




BANK FOR INTERNATIONAL SETTLEMENTS



BIS Working Papers

No 886

Price search, consumption inequality, and expenditure inequality over the life-cycle

by Yavuz Arslan, Bulent Guler and Temel Taskin

Monetary and Economic Department

September 2020

JEL classification: D10, D91, E21.

Keywords: Consumption inequality, price search, incomplete markets, life-cycle models, partial insurance.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2020. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Price Search, Consumption Inequality, and Expenditure Inequality over the Life-Cycle*

Yavuz Arslan[†]

Bulent Guler[‡]

Temel Taskin[§]

August 2020

Abstract

In this paper, we differentiate consumption from expenditure by incorporating price search decision into an otherwise standard life-cycle model. With this extension, households can pay lower prices for the same consumption good if they allocate more time for price search. We first analytically show that under very general conditions, poorer (in both income and/or wealth) households search more and pay lower prices compared with wealthier ones. As a result, consumption inequality is smaller than expenditure inequality, and the gap between them increases over the life-cycle. Next, we quantify these mechanisms by calibrating our model to the US data. We find that the life-cycle increase in consumption inequality is about 30 percent lower than the increase in expenditure inequality. We also show that the availability of a price search option provides a large insurance mechanism against adverse income shocks and increases the welfare of a new born by 3.9 percent in consumption equivalent terms.

Keywords: Consumption inequality, price search, incomplete markets, life-cycle models, partial insurance.

J.E.L. Classification: D10, D91, E21.

*We would like to thank Arpad Abraham, Mark Aguiar, Mark Bilal, Yongsung Chang, Fatih Guvenen, Tarik Kara, Fatih Karahan, Burhanettin Kuruscu, Cedric Okou, Serdar Ozkan, Fabrizio Perri, seminar participants at the University of Rochester, the European University Institute, the Central Bank of the Republic of Turkey, the Midwest Macro Meeting in Bloomington, Indiana, the Society for Economic Dynamics Meeting in Istanbul, the Royal Economic Society Meeting in London, and the Latin American Meeting of the Econometric Society in Santiago, Chile. The authors acknowledge the Indiana University Pervasive Technology Institute for providing HPC resources that have contributed to the research results reported within this paper. All errors are our own. The views expressed here are those of the authors and do not reflect those of the World Bank or the Bank for International Settlements.

[†]Bank for International Settlements. Email: yavuz.arslan@bis.org

[‡]Indiana University. Email: bguler@iu.edu

[§]World Bank. Email: ttaskin@worldbank.org

1 Introduction

Most of the earlier literature that has studied the life-cycle dynamics of consumption inequality implicitly assumed that consumption and expenditure are equivalent.¹ On the empirical side, it has been common to use expenditure data to measure consumption. On the theoretical side, the price of consumption was assumed to be 1. However, there is a growing literature that documents a significant dispersion in prices paid for identical goods that implies a difference between consumption and expenditure.² Our hypothesis is that to the extent that prices paid covaries with age, wage and wealth, consumption and expenditure will diverge over the life-cycle.

To quantify the significance of this price search channel, we use a standard life-cycle model with incomplete markets and endogenous labor supply. We assume that households face idiosyncratic wage shocks. However, they can insure themselves against these shocks through a risk-free asset and flexible labor supply. Different from the standard models, we also allow households to search for lower prices to mitigate the effects of adverse shocks. When households search more, they pay less and consume more. However, they enjoy less leisure due to their time constraint.

First, we analytically show that when labor supply is exogenous or zero, price search is a decreasing function of wealth. When labor supply is endogenous, price search becomes negatively correlated with wage. We then analyze the effect of price search on expenditure inequality and consumption inequality. The model allows us to get a closed-form expression for the gap between expenditure inequality and consumption inequality. This gap crucially depends on the return to price search technology, wage inequality, and the covariance of consumption and wage, which are both positive and increasing over the life-cycle. This implies that the gap between expenditure inequality and consumption inequality widens over the life-cycle.

We calibrate the benchmark model to the U.S. economy to quantify the magnitude of the difference between consumption and expenditure. Our results show that the increase in the variance of consumption is 30 percent lower than the increase in the variance of expenditure over the life-cycle. The rise in the covariance of wage and consumption over the life-cycle can explain about 85 percent of this difference. The rest is due to the rise in the variance of wages.

We show that the rise in the variance of prices and search effort implied by the model is sup-

¹In this literature, studies usually estimate the life-cycle profile of inequality, and conclude that it rises around 30 log points over the life-cycle. A partial list of studies in this line of research includes Deaton & Paxson (1994), Blundell & Preston (1998), Gourinchas & Parker (2002), Storesletten et al. (2004), Krueger & Perri (2006), Guvenen (2007), Blundell et al. (2008), Kaplan & Violante (2010), Heathcote et al. (2014), and Aguiar & Bils (2015).

²Baye et al. (2006) provide a detailed survey of the dispersion in prices paid for identical goods. Aguiar & Hurst (2007) show that richer people pay higher prices for identical goods in the U.S. Also, they report that prices paid for identical goods change over the life-cycle, which is a result of a change in price search due to a change in the cost of time. Similar results are also shown by Kaplan & Menzio (2015). Using the U.S. data, Sorensen (2000) reports a significant dispersion in prices paid for the same medicine. Dahlbay & West (1986) show that there is substantial price dispersion among automobile insurance companies in Canada. Pratt et al. (1979) document price dispersion for several categories of goods. Woodward & Hall (2012) find that mortgage borrowers loose at least \$1000 in brokerage fees by visiting too few brokers. Kaplan & Menzio (2016) show that the existence of price search can result in self-fulfilling unemployment fluctuations.

ported by data. We use A.C. Nielsen data and show that the variance of price paid is small (compared to the variance of consumption and expenditure) and increases over the life-cycle, as the model predicts. This data set does not have a perfectly comparable variable for the search effort. Therefore, we use proxies, such as shopping frequencies, number of different stores visited, trips per stores, and coupon usage. We show that the variances of all proxies increase over the life-cycle.

Our findings suggest that the link between average income and average consumption growth is significantly weaker than the one between income and expenditure (Carroll & Summers, 1989; and Guvenen, 2007). The reason is that households with higher wage growth will pay increasingly higher prices over their life-cycle. As a result, their expenditure will grow faster than their consumption. To quantify the mechanism, we compare two groups of households, one with a 1 percent higher annual wage growth rate than the benchmark economy and the other with a 1 percent lower wage growth rate than the benchmark economy over the life-cycle. We show that the difference between the average consumption differential at the age 65 is 15 percent lower than the expenditure differential.

In addition to its impact on consumption and expenditure, price search enables consumers to insure themselves against stochastic income shocks. We follow the formulation of Blundell et al. (2008) and Kaplan & Violante (2010), and quantify the partial insurance role of price search by computing the consumption insurance coefficients with respect to transitory and persistent shocks. The computed insurance coefficient of consumption for persistent shocks is 29 percentage points higher in the model with price search compared to the one without price search. The same coefficient is 10 percentage points higher for transitory shocks. As suggested by higher insurance coefficients, price search improves the welfare of the households. A comparison of the benchmark economy and the one with no price dispersion shows that the new-born households gain 3.9 percent in consumption equivalent terms by choosing the benchmark economy.

Surprisingly, our results suggest that the availability of price search increases asset holdings (by 14 percent). Absent any other mechanisms one would expect that, as price search improves insurance, it would reduce the need to hold a large amount of assets. However, price search modifies the return dynamics of the risk-free bond and transforms it into a more favorable asset. Intuitively, a household would search more and pay lower prices when he/she receives a low income shock. Given an asset level, he/she can afford more consumption with lower prices. Therefore, price search generates idiosyncratic return dynamics for the bond that has a negative correlation with his/her income, which increases the demand for the asset. This mechanism outweighs the decrease in asset demand due to improved insurance, and asset holdings increase in equilibrium.

Our results are robust to several perturbations of the model, but there are two parameters that deserve discussion. First, the return to search effort seems to change our results significantly. In the benchmark, we use the estimates provided by Aguiar & Hurst (2007) to calibrate this parameter. However, Aguiar & Hurst (2007) use micro data on grocery store transactions. Given that

grocery shopping composes a limited part of consumption bundle and that the return to search could vary across categories, the value of this parameter requires further attention.³ Second, in our model, time cost of search effort increases with the size of the consumption bundle. The benchmark calibration assumes a linear relationship between the two to be on the conservative side, whereas some empirical studies (e.g., Stroebel & Vavra, 2019) suggest a convex one without a point estimate pertinent to our case. Given the lack of empirical consensus, we solve the model with alternative values of these two parameters for robustness.

The findings we present in this paper have several implications. First, they suggest that consumption inequality is not as high as expenditure inequality. However, to have a more precise estimate of consumption inequality, it is necessary to analyze price and quantity data jointly for the full set of goods. Second, a better access to search technologies, such as internet, would improve price search options and lower consumption inequality further. Third, the increase in internet-based sales during the last two decades would again lower consumption inequality.

Related Literature Among many other studies in the quantitative life-cycle literature, our paper is closely related to Guvenen (2007), Storesletten et al. (2004), and Karahan & Ozkan (2010). Those papers study the role of income processes in the life-cycle profile of consumption inequality. Kaplan (2012) extends a similar model with unemployment risk to better match life-cycle profiles of variance of consumption and labor allocations. There is a common implicit assumption in these models that the price of a consumption good is unique, and therefore consumption is equal to expenditure. However, as we mentioned above, there is also a large empirical literature that rejects this assumption. Our paper differs from standard life-cycle studies in that it differentiates consumption from expenditure.

In a contemporaneous study, Pytka (2018) follows a similar strategy to ours and studies the search and consumption decisions jointly. While in his framework price dispersion is endogenous, we take the distribution of retail prices as given. A distinguishing feature of our framework is that labor supply is endogenous. We consider it to be an important margin when studying the life-cycle dynamics of prices and consumption.

The paper continues as follows. In Section 2, we present the model and our theoretical analysis. In Section 3, we explain the details of the calibration strategy. We discuss our results in Section 4. Section 5 presents the role of price search as an insurance mechanism. We report the robustness results in Section 6 and conclude in Section 7.

