

Aguiar, Mark; Hurst, Erik

Working Paper

Lifecycle prices and production

Public Policy Discussion Papers, No. 05-3

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Aguiar, Mark; Hurst, Erik (2005) : Lifecycle prices and production, Public Policy Discussion Papers, No. 05-3, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/59240>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Lifecycle Prices and Production

Mark Aguiar and Erik Hurst

Abstract:

Using scanner data and time diaries, we document how households substitute time for money through shopping and home production. We find evidence that there is substantial heterogeneity in prices paid across households for identical consumption goods in the same metro area at any given point in time. For identical goods, prices paid are highest for middle-aged, rich, and large households, consistent with the hypothesis that shopping intensity is low when the cost of time is high. The data suggest that a doubling of shopping frequency lowers the price paid for a given good by approximately 10 percent. From this elasticity and observed shopping intensity, we impute the shopper's opportunity cost of time, which peaks in middle age at a level roughly 40 percent higher than that of retirees. Using this measure of the price of time and observed time spent in home production, we estimate the parameters of a home production function. We find an elasticity of substitution between time and market goods in home production of close to 2. Finally, we use the estimated elasticities for shopping and home production to calibrate an augmented lifecycle consumption model. The augmented model predicts the observed empirical patterns quite well. Taken together, our results highlight the danger of interpreting lifecycle expenditure without acknowledging the changing demands on time and the available margins of substituting time for money.

Keywords: Lifecycle Consumption, Permanent Income Hypothesis, Home Production, Time Use

JEL Classifications: E2, D1, D4

Mark Aguiar is a Senior Economist at the Federal Reserve Bank of Boston. Erik Hurst is an Associate Professor of Economics and the John Huizinga Faculty Fellow at the Graduate Business School of the University of Chicago. Their e-mail addresses are mark.aguiar@bos.frb.org and erik.hurst@gsb.uchicago.edu, respectively.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bos.frb.org/economic/ppdp/index.htm>.

The views expressed in this paper are solely those of the authors and do not reflect official positions of the Federal Reserve Bank of Boston or the Federal Reserve System.

We are particularly grateful to Jean-Pierre Dube for assistance with the Homescan data. We would also like to thank seminar participants at the Boston Fed, University of California (Berkeley), University of Chicago, Columbia University, MIT, the NBER Summer Institute, Northwestern University, and the University of Rochester. Hurst thanks the University of Chicago Graduate School of Business for financial support.

This version: July 2005

1. Introduction

This paper studies how households substitute time and money. The vast majority of the literature on this question focuses on labor supply decisions. However, such an exclusive focus overlooks a number of other mechanisms that households use to substitute time for money. In this paper, we use a novel data set to document that households that shop more intensively pay lower prices for *identical* goods. We merge this data set with a new data set on time use to estimate parameters of the shopping and home production technologies households employ to minimize the total cost of consumption. Then, using a quantitative lifecycle model, we show that observed household behavior, in terms of expenditure, time use, and prices, is consistent with standard economic principles, once we allow households to access the shopping and home production technologies. When households have access to home production and shopping technologies, we show that market expenditure is a poor proxy for actual consumption, when the value of the household's time is changing.

The economic theory that motivates this paper originated in two seminal works of the 1960s. Becker (1965) formalized the notion that consumption is the output of a production function that combines market goods and time.¹ Such a "home production" function allows households to substitute time for expenditures optimally in response to fluctuations in the relative cost of time. A similar implication lies behind Stigler's (1961) model of search. In the presence of informational frictions, the same good may sell for different prices at a given point in time. By shopping more intensively, a household can lower the market price for a given basket of goods.

These theoretical insights are now familiar. However, the quantitative importance of these margins is difficult to pin down.² The first contribution of this paper is to explore how prices for goods vary across households in practice, and to what extent this variation accords

¹ An even earlier reference is Reid (1934).

