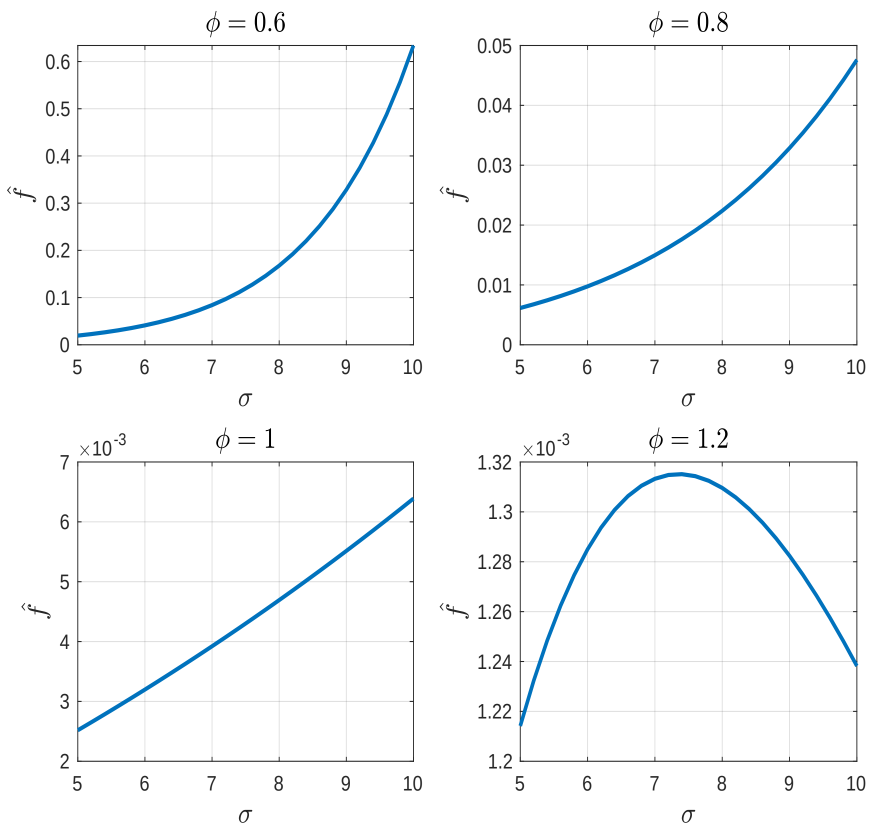

Figure 3: $\bar{P} = 0.95P^*$

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

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

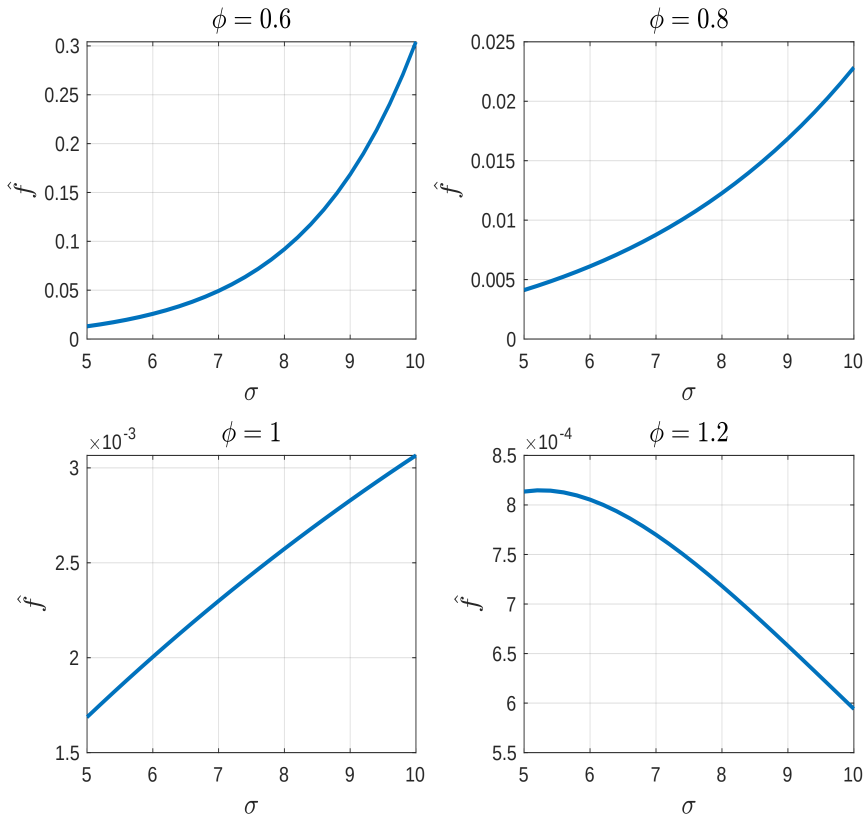

Figure 4: $\bar{P} = 1.05P^*$

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.

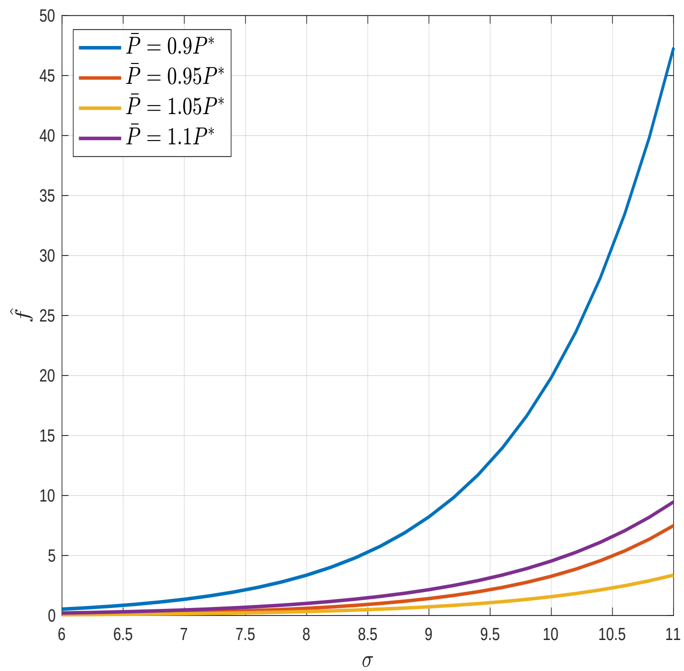

Figure 5: $\hat{f}$ as a Function of σ

C Summary of Data

Table 2 summarizes the NielsenIQ Homescan data sample in the main analysis. Table 3 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 2: 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 3: 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 4 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 4: 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 5 gives the estimates of the elasticity of substitution within product modules by each age group and each year in our sample.

Table 5: 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.