³For example, one can argue that the return to search is higher for online shopping, as consumers can check prices without visiting stores.

2 Model

Our model is an extension of a standard life-cycle consumption-saving problem with endogenous labor supply and price search technology. The price search technology allows households to search for lower prices and partially insure themselves against adverse income shocks.⁴ The asset markets are incomplete due to uninsurable idiosyncratic wage shocks. The population consists of a continuum of households who work for T periods and afterwards enjoy retirement until period T^* . We study the age-inequality profiles of consumption and expenditure in this environment. Each component of the model is explained in detail below.

2.1 Households

In each period, households make two decisions: (1) the consumption/saving decision, and (2) the time allocation between leisure, labor supply and the price search decision. Households can enjoy higher consumption by searching for lower prices in exchange of less time for leisure and labor supply. They maximize the expected lifetime value of discounted utility:

$$E \sum_{t=0}^{T^*} \beta^t u(c_t, l_t),$$

subject to the budget and borrowing constraints:

$$\begin{aligned} p(s_t)c_t + a_{t+1} &= w_t n_t + (1+r)a_t, \\ a_{t+1} &\geq 0, \end{aligned}$$

where, $u(\cdot)$ is periodic utility, β is time discount factor, c_t is consumption, l_t is leisure time for the household at period t , a_{t+1} is saving, and w_t is the wage of the household at age t . We use a utility function similar to the one used in Kaplan (2012), which is specified as follows:

$$u(c_t, l_t) = \frac{c_t^{1-\sigma}}{1-\sigma} - \phi_t \frac{(1-l_t)^{1+\gamma}}{1+\gamma},$$

where σ is risk aversion and γ is the inverse of the Frisch elasticity.⁵ We assume that the disutility from labor, ϕ_t , evolves with experience.⁶

⁴See Becker (1965) and Becker and Ghez (1975) for a benchmark theory on the role of time allocation over the life-cycle.

⁵To be more precise, due to the presence of price search, model implied Frisch elasticity is not equal to $1/\gamma$. However, we computed the implied Frisch elasticity in the model with the benchmark calibration and find it to be very close to $1/\gamma$. Despite this caveat, for expositional reasons, we refer to γ as the measure of Frisch elasticity in the remainder of the paper.

⁶We make this assumption to match the model implied life-cycle labor supply dynamics with the data.

The main departure of our model from standard life-cycle models is that we distinguish consumption from expenditure. Expenditure is the product of price paid and the size of the consumption bundle (i.e., $e_t = p_t c_t$, where c_t is consumption, and e_t is expenditure, and p_t is the price of the consumption good). In standard models, p is normalized to 1 since every other price is relative to the price of the consumption good.

We assume that households can lower the price of their consumption bundle by searching for lower prices. We follow Aguiar & Hurst (2007) and assume that price and search effort have a loglinear and negative relationship:

$$\log(p_t) = \theta_0^t + \theta_1 \log(s_t), \quad (1)$$

with $\theta_1 < 0$.^{7,8}

This particular price technology implies that doubling the search effort lowers the price paid by θ_1 percent. While a higher search effort decreases prices paid, it is costly. Specifically, we assume that it takes time effort to find lower prices. A household with one unit/variety of consumption that puts effort s to search prices needs to dedicate s amount of her time in a period. Extending this to a whole consumption basket, a household with a consumption level of c_t spends $s_t c_t^\psi$ of her time for searching. Here, ψ measures how much search time increases with the size of the consumption bundle. Together with the time spent on working, this implies the following time constraint on households:

$$s_t c_t^\psi + n_t + l_t = 1. \quad (2)$$

Modeling time cost of price search in this way has several advantages. First, as the variety of goods increases with the size of the consumption bundle (Li, 2013), it is more realistic to assume that one should devote additional time for each variety to get lower quotes. For example, one should search among grocery stores to pay lower prices for groceries and search electronics stores to pay lower prices for computers. Second, if we were to assume that the time cost of search is independent of the size of the consumption basket, the model would have unrealistic implications. In that case, as consumption increases, price search will also increase. The reason is that with the same amount of price search, households will lower the price of a larger consumption bundle. Hence, they will have a higher return on price search.

The functional form of time cost of price search used in this paper also generates consistent results with the empirical counterparts documented in Pytka (2018): among employed, higher earners spend more time on shopping. However, there is an important difference. While Pytka

⁷Notice that we assume that θ_0^t is age dependent. This assumption is useful in matching the life-cycle profiles of average prices paid.

⁸See Appendix for a justification of this reduced form functional form.

(2018) assumes that higher shopping effort has a direct utility cost, in our specification, it increases the time cost. Therefore, it lowers the available time for both labor and leisure. Given that labor supply decision is endogenous in our model, we find the current specification–time cost rather than utility cost for search effort– to be a better fit for our question.

We follow the literature for the evolution of the wage, w_t , over the life-cycle. In each period, a household receives both persistent and transitory wage shocks. The log wage follows:

$$\log(w_t) = \beta_0 + \beta_1 t + \beta_2 t^2 + z_t + \epsilon_t \text{ where } \epsilon_t \sim N(0, \sigma_\epsilon^2),$$

where β_0 is a scale parameter, β_1 is the return to experience, t is age, z_t is the persistent, and ϵ_t is the transitory components of the wage shock. The persistent wage shocks follow an AR(1) process where ρ measures persistence:

$$z_t = \rho z_{t-1} + v_t, \text{ where } z_0 = 0 \text{ and } v_t \sim N(0, \sigma_v^2).$$

2.2 Theoretical Analysis

To get some insight of the mechanisms in the model through the lens of closed form theoretical relationships, we focus on the static part of the model where a household chooses the optimal level of labor supply and price search given the current wealth and the next period saving choice. We denote y as the change in wealth: $y_t = a_t(1+r) - a_{t+1}$.

2.2.1 Price Search

We first analyze how search effort changes with wage and wealth. We study positive and zero labor supply cases separately.

Positive labor supply When households have wealth below a threshold and/or do not have a large disutility from working, they choose to have a positive labor supply. The following lemma characterizes the price search as a function of wealth and wage when labor supply is positive.

Lemma 1 *When labor supply is endogenous and positive, the equation characterizing price search becomes:*

$$p'(s) = -wc^{\psi-1}. \quad (3)$$

If consumption is a normal good and $p''(s) > 0$, then we have:

$$\frac{ds}{dy} \begin{cases} > 0 & \text{if } \psi < 1 \\ = 0 & \text{if } \psi = 1 \\ < 0 & \text{if } \psi > 1. \end{cases}$$

With positive labor supply, the effect of wealth on price search depends critically on the parameter, ψ , which determines how the size of consumption bundle affects the time cost of price search. The marginal benefit of price search is the product of marginal change in price, consumption, and the marginal utility of consumption, and the marginal cost of price search is the product of marginal utility of leisure and marginal increase in time used for price search. When wealth increases, through the usual wealth effect, consumption increases and marginal utility of consumption decreases. The drop in the marginal utility of consumption induces a drop in the marginal utility of leisure through a decrease in labor supply. When labor supply is endogenous, these two effects change proportionally. When $\psi = 1$, since both the marginal cost and the marginal benefit of price search are proportional to the change in consumption and both marginal utility of consumption and marginal utility of leisure change proportionally, price search becomes independent of wealth. However, if $\psi < 1$, time cost of price search becomes smaller, and the marginal cost of price search increases less. Hence, price search becomes an increasing function of wealth. On the other hand, if $\psi > 1$, the increase in the time cost of higher price search becomes more costly, and price search becomes a decreasing function of wealth. It is important to emphasize that this result does not depend on the functional form of the preferences and price search technology.

Given equation 1, we can also express the price search explicitly as a function of wage and consumption:

$$\log(s) = \frac{\log(-\theta_1) + \theta_0 - \log(w) - (\psi - 1)\log(c)}{1 - \theta_1}. \quad (4)$$

It is possible to infer several insights from equation 4. When $\psi \geq 1$, equation 4 implies a loglinear negative relationship between search and wage, the elasticity being inversely related to the return of price search, θ_1 . The intuition is that when the labor supply is positive, the opportunity cost of search time is foregone earnings. As a result, an increase in wages lowers search time. At the same time, search effort becomes less sensitive to the changes in wages when the return on search, θ_1 , increases. In addition, when $\psi > 1$, an increase in consumption lowers search time, as it increases the time cost of search.

Zero labor supply Next, we analyze the case in which there is no labor supply, $n = 0$. This corresponds to the cases in which households have sufficiently large wealth, are retired, and/or have a large disutility from labor. In this case, it is possible to solve for price search explicitly as a function of fundamental parameters of the model:

$$\log s = A - \frac{\sigma - 1 + \psi(1 + \gamma)}{1 + \gamma - \theta_1(\sigma - 1 + \psi(1 + \gamma))} \log(y), \quad (5)$$

where $A = \frac{\theta_0(\sigma - 1 + \psi(1 + \gamma)) - \log(\phi(\psi - \frac{1}{\theta_1}))}{1 + \gamma - \theta_1(\sigma - 1 + \psi(1 + \gamma))}$.

Equation 5 implies that when labor supply is zero and the return to price search is positive (i.e., $\theta_1 < 0$), then price search decreases with wealth (i.e., $\frac{ds}{dy} < 0$ if $\psi \geq 1$ or $\sigma > 1$ and $\psi \geq 0$). There are three reasons for this. First, there is the usual wealth effect; as wealth increases, leisure increases and that, in turn, decreases price search. Second, there is the substitution effect; a decrease in price search increases price paid and results in lower consumption. When $\psi = 0$ and $\sigma = 1$, these two effects cancel out, and price search becomes independent of wealth. However, when $\sigma > 1$, wealth effect dominates, and price search decreases as wealth increases. Third, when $\psi > 0$, an increase in wealth increases consumption, which raises the time cost of price search. Consequently, price search declines. When $\psi \geq 1$, price search becomes a decreasing function of wealth as long as $\sigma \geq 0$.