² Recent empirical papers documenting price dispersion and returns to search in retail prices include Sorensen (2000) and Brown and Goolsbee (2002). Using Argentine scanner data, McKenzie and Schargrodsky (2004) find that shopping increased and transaction prices fell during the 2002 Argentine crisis. Rupert, Rogerson, and Wright (1995, 2000) and Aguiar and Hurst (forthcoming) use micro data to document the importance of home production. McGrattan, Rogerson, and Wright (1997), Campbell and Ludvigson (2001), and Chang and Schorfheide (2003) calibrate or estimate home production parameters using aggregate data.

with standard theory. To do this, we use data from ACNielsen's Homescan survey. This survey collects grocery scanner data at the level of the household. Each purchase in the database records the actual price paid by the household at the level of the Universal Product Code (UPC) bar code. The database is novel in the sense that it has detailed demographics about the household making the purchases and it tracks the household purchases across multiple retail outlets. Because the data set also includes information about the shopping trip, we can infer the household's shopping intensity.

We find that the price paid for a particular good is an increasing function of income. Specifically, households with annual income over \$70,000 pay 5 percent more on average for an identical good (defined by UPC) than households earning less than \$30,000. This result is consistent with the fact that high-income households face a higher opportunity cost of time. Additionally, we find that households with more children pay higher prices than households with fewer or no children. This effect is robust to controls for income. Given the additional time demands associated with having children, households with more children will have higher opportunity costs of time compared with households with no children, all else being equal.

One of our most striking results is that prices paid by households are hump-shaped over the life cycle. Households with heads in their early forties pay, on average, between 6 percent and 8 percent more for identical goods than either households with heads in their early twenties or households with heads in their late sixties. Households with heads in their forties face the highest market opportunity cost of time (the highest wages) as well as the highest nonmarket demands on time (the most children). Also, we document that there is a lifecycle profile to the dispersion of prices paid for identical goods. That is, along with the higher mean price, middle-aged shoppers also pay a wider variety of prices over time for a particular good. This is consistent with standard search-theoretic intuition. Busy, middle-aged shoppers pay whatever price happens to prevail at the time and place of purchase. Retirees, on the other hand, search more intensively, and, in the process, generate a tighter price distribution around a lower mean.

Given the price data, as well as information on shopping frequency in the Homescan data set, we are able to estimate a "shopping function" that maps time and quantity purchased into price. Holding constant the quantity of goods purchased, we find that households that

shop more frequently pay lower prices. Specifically, all else being equal, households that double their shopping frequency will pay prices that are 5 to 12 percent lower, on average. Likewise, holding shopping frequency constant, households that purchase more goods pay higher prices.

Optimality implies that the shopper equates the marginal value of additional shopping for lower prices to the opportunity cost of time. With this in mind, we use the observed shopping behavior, as well as the estimated shopping function, to calculate the shopper's opportunity cost of time for each household. We show that the cost of time is hump-shaped over the life cycle, but in a manner that differs from the wage of the household head. This reflects the reality that the shopper may not face the same wage as the household head and/or that the household may not be able to adjust labor hours at the margin.

A second contribution of the paper is that we use the price data to estimate the parameters of a home production function. The identification assumption is that the opportunity cost of time of the shopper is the same as that of the person undertaking home production. Under this assumption, the marginal rate of transformation (MRT) of time to dollars in shopping is equated to that in home production. Using detailed data on time spent in home production from the recent American Time Use Survey (ATUS), we can use the first-order condition between shopping and home production to estimate the parameters of the home production function. The advantage of this approach is that we do not need to assume that the cost of time in home production is the market wage. This allows us to calculate a price of time for retirees and for married households with only one worker. We estimate an elasticity of substitution between time and goods in home production of close to 2, and, in all specifications, we can reject the null hypothesis that the elasticity is 1 or less.

With the home production function in hand, we calculate implied household consumption, using observed inputs of time and market goods. We document that this series varies over the life cycle in a manner distinct from household expenditures. Specifically, the ratio of implied consumption to expenditures declines as household heads enter middle age and then rises rapidly through their retirement. The lifecycle profile of this ratio reflects the changing cost of time as household heads age and highlights the danger of inferring the lifecycle profile of consumption directly from expenditures.

Finally, we incorporate into an otherwise-standard model of lifecycle consumption the fact that households can shop for bargains and undertake home production. We find that our simple model augmented with home production and shopping can match the data quantitatively along a variety of dimensions. In particular, our model generates a hump-shaped profile in household expenditure over the life cycle of similar magnitude to the profile found in the data. Additionally, our model matches the empirical lifecycle profiles of time spent shopping, time spent in home production, and prices paid. In this sense, the empirical pattern of shopping intensity is consistent with optimality, given the observed dispersion of prices.