Equation 5 provides further insights about the importance of some of the key parameters of the utility function. For instance, when the risk aversion, σ , increases, the elasticity of search to wealth increases. Similarly, a rise in γ , the inverse of Frisch elasticity, lowers this elasticity. One implication is that these parameters can potentially be important in differentiating consumption from expenditure over the life-cycle. We will illustrate the significance of these parameters in the robustness analysis. To proceed in our theoretical analysis, we assume that $\psi = 1$. In the calibration section, we discuss our parameter choice, and in the robustness section, we do several analyses to determine the quantitative importance of this parameter.

2.2.2 Expenditure vs. Consumption Inequality

After deriving the optimality conditions for price search, we can now discuss the effects of price search on consumption and expenditure inequality. Since expenditure is equal to price times consumption, it is possible to express the connection between expenditure and consumption inequality as follows:

$$Var(\log e) = Var(\log c) + Var(\log p) + 2Cov(\log p, \log c).$$

This relationship suggests that the difference between expenditure inequality and consumption inequality depends on the variance of log prices paid and the covariance of log prices and log consumption. We next analyze these terms for exogenous and endogenous labor supply cases separately.

Endogenous Labor Supply When labor supply is endogenous and $\psi = 1$, the relation between consumption and expenditure inequality becomes:

$$Var(\log e) = \left(\frac{\theta_1}{1 - \theta_1} \right)^2 Var(\log w) + Var(\log c) - 2 \frac{\theta_1}{1 - \theta_1} Cov(\log c, \log w). \quad (6)$$

The statistical properties of the estimated wage process imply that $Var(\log w)$ is an increasing function over the life-cycle. Similarly, the covariance of consumption and wages, $Cov(\log c, \log w)$, is positive and increases over the life-cycle. As a result, the difference between $Var(\log e)$ and $Var(\log c)$ should be positive and increasing over the life-cycle.

Zero Labor Supply In this relatively simpler case, it is possible to derive an analytical expression defining the relation between expenditure and consumption inequality. When $\psi = 1$, expenditure inequality, defined by the variance of log expenditures, becomes a constant proportion of consumption inequality:

$$Var(\log e) = \left(1 - \theta_1 \frac{\sigma + \gamma}{1 + \gamma}\right)^2 Var(\log c). \quad (7)$$

This finding suggests that when labor supply is zero, the variance of log expenditures is higher than the variance of log consumption as long as the return to price search is positive (i.e., $\theta_1 < 0$). This result is an immediate consequence of equation 5, from which we know that price search is a decreasing function of wealth. That in turn implies that wealthier households (or, equally, high consumption households) pay higher prices compared to poorer ones. This mechanism increases the difference between expenditure levels when compared to consumption, which makes the expenditure variance larger. Equation 7 also implies that as risk aversion, σ , and the Frisch elasticity, $1/\gamma$, increase, the gap between expenditure and consumption inequality widens. Moreover, as the return to price search increases (i.e., absolute value of θ_1 , increases), again, the gap between expenditure and consumption inequality increases. We will return to these observations in Section 6 where we conduct several robustness analyses with respect to important parameters of the model.

Taking Stock The model shows that price search decreases as wealth increases when labor supply is zero. That is, wealthier households search less for the prices and pay more for the same consumption bundle. When labor supply is positive and $\psi = 1$, we find that price search is independent of wealth but decreases with wage. This also implies that the gap between expenditure and consumption inequalities is mainly driven by the evolution of wage and the covariance of consumption and wages, both of which increase over the life-cycle in standard models. Given these theoretical observations, in the following sections we aim to quantify the difference between expenditure and consumption inequalities using a calibrated version of the life-cycle model.

3 Calibration

We calibrate the model in two stages.

External Calibration In the first stage, we directly use the values of certain parameters that are well established in the related literature. This gives us an opportunity to understand the role of price search in standard life-cycle models. The model period is set to one year, and each household starts working at age 21, retires at age 65, and dies at age 90. Each household starts her/his working life with zero assets. We set $r = 0.03$. The value of the relative risk-aversion parameter for consumption, σ , is set to 2, a standard value in the literature. The value of parameter γ , which pins down the intratemporal elasticity of substitution for leisure, is usually set greater than the value of σ in the literature. Therefore, we set the value of γ to 4, which implies Frisch elasticity of 0.25, consistent with the estimates reported by Chetty et al. (2011). We borrow the parameters of the deterministic and stochastic component of the wage process - $\beta_0, \beta_1, \beta_2, \rho, \sigma_v^2, \sigma_\epsilon^2$ from Kaplan (2012). For the pension process, we follow Guvenen (2007), which mimics the U.S. Social Security system. After retirement, the pension of each household is determined by the ratio of his wage in the last working period to the average wage in the last working period.

We set the return on price search parameter, $\theta_1 = -0.19$, as reported in Aguiar & Hurst (2007).⁹ One concern with this estimate is that it reflects the return on price search for only grocery store products. However, we have a composite consumption good in the model, and the corresponding return may be different. We do a robustness analysis to address some of these concerns.

An important parameter that we do not have precise information for is ψ . It has strong implications about how the time cost of search effort changes with the size of the consumption bundle. While we do not have point estimates, there is strong evidence that it is larger than 1. Stroebel & Vavra (2019) document from a large panel of households that homeowners reduce their search effort (i.e., the use of coupons and sales) when house prices increase. In our model, this happens when ψ is larger than 1 (Lemma 1). For the quantitative analysis, we choose $\psi=1$. This choice is on the prudent side, as larger values amplify our results (see the robustness analysis).

Internal Calibration In the second stage, we calibrate the remaining parameters to match particular moments in the data. Two specific assumptions we make needs further discussion. First, we assume that the constant in the price search technology, θ_0 , and the disutility of time cost, ϕ , change deterministically over the working period. We impose these assumptions to better match the life-cycle evolution of prices paid and the labor supply we observe in the data. These assumptions can be rationalized due to several factors, such as experience, access to better technologies, or higher time endowment due to marriage. We should note that the model with constant price search and disutility parameters produces almost identical results for the evolution of consumption and expenditure inequality. However, it generates an increasing price paid and decreasing labor supply over the life-cycle.¹⁰

⁹Appendix provides further justification for this parameter supported by theory and empirical evidence on price dispersion reported by Kaplan & Menzio (2015).

¹⁰See Section 6 for details.

Table 1: Benchmark Model Parameters

Parameter	Explanation	Value
Externally Set		
r	Risk-free interest rate (U.S. data)	0.03
σ	Intertemporal elasticity of substitution for consumption (various sources)	2
γ	Intertemporal elasticity of substitution for leisure (various sources)	4
θ_1	Return to price search	-0.19
ψ	Elasticity of search time wrt. the size of consumption	1
β_0	Constant of the wage process	0.4
β_1	Return to experience	0.035
β_2	Return to experience-squared	-0.00042
ρ	Persistence of wage process	0.958
σ_v^2	Variance of persistent wage shock	0.017
σ_ϵ^2	Variance of transitory wage shock	0.081
Internally Calibrated		
β	Discount factor	0.967
ϕ^1	Disutility-age 1	372.1
ϕ^T	Disutility-age T	72.0
ϕ^P	Disutility-curvature	7.56
θ_0^1	Price search constant-age 1	-0.267
θ_0^T	Price search constant-age T	-0.444
θ_0^P	Price search constant-curvature	2.25

We chose to match the life-cycle prices paid to avoid any confusion that may result from unrealistic life-cycle price dynamics. To keep the number of parameters small, we assume that θ_0 and ϕ follow a monotone path over the life-cycle and represent these functions with only three parameters that govern the curvature, and the levels at the beginning and end of the working stage of the household. More specifically, we impose the following functional form for these parameters as a function of age:

$$x_t = \begin{cases} x_1 + (x_T - x_1) \left(\frac{t-1}{T-1} \right)^{x^p} & \text{if } t \leq T \\ x_T & \text{if } t > T, \end{cases}$$

where x_1 is the level of the parameter at the beginning of the period, x_T is the level at the time of retirement, and x^p is the curvature parameter.¹¹ This gives us six parameters, $\phi^1, \phi^T, \phi^p, \theta_0^1$, and θ_0^T, θ_0^p , to be calibrated. We assume that after retirement these parameters stay constant at their corresponding levels at the time of retirement.

In total, we jointly calibrate seven parameters, $\beta, \phi^1, \phi^T, \phi^p, \theta_0^1, \theta_0^T, \theta_0^p$ to match the following moments: the wealth-to-income ratio (3.5) as documented in Heathcote et al (2010), the average price paid (1.0) as normalization, the life-cycle profiles of hours worked and prices paid, as documented in Kaplan & Menzio (2015).^{12,13}

Price Index We follow Aguiar & Hurst (2007) to define an average price measure for each household. Specifically, let the total expenditure during month m is defined as follows:¹⁴

$$X_m^j = \sum_{i \in I, t \in m} p_{i,t}^j q_{i,t}^j, \quad (8)$$

where $p_{i,t}^j$ is the price of good i purchased by household j on shopping trip t , and $q_{i,t}^j$ is the corresponding quantity. Price paid to the same good at a given time might vary across households. Therefore, average price paid for a given good during a month is calculated by taking the weighted (by quantity purchased) average across households:

$$\bar{p}_{i,m} = \sum_{j \in J, t \in m} p_{i,t}^j q_{i,t}^j / \bar{q}_{i,m}, \quad (9)$$

where the denominator of the weights is calculated as follows:

¹¹The term x corresponds to either θ_0 or ϕ . This particular form restricts x to be a monotonic and either convex ($x^p > 1$) or concave function ($x^p < 1$) of t . x_1 and x_T correspond to the two end points of x .

¹²We construct five-year averages of hours worked and price paid for the life-cycle (before retirement) and target these moments. This will result in nine separate moments for hours-worked, and nine separate moments for price-paid.

¹³Following Aguiar & Hurst (2007), the life-cycle profile of prices reflects estimated coefficients of age-group dummy variables in a regression in which log prices are regressed on age-group dummy variables, several observables, and shopping needs.

¹⁴See Aguiar & Hurst (2007) for further details about this price index. An alternative price index, defined as the average of the ratio of price to average price for each good, is highly correlated (0.8) with the benchmark price index.

$$\bar{q}_{i,m} = \sum_{j \in J, t \in m} q_{i,t}^j. \quad (10)$$

Next we calculate the monthly expenditures of a household at average price for the consumed bundle as follows:

$$Q_m^j = \sum_{i \in I, t \in m} \bar{p}_{i,t} q_{i,t}^j. \quad (11)$$

Then, the price index for household j is computed as the ratio of her/his actual expenditure to the cost of his/her consumption bundle at average prices:

$$\bar{p}_m^j = X_m^j / Q_m^j, \quad (12)$$

which is normalized by the average price index across households within the month:

$$p_m^j = \frac{\bar{p}_m^j}{\frac{1}{J} \sum_{k \in J} \bar{p}_m^k}, \quad (13)$$

equating average price to unity. If a household's price index is higher (lower) than 1, s/he pays higher (lower) than average to the identical consumption bundles.

An important feature of this formulation is that the differences in price index between households do not reflect differences in the quality of goods purchased. Variation in prices is measured at UPC (universal product code) level, which is unique for each product.


Calibration Results Table 2 presents the model-generated moments against the data counterparts. The model matches the targeted moments. The left panel of Figure 1 shows the life-cycle profile of average prices paid by households in the model. This pattern is close to the estimates documented in Aguiar & Hurst (2007) and Kaplan & Menzio (2015).

The right panel of Figure 1 shows the life-cycle profile of average hours worked. The model-generated hours worked follows the empirical counterpart closely and matches the peak at age group 41-45. The drop in hours worked toward retirement in the model falls slightly short of the empirical counterpart. Factors, such as early retirement options, can help reconcile the model output with the empirical data; however, we find the model's current match sufficient and do not pursue further improvement.

Table 2: Benchmark Model-Moments

Moment	Data	Model
Wealth-to-income ratio	3.5	3.5
Average price paid	1.0	1.0

Figure 1: Price Paid and Hours over the Life-Cycle: Model vs. Data



Notes: The figure on the left panel plots the life-cycle profile of average log prices paid both in the model and the data. The figure on the right panel plots the life-cycle profile of average log hours worked in the model and the data. Both figures show the deviation from their average level at age group 21-25.

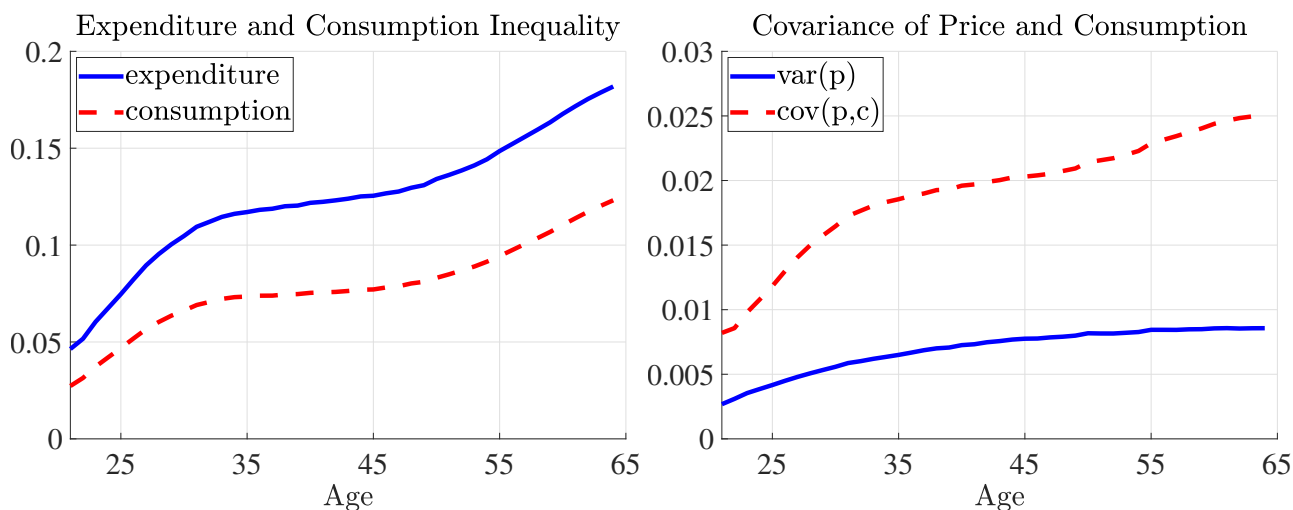
4 Results

As we showed in the theoretical section, the presence of price search can potentially generate a difference between consumption and expenditure inequality profiles for a given saving and the labor supply decision. However, price search can also alter the asset accumulation and labor supply decisions of households, since it can serve as an additional insurance mechanism for adverse shocks. In this section, we quantitatively evaluate the significance of price search.

4.1 Age-Inequality Profiles of Consumption and Expenditure

Our model predicts a substantially lower consumption inequality than expenditure inequality throughout the life-cycle. As Figure 2 shows, the cross-sectional variance of log expenditure increases by about 13.6 log points over the life-cycle, whereas the variance of log consumption increases by about 9.5 log points. This is close to a 30 percent reduction in the rising inequality over the life-cycle.

Figure 2: Expenditure and Consumption Inequality



Notes: The figure on the left plots the life-cycle profile of consumption and expenditure inequality measured by the variance of log of the variables. The figure on the right decomposes the expenditure inequality into two components: the variance of price paid and the covariance between consumption and price paid.

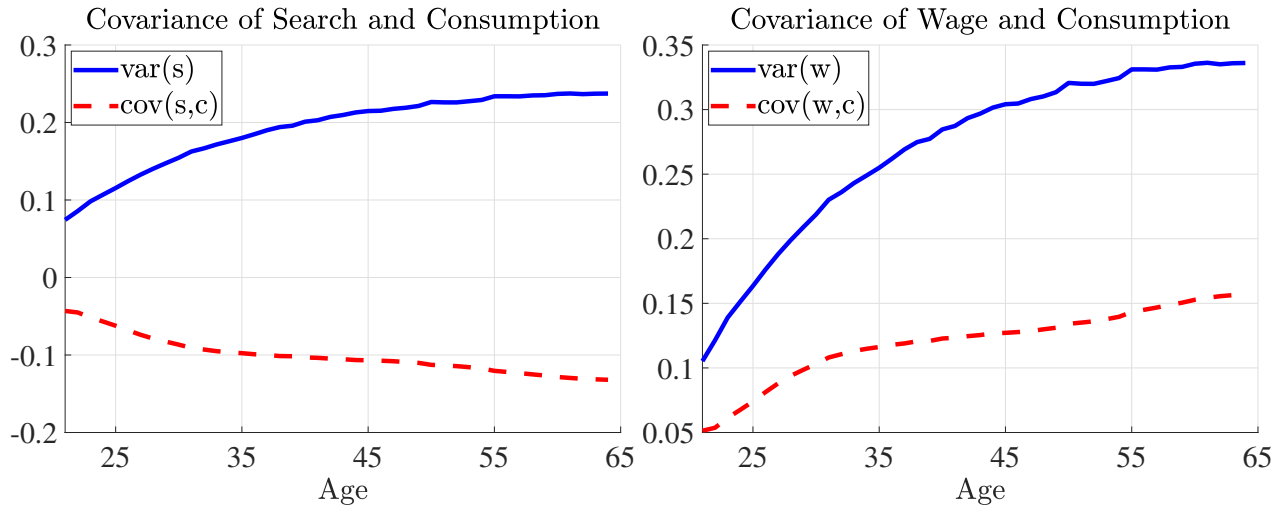
To understand the drivers of the gap between consumption and expenditure inequalities throughout the life-cycle, we decompose the variance of expenditure as follows:

$$var(\log e) = var(\log c) + var(\log p) + 2cov(\log c, \log p).$$

We calculate each component of $var(\log e)$ from the model-generated data. As can be seen from the right panel of Figure 2, the bulk of the difference between expenditure and consumption inequality is due to the increase in the covariance between consumption and price paid. The variance of price paid almost triples over the life-cycle, but the level is very small. As it starts from a low level, it contributes around 15 percent to the increase in the difference between expenditure and consumption inequality. These findings suggest that the gap between expenditure and consumption inequality arises primarily due to the increase in the covariance of price paid and consumption.

It is also possible to understand the dynamics of the gap between expenditure and consumption inequality from the perspective of the search behavior. Figure 3 plots the life-cycle profile of the covariance between consumption and price search. As would be expected, it is negative and declining over the life-cycle. This is because average wealth and wealth inequality increase with age. The number of people who decrease their search as their wealth increases is larger among older people (Figure 3).

Figure 3: **Decomposition of Expenditure Inequality**



Notes: The figure on the left plots the life-cycle profile of the variance of log price search and covariance of log of price search and log of consumption. The figure on the right plots the life-cycle profile of the variance of log of wage and covariance of log of wage and log of consumption.

Lastly, in the theory section, we showed that the variance of log expenditure can be decomposed into the variance of log consumption, the variance of log wages, and the covariance of log wages and log consumption. The right panel in Figure 3 shows the evolution of wage inequality and the covariance of wages and consumption over the life-cycle. Wage inequality substantially increases over the life-cycle, similar to the data. On top of that, the covariance of wages and consumption is also positive and increases significantly over the life-cycle. However, given the

calibration of $\theta_1 = -0.19$, the coefficient of the contribution of the change in wage inequality is multiplied by a small number, $\left(\frac{\theta_1}{1-\theta_1}\right)^2 = 0.03$, whereas the change in the covariance of the wage and consumption is multiplied by a much larger number, $-2\frac{\theta_1}{1-\theta_1} = 0.32$. Our results reveal that, in particular, the covariance of consumption and wages is much more important in the evolution of the gap between consumption and expenditure inequality. We find that the contribution of the increase in the covariance of consumption and wage on the gap between expenditure inequality and consumption inequality is 85 percent, whereas the contribution of the increase in the variance of wage is 15 percent.