There is a growing interest in the role of nonmarket activities and the allocation of work between the market and the household. The insights from modeling household production have already proved fruitful in explaining, for example, the baby boom [Greenwood et al. (2005)], business cycles [Benhabib, Rogerson, and Wright (1991)], and the excess sensitivity of consumption to predictable income changes [Baxter and Jermann (1999)]. This paper adds to this literature by documenting quantitatively how home production and shopping behavior drive a wedge between household market expenditures and actual household consumption. This wedge increases as the price of time increases. As a result, holding family structure constant, households with middle-aged heads will have higher expenditures and lower consumption than either their younger or older counterparts.

The lifecycle profile of expenditures has been well documented.³ Heckman (1974) interprets the hump shape in expenditure over the life cycle as being evidence that household utility is nonseparable in consumption and leisure. When leisure is low (during middle age), households compensate by increasing their expenditures. Attanasio et al. (1999) and Blundell, Browning, and Meghir (1994) attribute a portion of the lifecycle profile of expenditure to changing preferences associated with changing household structure. Our data, and the accompanying model, provide a microfounded story of how the ability to home produce and shop implies a nonseparability between *expenditure* and leisure even when utility is separable over *consumption* and leisure.

³ See, for example, Heckman (1974), Blundell, Browning, and Meghir (1994), Attanasio et al. (1999), Gourinchas and Parker (2002), and Krueger and Fernandez-Villaverde (forthcoming).

2. Prices Paid Over the Life Cycle

2.1. Data

Our price data are from ACNielsen Homescan Panel. The Homescan database is designed to capture all consumer grocery packaged goods purchased by the household at a wide variety of retail outlets. We use the Homescan database for Denver, covering the period January 1993 through March 1995.⁴ The survey is designed to be representative of the Denver metropolitan statistical area, and summary demographics line up well with the estimates from the 1994 Panel Study of Income Dynamics (PSID) (see Table A1).

Respondents in the Homescan survey remain in the survey for upwards of 27 months. The survey is implemented at the household level and contains detailed demographics, which are updated annually. Specifically, we know the characteristics of the household: the head's age, sex, race, education, and employment status; and family composition and household income. Employment status and household income are broadly measured as categorical variables.

Households selected for the Homescan sample are equipped with an electronic home scanning unit. After every shopping trip, the shopper scans the UPC bar codes of all the purchased packaged goods.⁵ The shopper provides three additional pieces of information regarding each transaction: the date, the store, and the total amount of discounts due to promotions, sales, or coupons. The scanners are programmed to include all the stores in the household's shopping area (including grocery stores, convenience stores, specialty stores, super centers, and price clubs). If a household member shops at a store outside the defined shopping area, the household can manually input the store information. ACNielsen maintains a database of current prices for all stores within the metropolitan area. Given the store and date

⁴ For proprietary reasons, the ACNielsen Company is reluctant to release any of the Homescan data. However, in the late 1990s, they did make the Denver data available to academics for research. For this reason, we have access only to the Denver data from the early 1990s. We thank Jean-Pierre Dube for providing us with the data.

⁵ Each type of packaged good has a unique UPC bar code printed on its packaging. The codes are very specific. A liter bottle of Pepsi, a six pack of Pepsi cans, and a twelve pack of Pepsi cans, all have distinct UPC bar codes.

information, ACNielsen can link each product scanned by the household to the actual price for which it was selling at the retail establishment.⁶ In terms of associated demographics and coverage of multiple outlets, the Homescan database is superior to retail-based scanner data for lifecycle analysis.

Within the Homescan data set, we have 2,100 separate households and over 950,000 transactions. For our analysis, we focus on households in which the average age of the “primary shopper” is between the ages of 24 and 75, and, unless otherwise noted, we restrict the age of the household head to be at least 25.⁷ This restriction leaves us with just over 2,000 households.