Price and search inequality over the life-cycle in the data We document the life-cycle profile of inequality in prices and search using the A.C. Nielsen Homescan data. We follow Aguiar & Hurst (2007) to create a household level price index using the shopping transactions. The price index is described in detail in Section 3. Inequality is measured as the standard deviation of log price after controlling for shopping needs, specifically the number of product categories, number of UPCs, and the size of the consumption basket (i.e., expenditures at average prices, Q). The left panel of Figure 4 illustrates the increased life-cycle profile of inequality in prices paid across households. We use a similar approach to measure the life cycle profile of inequality in price search in the data.

The model implied variance of price and search behavior are in line with empirical data.¹⁵ Data suggests that the cross-sectional variance of prices paid increases over the life-cycle (the top-left panel of Figure 4). Despite the rise and similar to the model, the variation in price paid is smaller compared to the expenditure inequality. As a result, it can explain only a small part of the rise in expenditure inequality.

Individual prices are observed in the scanner data, however search time and effort is only available as proxies. In Figure 4, we report the dynamics of several of them. For instance, the variance of shopping frequency increases until 35-39 and then has a declining trend. The cross-sectional variance of trips to different stores and trips per store keeps rising over the life-cycle (middle panel of Figure 4). All these search measures are obtained after controlling for shopping needs and, therefore, can be interpreted as shopping for lower prices. Similarly, the increase in the cross-sectional variance of coupon usage (arguably a less noisy measure of price search) over the life-cycle also supports the model output on price search. Taken altogether, data confirms model implications that the variance of time allocated to price search increases over the life-cycle (right panel of Figure 4).

¹⁵As the level of search and prices paid are already documented in the literature (Aguiar & Hurst, 2007; Kaplan & Menzio, 2015) we do not report them here.

Figure 4: Standard Deviation of Prices and Shopping Behavior over the Life-Cycle



Notes: These figures plot the variance of prices and search behavior over the life-cycle. We use AC Nielsen data from Aguiar and Hurst (2007). See the text for details.

4.2 Price Search, Expenditure and Consumption Growth over the Life-Cycle

While our main focus is on the dynamics of inequality, price search can also alter the life-cycle dynamics of average expenditure and consumption. The mechanism is that, as households get older, their wages and wealth increase, both of which make them pay higher prices. On the other hand, for the numerical analysis, we calibrated the search technology to match the changes in the life-cycle prices paid. Our results suggest that consumption grows by about 4 percent more than expenditure over the life-cycle in our baseline economy (Figure 5).


The difference between expenditure and consumption increases if one considers households with different wage growth patterns. Facing the same search technology, households with higher income growth will pay higher prices as their income and wealth increases. As a result, their expenditure will grow faster than their consumption. Our results depicted in Figure 5 suggest that this may be quantitatively important. To check this hypothesis, we solve the benchmark model with higher wage growth and lower wage growth profiles. In the benchmark economy, wage grows by 2 percent annually. In the high growth counterfactual, we assume wage grows by 3 percent annually, and in the low wage growth counterfactual we assume wage grows by 1 percent annually.

We find that, for high wage growth households, expenditure growth is about 4 percent greater than their consumption growth over the life-cycle. However, for low wage growth households, consumption grows 11 percent more than expenditure over the life-cycle. This creates more than a 15 percent difference in expenditure and consumption growth rates.

These findings may bring a partial explanation to the empirical observation that consumption growth parallels income growth in the data (Carroll & Summers, 1989; Guvenen, 2007). Since the empirical literature uses expenditure data rather than consumption, a systematic deviation of prices paid across different income growth patterns will cause an overestimation. Our results suggest that consumption growth still follows income growth, but relationship is weaker. As a

result, consumption growth patterns of different income growth patterns do not diverge as much as the expenditure (the left and middle panels of Figure 5).

Figure 5: Average Income, Expenditure and Consumption Growth over the Life-Cycle



Notes: The left and middle panels of the graph plots the dynamics of the deviation of the log expenditure and consumption from the levels at age 21 for different income growth economies. The average annual wage increase is 2 percent in the benchmark economy, 3 percent for the “high income growth” case, and 1 percent for the “low income growth” case. The right panel plots the difference between average expenditure and consumption.


4.3 The Determinants of Search Behavior

Wage, wealth, and age are the key drivers of search and prices paid in the model. Earlier in our theoretical analysis we showed that price search should be negatively correlated with wage. With respect to wealth, we showed that, when $\psi = 1$ and labor supply is positive, price search depends only on wage and is independent of wealth. However, when the labor supply becomes zero, which happens when wealth is sufficiently large, price search becomes negatively correlated with wealth. Figure 6 confirms these results. Consistent with the predictions of the theoretical model, the left panel of Figure 6 shows that, as wealth increases, price search initially stays constant. However, after a sufficiently large wealth, when the critical level of wealth depends on wage and age, it starts to decrease. The turning point coincides with the wealth level when the labor supply becomes zero. Figure 6 also shows the effect of the wage on price search.¹⁶ A higher wage implies lower price search, confirming our theoretical results. Another important observation is that, as wealth increases price search for different income groups converges to a low level. Therefore, the role of wealth becomes the dominant determinant of price search as households become wealthier.

The right panel of Figure 6 plots search effort as a function of wealth and age. It shows that as a household ages, price search declines. This is mainly due to the increasing wage profile over the life-cycle.

¹⁶Wages correspond to different levels of persistent shock. The plot with transitory shocks generates the same qualitative results.

Figure 6: The Determinants of Price Search



Notes: The figure on the left panel plots search time for different wealth and income groups. The figure on the right panel plots search time as function of wealth and age. In both figures, transitory shock is set to the median level. The figure on the left is for a household at the age of 45, and the figure on the right is for a household with median level of the persistent shock.

5 Price Search, Self-Insurance and Welfare

5.1 Quantifying the Partial Insurance with Price Search

Consumers can insure themselves against income risk in several ways. The most common method assumed in the literature is to accumulate precautionary savings and use labor supply to insure against adverse shocks. Our framework provides price search as an additional channel for insurance. In this section, we compute the size of partial insurance obtained through price search.

In order to quantify the partial insurance role of price search, we solve the model without price search. To separate the effect of price search, we recalibrate all the parameters of the model except the ones governing the price search technology to match the same statistics as in the benchmark calibration.¹⁷

Following Blundell et al. (2008) and Kaplan and Violante (2010), we calculate the insurance coefficients as follows:

$$\phi^\epsilon = 1 - \frac{\text{cov}(\Delta c_{it}, \epsilon_{it})}{\text{var}(\epsilon_{it})}, \quad (14)$$

where the insurance coefficient of shock ϵ is denoted by ϕ^ϵ . In the complete markets framework, $\phi^\epsilon = 1$ due to the fact that idiosyncratic shocks do not affect consumption, that is, $\text{cov}(\Delta c_{it}, \epsilon_{it}) = 0$. In an autarky environment, consumption is one-to-one mapped to idiosyncratic shocks, and thus $\phi^\epsilon = 0$. Under incomplete markets, one would expect the insurance coefficients to be between 0

¹⁷This calibration results $\beta = 0.962$, $\phi_1 = 819$, $\phi^T = 176$, and $\phi^P = 7.96$.

Table 3: Insurance Coefficients

	w/ Price Search		w/out Price Search	
	Expenditure	Consumption	Expenditure	Consumption
Transitory	.75	.91	.83	.83
Persistent	.38	.54	.42	.42

Notes: The insurance coefficients are calculated using equation 14 for the entire population. See the text for details

and 1.

Using U.S. micro-level consumption data, Blundell et al. (2008) estimate the insurance coefficients as 0.36 and 0.95 for permanent and transitory shocks, respectively.¹⁸ Kaplan & Violante (2010) generate similar coefficients in a simulated economy using a life-cycle model with incomplete asset markets, inelastic labor supply, and idiosyncratic income shocks.

We use equation 14 and first compute the insurance coefficients for the entire population in our benchmark economy. As shown in Table 3, in the case of no price search, the insurance coefficients are 0.42 (persistent shock) and 0.83 (transitory shock) for both consumption and expenditure.¹⁹ With price search, the corresponding insurance coefficient for consumption increases to 0.54 for the persistent shock and 0.91 for the transitory shock. The corresponding coefficients for expenditure are 0.38 and 0.75, respectively.²⁰ In the model with price search, the consumption insurance coefficients of persistent and transitory shocks are improved by 0.12 (28.6 percent) and 0.08 (9.6 percent), respectively. The main reason for such an increase in insurance coefficients is that price search can serve as an additional mechanism to insure against adverse shocks. In response to adverse wage shocks, not only can households use assets to avoid a drop in consumption, they can also increase price search and still be able to smooth consumption at the expense of leisure.

Next, we study the life-cycle properties of insurance coefficients. Overall, our findings confirm the findings of Kaplan & Violante (2010) and suggest that insurance coefficients rise over the life-cycle for both transitory and persistent income shocks (Figure 7). For the transitory income shocks, insurance coefficients are large and peak much earlier than the permanent shocks. The insurance coefficients for permanent shocks keep rising over the life-cycle.

Price search shifts insurance coefficients upwards for all ages. For instance, for the persistent shock case (Figure 7, the right panel) insurance coefficients for consumption and expenditure are 0.36 and 0.20 (a difference of 0.16), respectively. As consumers get older, the difference stays


¹⁸They use the Panel Study of Income Dynamics (PSID) and Consumer Expenditure Survey (CEX) data. See Blundell et al. (2008) for details.

¹⁹Note that consumption in this case is equal to expenditure by definition.

²⁰Assuming zero borrowing limits, Kaplan & Violante (2010) compute insurance coefficients for consumption in the range of 0.27-0.42 (depending on the persistence of the shocks) for persistent shocks and 0.81-0.82 for transitory shocks.

stable, but the relative contribution of price search to insurance declines. The reason for this is that when consumers are young, price search is the primary way that they can insure themselves, as they do not have much wealth. As they get older, they accumulate wealth, which then becomes a more important insurance mechanism.

Figure 7: Insurance via Price Search over the Life-Cycle



Notes: For each age group, we calculate the insurance coefficients: $\phi^\epsilon = 1 - \frac{\text{cov}(\Delta c_{it}, \epsilon_{it})}{\text{var}(\epsilon_{it})}$, for both expenditure and consumption. The left and the right panels plot the insurance coefficients for the transitory and permanent income shocks, respectively.

5.2 Price Search and Welfare

To study the welfare implications of price search, we solve the benchmark model by closing the price search channel. However, completely eliminating the price search from the model, as we did in Section 5.1, would not be a correct comparison for the role of price search. This is because, in that model without price search, there is no disutility of search, and, thus, households unambiguously get a higher welfare. Instead, we solve the benchmark economy by assuming each household searches in such a way that they all pay the same price of 1, which is the average price paid in the benchmark economy. Comparing the welfare of this economy and the benchmark economy, we find that new-born households gain by 3.9 percent in consumption equivalent terms in the benchmark economy, as opposed to the counterfactual without price search.

5.3 Price Search, Savings and Labor Supply

As shown above, the opportunity of price search is a very powerful mechanism to insure against income shocks. Therefore, one would expect that its availability should reduce the demand for

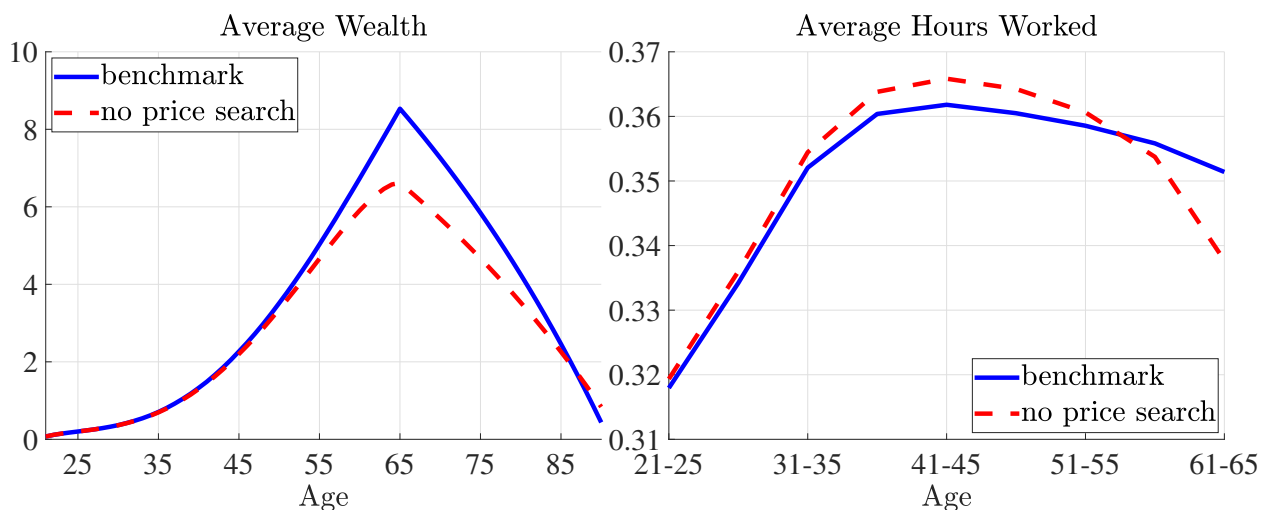
savings. However, there is an additional mechanism that works in the other direction. Consider a household with a low-income shock with some savings. With price search the household has the opportunity to pay a lower price for consumption. That means that, with the same amount of assets, he/she can buy more consumption. This corresponds to a higher return on savings during an economically bad time. As a result, the availability of the price search option makes the risk-free bond a more favorable asset. Similarly, it would be expected to affect labor supply decisions, as price search can also be seen as non-market activity that may reduce average price paid.

To quantify these mechanisms, we solve a variant of the model in which we close the price search opportunity while preserving the benchmark calibration. For this exercise, as shown in Section 5.2, we fix the search amount of each household to make them pay a price of 1, which is the average price paid in the benchmark economy. As a result, households in both economies, on average, pay the same prices.

Figure 8 plots the effect of price search on savings and labor supply. The availability of price search option increases saving through out the life-cycle (Figure 8, left panel). Comparing both economies, average wealth holdings decline around 14 percent once we close the price search option. These results suggest that the second mechanism mentioned above dominates.

The price search option affects labor supply decisions as expected. Once we close the price search flexibility, average hours increase, by about 0.5 percent (Figure 8, right panel).

Figure 8: The Effects of Price Search on the Saving and Labor Supply Behaviors




Notes: The figures compare the dynamics of wealth and labor supply with (benchmark) and without the price search option. The “no price search” case is calibrated to have an average price paid of 1, the same as in the benchmark economy.

6 Robustness Analysis

In this section, we provide several robustness analyses with respect to the parameters of price and utility functions. In all of these exercises, we recalibrate the discount factor (β), the disutility of labor parameters (ϕ^1, ϕ^T, ϕ^p), and parameters related to price search technology ($\theta_0^1, \theta_0^T, \theta_0^p$) to match the wealth-income ratio, moments related to hours worked and price paid over the life-cycle as in the benchmark economy.²¹

Figure 9: **Robustness: Expenditure Variance/Consumption Variance**



Notes: The figures plot various robustness analyses for the difference between expenditure and consumption inequality over the life-cycle. For all the exercises the parameters are recalibrated to match the same targets as in the benchmark economy. Both expenditure and consumption inequality are measured by the variance of log expenditure and consumption.

²¹The parameters calibrated for each case can be found in the Appendix.

Risk Aversion (σ) In the models with incomplete markets, risk aversion is an important parameter that determines the desire for consumption smoothing. It is well-known within these models that as risk aversion increases consumption inequality decreases. However, the same mechanism also decreases the expenditure inequality. As such, it is not clear what happens to the difference between expenditure and consumption inequality. That said, equation 6 implies that an increase in the risk aversion lowers the difference between expenditure and consumption inequality through its effect on the covariance of consumption and wages. Higher risk aversion makes households less responsive to wage shocks, which makes the covariance of consumption and wage drop.

The upper left panel of Figures 9 illustrates the role of risk aversion on the difference between expenditure and consumption inequality. The results confirm the intuition mentioned above: the difference decreases as risk aversion increases. However, the difference persists to be large that suggests that the role of price search remains quantitatively important across different risk aversion parameters.

Frisch Elasticity ($1/\gamma$) Frisch elasticity, which approximately corresponds to $\frac{1}{\gamma}$ in our model, determines how flexible the labor supply is in response to wage fluctuations. In turn, this flexibility will have an influence on the insurance role of the labor supply so as to avoid the fluctuations in consumption generated by exogenous and stochastic movements in wages. As Frisch elasticity increases, labor supply becomes more elastic, and, as a result, the labor supply can play a more significant role in insuring against wage fluctuations. Therefore, higher Frisch elasticity should imply lower consumption and expenditure inequalities, again through its effect on the covariance between consumption and wage. Higher elasticity makes consumption less responsive to wage shocks since the labor supply can play a more significant role in insuring against wage shocks.

The upper right panel of Figures 9 shows the quantitative results. The results suggest that a decrease in γ , which corresponds to an increase in the Frisch elasticity, decreases the difference between expenditure and consumption inequality. More importantly, and for the purposes of this exercise, the difference remains large and changes little across different Frisch elasticity parameters.

Return to search (θ_1) If $\theta_1 = 0$, there would be no return to price search, no one would search for prices, and everyone would pay the same price. In that case, consumption and expenditure would be equal to each other. This reasoning implies that return to search effort is crucially important for the results. For the quantitative exercise, we calibrate this parameter from the estimates provided in Aguiar & Hurst (2007). However, these estimates are obtained using data for grocery store items, not for the whole consumption basket. Therefore, it may not be perfectly fit for our analysis.

For instance, the return on price search may be even higher for non-grocery store items. One reason is that non-grocery store items are more likely to be sold online, where the return on search is probably larger.²² For example, Baye et al. (2004) show that the exact same item is offered at different prices on different websites. More importantly, the price dispersion is large. For instance, a consumer could pay about 12.5 percent less than the average price if he searches for the lowest price. Similarly, consumers can significantly lower their prices, by about 20-47 percent, by buying larger sized items of the same product (Nevo & Wong, 2018). On the other hand, it is possible that some market structures, e.g. monopolistic markets, may limit the returns to price search. For instance, there is not much price dispersion in local energy markets, and, as a result, there is limited room for consumers to search and pay less.

We perform a robustness analysis by solving the model with three different values of θ_1 . The results suggest quantitatively strong effects of return on price search on the difference between expenditure and consumption inequality (the lower left panel of Figure 9). As equation 6 predicts, as the return on price search rises the difference between consumption inequality increases.

These results also suggest that, in sectors or goods where there may be higher returns to search, the discrepancy between expenditure and consumption inequality would be higher. Similarly, when search technology improves and/or firms increase their online presence, as happened during the last 20 years, this discrepancy may again become larger. In addition, if, for example, different segments of the population, such as older people, do not have access to better search technologies, then expenditure and consumption inequalities may become different across different groups. In this paper, we choose not to delve deeper into these potentially important extensions.

The elasticity of search time wrt. the size of consumption (ψ) The parameter ψ measures how the time cost of price search increases with consumption. In the benchmark model, we assumed that ψ is 1. Our justification was that as consumption increases the variety of goods consumed also increases. Indeed, as mentioned earlier, Li (2013) shows that the variety increases almost linearly for food consumption. If consumers dedicate separate times for different goods, then it is reasonable to expect that ψ is close to 1. However, one can come up with other scenarios in which ψ is higher than 1. For example, if the variety of goods consumed increases it may be difficult to search for each variety separately if there are fixed costs. In addition, empirical evidence in Stroebe & Vavra (2019) suggest that ψ is larger than 1.

With these concerns in mind, we find it useful to do a robustness analysis for three different values of ψ : 1 (our baseline parameter), 0.5 and 1.5. Our results shows that ψ has large effects (the lower right panel of Figure 9). Expenditure inequality increases more and consumption inequality increases less for higher ψ . Combining these two effects results in an even larger role of ψ on the difference between consumption and expenditure inequality.

²²A consumer will not need to incur the cost of travel, for instance.