One should keep in mind that the database is essentially a cross-section at a given point in time (the panel dimension covers only 27 months). Therefore, when we discuss lifecycle patterns, we will be comparing different cohorts. This may, for example, overstate the decline in expenditure between households with middle-aged heads (richer cohorts) and households with older heads (poorer cohorts). Likewise, it could cause us to understate the increase in expenditure between households with young heads and those with middle-aged heads. However, this should not be as important an issue for the normalized variables we focus on, such as the ratio of consumption to expenditure, as it would be if we used non-normalized data.

In Appendix A, we discuss and quantify a number of potential data quality issues with the Homescan data. These issues include: representativeness of the households in the Homescan sample, coverage of the goods scanned by households in the sample, sample attrition, and the importance of store and grocery-chain fixed effects.

⁶ Households may pay less than the stated store price if they use coupons or avail themselves of in-store discounts. This information is manually entered by household members. Given that this information is likely fraught with large amounts of measurement error, we do not use it when computing and reporting our price indices. However, as a robustness check, we repeated our analysis including the coupon information when computing price differences across households. As one would expect, the inclusion only strengthened our results, given the higher propensity of retired households to use coupons [see Cronovich et al. (1997)]. In other words, households with a low opportunity cost of time are also more likely to clip coupons.

⁷ The Homescan database records up to three ages for each household: that of the male head, the female head, and the primary shopper. The former two are categorical variables while the latter takes on all positive integers. The age of the primary shopper may change from shopping trip to shopping trip, depending on who did the shopping. For the remaining analysis, we focus on the age of the household head. When two heads are present, we follow standard practice (as in the PSID) and use the male head’s age. Given the fact that the heads’ ages are recorded in five-year blocks (for example, 25–29), the majority of married-couple households report the same age category for both heads. As a result, it makes little difference to our analysis whether we use the shopper’s age, the male head’s age, or the female head’s age.

2.2. Prices Paid and the Opportunity Cost of Time

Standard economics suggests that, all else being equal, households with a lower opportunity cost of time will be more likely to spend time searching/shopping to reduce the price paid for a given market good. There are many ways a household can do this. For example, the shopper may visit multiple stores to take advantage of store-specific sales, shop at superstores⁸—which may involve longer commutes and check-out lines than shopping at convenience stores—or clip coupons and mail in rebates.

Using the Homescan data, we can test the basic premise that households with a lower opportunity cost of time pay lower prices for identical goods. Given the fact that households buy a variety of different goods during each shopping trip, we need to define an average price measure for each household. To set notation, let $p_{i,t}^j$ be the price of good $i \in I$ purchased by household $j \in J$ on shopping trip (date) t . Let $q_{i,t}^j$ represent the corresponding quantity purchased. Total expenditure during month m is simply

$$X_m^j = \sum_{i \in I, t \in m} p_{i,t}^j q_{i,t}^j. \quad (2.1)$$

At the same point in time, there may be another household purchasing the same good at a different price. We average over households within the month to obtain the average price paid for a given good during that month, where the average is weighted by quantity purchased:

$$\bar{p}_{i,m} = \sum_{j \in J, t \in m} p_{i,t}^j \left(\frac{q_{i,t}^j}{\bar{q}_{i,m}} \right), \quad (2.2)$$

where

$$\bar{q}_{i,m} = \sum_{j \in J, t \in m} q_{i,t}^j. \quad (2.3)$$

The next task is to aggregate the individual prices into an index. We do so in a way that answers the question, how much more or less than the average is the household paying for its

⁸ Recently, Hausman and Leibtag (2004), using Homescan data, document that stores like Wal-mart offer prices between 5 percent and 55 percent lower than those for the same product in traditional grocery stores.

chosen basket of goods. That is, if the household paid the average price for the same basket of goods, the cost of the bundle would be

$$Q_m^j = \sum_{i \in I, t \in m} \bar{p}_{i,t} q_{i,t}^j. \quad (2.4)$$

We then define the price index for the household as the ratio of expenditures at actual prices to the cost of the bundle at average prices. We normalize the index by dividing through the average price index across households within the month, ensuring that for each month the index is centered around 1:

$$p_m^j = \frac{\tilde{p}_m^j}{\frac{1}{J} \sum_{j'} \tilde{p}_m^{j'}}, \quad (2.5)$$

where

$$\tilde{p}_m^j = \frac{X_m^j}{Q_m^j}. \quad (2.6)$$

The price index defined in (2.6) shares the typical feature (as with Laspeyres and Paasche indices) that the basket of goods is held constant as we vary the prices between numerator and denominator. To the extent that relative price movements induce substitution between goods, there is no reason to expect that the household would keep its basket constant.