The model allows us to derive an analytical expression for the relation between expenditure and consumption inequality as a function of wage inequality, covariance of wage and consumption, and two crucial parameters of the model, θ_1 and ψ :

$$\begin{aligned} Var(log e) = & \left(\frac{\theta_1}{1-\theta_1} \right)^2 Var(log w) + \left(1 + \left(\theta_1 \frac{\psi-1}{1-\theta_1} \right)^2 - \frac{2\theta_1(\psi-1)}{1-\theta_1} \right) Var(log c) \\ & + 2 \left((\psi-1) \left(\frac{\theta_1}{1-\theta_1} \right)^2 - \frac{\theta_1}{1-\theta_1} \right) Cov(log w, log c). \end{aligned}$$


As ψ gets larger, the role of the covariance of consumption and wage decreases, whereas the pass through from consumption inequality to expenditure inequality increases. The intuition is as follows. As ψ gets larger, it becomes costlier to search for higher prices for wealthier and higher income groups, as they are the ones who can consume more. Therefore, households with more wealth and higher income will spend less time on price search and will consume less compared to lower ψ levels. As a result, the level of consumption of the upper part of consumption distribution shifts down, making the consumption inequality smaller compared to expenditure inequality. The same reasoning also implies larger expenditure inequality, as prices paid by higher wealth and income groups will increase with ψ . Our analysis of ψ reveals that, to precisely quantify the role of price search on expenditure and consumption inequality, the value of ψ is crucial.

Indivisible Labor In our baseline economy, we assume flexible labor where households can adjust their time to work, search, and enjoy leisure to maximize their utility. However, in practice, this is not necessarily the case. Very often, workers cannot adjust their hours in a way that suits them the best.²³ In this section, to explore the consequences of our flexible labor formulation, we assume that households only have three labor supply options: work full time (40 percent of time endowment), part time (20 percent of time endowment) or not work at all. The rest of the time allocation problem stays the same as the benchmark. With these indivisible labor supply choices, we recalibrate the model with the same targets that we use in our benchmark analysis.

Our main conclusion from this analysis is that the way we model labor supply does not alter the results. Figure 10 plots our findings. The left panel shows the mean hours over the life-cycle. Other than the early periods of working life, the model matches the labor supply behavior in the data: the mean hours worked peaks much earlier than the baseline economy. The middle panel shows that despite differing life-cycle dynamics, both models imply very similar increases in the difference between consumption and expenditure inequalities. The right-panel shows that the covariance of consumption and wages, an important determinant of the difference between consumption and expenditure, is almost the same in both cases. We conclude that our results are robust to the flexible labor supply modeling choice.

²³See Hansen (1985) for an early reference for the role of indivisible labor on the business cycle dynamics.


Figure 10: Expenditure and Consumption Inequality if Labor Supply were Inflexible



Notes: The figures plot the results of a version of the model in which the labor supply is indivisible. Households have three options: work for full time (40 percent of time endowment), part time (20 percent of time endowment) or not work at all.

Age-Dependent Parameters Our model assumes that disutility of time-spent outside of leisure, ϕ , and the price-search return parameter, θ_0 , are both age dependent. We impose these assumptions to better match the life-cycle patterns for hours worked and price paid. To test whether these assumptions have any effect on our main results, we recalibrate the model with constant disutility parameter and constant price-search return parameter, separately. Figure 11 shows the comparison of these two counterfactual with the benchmark results. In both cases, the difference between expenditure inequality and consumption inequality has no significant change from the benchmark economy. However, with a constant disutility parameter, as shown in Figure 11 (middle panel), the model predicts a monotonically decreasing labor supply over the life-cycle. Similarly, when we assume a constant θ_0 over the life-cycle, price paid counterfactually increases with the household age, as seen in Figure 11 (right panel). These results show that our assumptions about these age-dependent parameters allow us to capture the life-cycle dynamics of hours worked and price paid without any significant effects on our main results of the paper.

Figure 11: Significance of Matching the Life-Cycle Dynamics of Hours and Prices



Notes: The figures plot the results of a version of the model in which we do not match the life-cycle dynamics of prices and hours. In the benchmark analysis we assumed ϕ to be age dependent to match the model's labor supply dynamics to the data. For the prices, we assumed age dependent θ_0 . The robustness analysis fixes both parameters to a constant.

7 Discussion and Conclusion

In this paper, we study the role of price search on the age-inequality profiles of consumption and expenditure by introducing price search decisions into a life-cycle model. Our quantitative exercise, using an estimated wage process and price search functions from the literature, predicts that consumption inequality is significantly lower than expenditure inequality when households are allowed to search for prices. A plausibly calibrated version of our model implies that the increase in cross-sectional variance of log consumption is about 30 percent smaller than that of the log expenditure throughout the life-cycle. This result is robust to the various specifications we study in the paper.

We find that average expenditure and consumption growth deviate from each other over the life-cycle due to the dynamics of prices paid. This finding has particularly important implications for the literature that documents and studies the joint dynamics of income and consumption. We argue that if consumption data were used in those studies, consumption would still move with income, but the comovement would be weaker. The reason is that, relative to low-income growth households, high income growth ones will increase their expenditure more relatively to their consumption as the prices they pay increase more. As a result, the consumption differential will be less than that of expenditure.

Our results also suggest that the availability of price search allows for better insurance for households and increases their welfare. It also affects labor supply and consumption/saving decisions. The price search option also lowers the average hours worked since households can increase their consumption by lowering prices via search, which acts like non-market work.

The model can potentially be extended further to explain other empirical observations. For instance, Aguiar & Hurst (2013) document different inequality patterns for different expenditure categories. If, for example, different categories of goods feature different returns to search, then price search could be helpful in explaining the different expenditure patterns. It is also possible to analyze how improving search technologies may change the model's implications. For a proper analysis, however, it is necessary to model firms, as such changes in search technology may have an influence on price setting behavior.

References

- [1] Aguiar, M. and M. Bilal (2015), "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, 105(9), 2725-56.
- [2] Aguiar, M. and E. Hurst (2005), "life-cycle Prices and Production," *Boston Federal Reserve Bank Discussion Paper No: 05-3*.
- [3] Aguiar, M. and E. Hurst (2007), "Life-Cycle Prices and Production," *American Economic Review*, 97(5), 1533-1559.
- [4] Aguiar, M. and E. Hurst (2013), "Deconstructing Life-Cycle Expenditure," *Journal of Political Economy*, 121(3), 437-492.
- [5] Aiyagari, S. R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109(3), 659-84.
- [6] Baye, M.R., J. Morgan and P. Scholten (2004), "Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site," *Journal of Industrial Economics*, 52(4), 463-496.
- [7] Baye, M. R., J. Morgan, and P. Scholten (2006), "Information, Search, and Price Dispersion," *Handbook on Economics and Information Systems*.
- [8] Becker, G. (1965), "A Theory of the Allocation of Time," *The Economic Journal*, 75, 493-517.
- [9] Becker, G. and G. R Ghez (1975), "The Allocation of Time and Goods over the life-cycle," NBER Working Paper.
- [10] Blundell, R and I. Preston (1998), "Consumption Inequality and Income Uncertainty?," *Quarterly Journal of Economics*, 113 (2), 603-640.
- [11] Blundell, R., L. Pistaferri, and I. Preston (2008), "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887-1921.
- [12] Carroll, C. and L. Summers (1989), "Consumption Growth Parallels Income Growth: Some New Evidence," NBER Working Paper.
- [13] Chang, Y. and S. Kim (2007), "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations," *American Economic Review*, 97(5), 1939-56.
- [14] Dahlbay B. and D. S. West (1986), "Price Dispersion in an Automobile Insurance Market," *Journal of Political Economy*, 94(2), 418-38.

- [15] Deaton, A. and C. Paxson (1994), "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102(3), 437-67.
- [16] Gourinchas, P. and J. Parker (2002), "Consumption over the life-cycle," *Econometrica*, 70(1), 47-89.
- [17] Guner, N., G. Ventura and X. Yi (2008), "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, 11(4), 721-744.
- [18] Guvenen, F. (2007), "Learning Your Earning, Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, 97(3), 687-712.
- [19] Guvenen, F. (2009), "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12(1), 58-79.
- [20] Hansen, G (1985), "Indivisible labor and the business cycle," *Journal of Monetary Economics*, 16(3), 309-327.
- [21] Heathcote, J., K. Storesletten and G. Violante (2010), "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, 13(1), 15-51.
- [22] Heathcote, J., K. Storesletten and G. Violante (2014), "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 104(7), 2075-2126.
- [23] Kaplan, G. (2012), "Inequality and the Lifecycle," *Quantitative Economics*, 3(3), 446-512.
- [24] Kaplan, G. and G. Menzio (2016), "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy*, 124(3), 771-825.
- [25] Kaplan, G. and G. Menzio (2015), "The Morphology of Price Dispersion," *International Economic Review*, 56(4), 1165-1204.
- [26] Kaplan, G. and G. Violante (2010), "How Much Consumption Insurance Beyond Self-Insurance?," *American Economic Journal: Macroeconomics*, 2(4), 53-87.
- [27] Karahan, F. and S. Ozkan (2013), "On the Persistence of Income Shocks over the life-cycle: Evidence and Implications," *Review of Economic Dynamics*, 16(3), 452-476.
- [28] Krueger, D. and F. Perri (2006), "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73(1), 163-193.
- [29] Li, Nicholas (2013), "An Engel Curve for Variety," Working Paper.

- [30] MaCurdy, T. E. (1982), "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics*, 18(1), 83-114.
- [31] Nevo, A. and A. Wong (2018), "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession," *International Economic Review*, Vol.60, No.1, 2019.
- [32] Pratt, J.W., D.A. Wise and R. Zeckhauser (1979), "Price Differences in Almost Competitive Markets," *Quarterly Journal of Economics*, 93(2), 189-211.
- [33] Pytka, K (2018), "Shopping Effort in Self-Insurance Economies," Working Paper.
- [34] Sorensen, A. (2000), "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy*, 108(4), 833-850.
- [35] Stroebel, J. and J. Vavra (2019), "House Prices, Local Demand, and Retail Prices," *Journal of Political Economy*, 127(3), 1391-1436.
- [36] Stigler, G. J. (1961), "The Economics of Information," *Journal of Political Economy*, 69(3), 213-225.
- [37] Storesletten, K., C. Telmer and A. Yaron (2004), "Consumption and Risk Sharing over the life-cycle," *Journal of Monetary Economics*, 51(3), 609-633.
- [38] Woodward. S., E. and R. Hall (2012), "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence," *American Economic Review*, 102(7), 3249-76.

Appendix

A Proofs

Proof of Lemma 1 When labor supply is positive, the FOCs with respect to price search and labor supply become

$$\begin{aligned} s &: u_2 c^\psi = \lambda p'(s) c \\ n &: u_2 = -w\lambda. \end{aligned}$$

Combining these two equations results in

$$p'(s) = -wc^{\psi-1}.$$

If we take the total derivative of this equation with respect to s and c , we arrive at

$$p''(s) ds = (1 - \psi) wc^{\psi-2} \frac{dc}{dy} dy,$$

which can be written as

$$\frac{ds}{dy} = (1 - \psi) \frac{wc^{\psi-2}}{p''(s)} \frac{dc}{dy}.$$

If consumption is normal good, then $\frac{dc}{dy} > 0$, and if $p''(s) > 0$, then the sign of $\frac{ds}{dy}$ is equal to the sign of $1 - \psi$.

Derivation of Equation 7 Expenditure is defined as

$$\begin{aligned} e &= pc \\ \log e &= \log p + \log c. \end{aligned}$$

Substituting equation (1) into the above equation, we have

$$\log e = \theta_0 + \theta_1 \log s + \log c.$$

Taking the variance of both sides, we get

$$Var(\log e) = \theta_1^2 Var(\log s) + Var(\log c) + 2\theta_1 Cov(\log s, \log c). \quad (15)$$

The FOC for price search evaluated at $n = 0$ results in

$$\log s = -\frac{\log\left(\phi\alpha^{1+\gamma}\left(1 - \frac{1}{\theta_1}\right)\right)}{1 + \gamma} - \frac{\sigma + \gamma}{1 + \gamma} \log(c).$$

Using this equation, we can write the $Var(\log s)$ as

$$Var(\log s) = \left(\frac{\sigma + \gamma}{1 + \gamma}\right)^2 Var(\log c),$$

and $Covar(\log s, \log c)$ as

$$Cov(\log s, \log c) = -\frac{\sigma + \gamma}{1 + \gamma} Var(\log c).$$

Substituting these expressions into (15), we get

$$\begin{aligned} Var(\log e) &= \left(\left(\theta_1 \frac{\sigma + \gamma}{1 + \gamma} \right)^2 + 1 - 2\theta_1 \frac{\sigma + \gamma}{1 + \gamma} \right) Var(\log c) \\ &= \left(1 - \theta_1 \frac{\sigma + \gamma}{1 + \gamma} \right)^2 Var(\log c). \end{aligned}$$

Derivation of Equation 6 When labor supply is positive, as expressed in equation (4), we have

$$\log(s) = \frac{\theta_0 + \log(-\theta_1)}{1 - \theta_1} - \frac{\log(w)}{1 - \theta_1}.$$

Taking the variance of both sides, we arrive at

$$Var(\log s) = \left(\frac{1}{1 - \theta_1} \right)^2 Var(\log w).$$

Substituting this into (15) we have

$$Var(\log e) = \left(\frac{\theta_1}{1 - \theta_1} \right)^2 Var(\log w) + Var(\log c) - 2\frac{\theta_1}{1 - \theta_1} Cov(\log c, \log w).$$

B Price Search Technology

While there are several empirical studies that estimate effects of price search on prices paid, literature provides limited guidance on how households search for prices in the market. Here, we provide a justification for the reduced form price search technology we assumed in equation 1 based on exogenous price dispersion posted by producers in the market. We follow the seminal paper by Stigler (1961).

We assume that households can affect the price they pay by receiving price quotes from multiple stores. More specifically, we assume that the distribution of prices is given by an exogenous commonly known non-degenerate distribution $F(p)$ on $[\underline{p}, \bar{p}]$. At the beginning of each period, households decide on the number of stores, s , that they will visit.²⁴ That is, the price households pay upon receiving s quotes is $\min\{p_1, p_2, \dots, p_s\}$. It can be shown that the distribution of the lowest of s draws is $F_{\min}(p; s) = 1 - [1 - F(p)]^s$. Expected price paid by the household upon obtaining s quotes becomes $E(p_{\min}; s) = \int_{\underline{p}}^{\bar{p}} p dF_{\min}(p; s) = \underline{p} + \int_{\underline{p}}^{\bar{p}} [1 - F(p)]^s dp$, in which the second equality is obtained from integration by parts. We assume that, in a given period, households pay the following price:²⁵

$$p(s) = E(p_{\min}; s) = \underline{p} + \int_{\underline{p}}^{\bar{p}} [1 - F(p)]^s dp. \quad (16)$$

Using detailed grocery data, Kaplan & Menzio (2015) find that the typical price distribution households face is symmetric, unimodal, and leptokurtic. They estimate the standard deviation of prices to be between 19 percent to 36 percent depending on how goods are classified. At the UPC level it is 19 percent whereas when products are aggregated with different brands and sizes it becomes 36 percent. If we assume F follows a normal distribution with mean 1 and standard deviation 0.235, a regression of $\log(p(s)) = \beta_0 + \beta_1 \log(s)$ results in $\beta_1 = -0.19$ with an $R^2 = 0.999$. If we keep the same standard deviation but increase the kurtosis of the distribution by fixing the skewness at zero (which is essentially the observed price distribution reported by Kaplan & Menzio, 2015), the estimated return to search increases.²⁶

C Robustness Calibration

In each of the robustness analyses, we recalibrate the parameters to match the same moments as in the benchmark economy. Only in the case of the model without price search, fixed disutility and fixed price search technology, the set of the parameters, and their corresponding moments to

²⁴We restrict s to be a finite integer. Households receive a price quote from each of the s stores, and pick the lowest price to purchase the consumption bundle.

²⁵This assumption can be supported by assuming households repeat the process of obtaining a fixed amount of price quotes sufficiently many times within a period. As long as each quote is independent from each other, households pay this expected price within a period.

²⁶This can be achieved by assuming Pearson type VII distribution such as

$$f(p) = \frac{1}{\alpha B(m-0.5, 0.5)} \left(1 + \left(\frac{p - \mu_p}{\alpha} \right)^2 \right)^{-m},$$

where B is the Beta function $\alpha = \sigma_p \sqrt{2m-3}$ and $m = \frac{5}{2} + \frac{3}{\zeta}$ with ζ as the excess kurtosis, μ_p as the mean of the distribution and σ_p as the standard deviation of the distribution. When ζ approaches 0, the distribution converges to normal distribution. If we use set $\mu_p = 1$, $\sigma_p = 0.235$ but $\zeta = 10$, then the same regression results $\beta_1 = -0.24$ with an $R^2 = 0.986$.

Table 4: Parameters: Robustness

	Risk Aversion		Frisch		ψ	Search Return		Fixed Labor	No Search	Constant ϕ	Constant θ_0
	$\sigma = 1$	$\sigma = 3$	$\gamma = 2$	$\gamma = 6$	$\psi = 1.5$	$\theta_1 = -0.1$	$\theta_1 = -0.3$				
β	0.965	0.970	0.966	0.968	0.967	0.966	0.968	0.967	0.962	0.968	0.966
ϕ^1	169.3	781.2	49.8	2746	337.7	545.4	239.8	552.3	819.5	92.5	337.8
ϕ^T	69.3	75.6	13.0	390	66.4	109.1	44.7	72.9	176.3	-	72.7
ϕ^P	10	9.3	6.8	8.3	7.4	4.0	7.5	7.4	7.96	-	7.7
θ_0^1	-0.27	-0.27	-0.28	-0.26	-0.19	-0.18	-0.33	-0.33	-	-0.27	-0.41
θ_0^T	-0.44	-0.45	-0.46	-0.44	-0.45	-0.28	-0.60	-0.44	-	-0.45	-
θ_0^P	2.25	2.25	2.24	2.25	2.52	2.26	2.24	2.81	-	2.25	-

be matched change. Table 4 presents the parameters calibrated for each robustness case.

Previous volumes in this series

885 September 2020	Credit supply driven boom-bust cycles	Yavuz Arslan, Bulent Guler and Burhan Kuruscu
884 September 2020	Retailer markup and exchange rate pass-through: Evidence from the Mexican CPI micro data	Fernando Pérez-Cervantes
883 September 2020	Inflation at risk in advanced and emerging market economies	Ryan Banerjee, Juan Contreras, Aaron Mehrotra and Fabrizio Zampolli
882 September 2020	Corporate zombies: Anatomy and life cycle	Ryan Banerjee and Boris Hofmann
881 September 2020	Data vs collateral	Leonardo Gambacorta, Yiping Huang, Zhenhua Li, Han Qiu and Shu Chen
880 August 2020	Rise of the central bank digital currencies: drivers, approaches and technologies	Raphael Auer, Giulio Cornelli and Jon Frost
879 August 2020	Corporate Dollar Debt and Depreciations: All's Well that Ends Well?	Julián Caballero
878 August 2020	Which credit gap is better at predicting financial crises? A comparison of univariate filters	Mathias Drehmann and James Yetman
877 August 2020	Export survival and foreign financing	Laura D'Amato, Máximo Sangiácomo and Martin Tobal
876 August 2020	Government banks, household debt, and economic downturns: The case of Brazil	Gabriel Garber, Atif Mian, Jacopo Ponticelli and Amir Sufi
875 August 2020	The Impact of credit risk mispricing on mortgage lending during the subprime boom	James A Kahn and Benjamin S Kay
874 August 2020	Demographic origins of the decline in labor's share	Andrew Glover and Jacob Short
873 July 2020	Effects of Fed policy rate forecasts on real yields and inflation expectations at the zero lower bound	Gabriele Galati and Richhild Moessner
872 July 2020	Effects of credit restrictions in the Netherlands and lessons for macroprudential policy	Gabriele Galati, Jan Kakes and Richhild Moessner

All volumes are available on our website www.bis.org.