One subtle difference does exist between the substitution bias inherent in our index and that presented by the typical price index. In a standard price index, the relative price of two goods may differ across time periods. In our framework, the distribution of prices for any two goods is the same across households, but the relative price of time varies. This results in variation in the relative *purchase* price of goods. However, it is, in theory, feasible for household j to purchase goods at the prices paid by household j' and vice versa. This is not true in intertemporal price comparisons, such as the CPI. By revealed preference, households in our sample would never be better off if they paid prices (inclusive of time shopping) recorded by other households that period, including the average price.

We interpret a price index greater than 1 as reflecting a household that pays higher prices on average, and vice versa for an index less than 1. It is important that the price premium

is not reflecting higher quality. With our index, the price differentials are for the *identical* goods as defined by the UPC.^{9,10}

Using our price index, we can revisit whether prices paid for the same goods vary across households with different costs of time. One measure of the opportunity cost of time is the market wage. In the Homescan data, we do not have wages; we have only categorical measures of household income. Using this data set, we aggregate up to four income categories: income below \$30,000, income between \$30,000 and \$50,000, income between \$50,000 and \$70,000, and income greater than \$70,000. In Table 1, Column 1, we report the mean price index for households within the four income categories. The results are striking. Households that earn less than \$30,000 a year, on average, pay 5 percent lower prices than households that earn over \$70,000 (p-value of difference < 0.01).^{11,12} Households that earn between \$30,000 and \$50,000 pay, on average, 3 percent lower prices than households that earn over \$70,000 (p-value of difference < 0.01). The difference in prices paid between households that earn less than \$30,000 and households that earn between \$30,000 and \$50,000 is also statistically significant (p-value of difference = 0.04). There is no significant difference in prices paid by households earning between \$50,000 and \$70,000 and those households earning above \$70,000 (p-value = 0.66). Overall, we find that, for a given basket of goods, low-income households pay lower prices than high-income households.¹³

A second influence on the opportunity cost of time is the large time demand associated with raising children. In Column 2 of Table 1, we see that households with larger families pay

⁹ See Appendix A for a discussion of how we redefined our price index to account for grocery-chain fixed effects. Our results were robust to this modification.

¹⁰ An alternative price index could be constructed by forming the ratio of price to average price for each good and averaging across the household's basket. The difference between that measure and the one defined above is not substantial in practice — they share a correlation coefficient of 0.8.

¹¹ The p-value represents the probability that a given result would be expected to occur entirely by chance. For example, if the p-value is less than .01, then we would expect such a result by chance in fewer than 1 percent of repeated outcomes. Since this probability is small, we reject the hypothesis that the difference is zero.

¹² Technically, the difference in the price index is 0.05 points. We refer to this difference as an approximately 5 percent increase due to the normalization of the price index to 1. A similar caveat holds throughout.

¹³ There is mixed evidence that prices are higher in poor neighborhoods [see the survey by Kaufman et al. (1997)]. These poor neighborhoods are usually associated with households having incomes much lower than \$30,000 a year. In our data, households in the poorest income bracket (income below \$5,000) do pay slightly higher prices on average than those closer to \$30,000. However, the small number of extremely low-income households makes it difficult to characterize this potential nonmonotonicity precisely.

higher prices than households with smaller families. Specifically, households with only one household member pay 10 percent less for an item than households with family size greater than or equal to five (p-value < 0.01). Similarly, Column 3 of Table 1 reports that single females with no children pay 7 percent lower prices than married couples with children (p-value < 0.01), while single males without children pay 4 percent lower prices than married couples with children (p-value < 0.01). These differences persist after controlling for household income. When we regress the price index on income categories and family-size categories together, both sets of regressors enter significantly (results not reported, but available from authors upon request).

Of course, more than the price of time varies across the income and household-size categories. In particular, middle-aged households (with higher incomes and larger household size) are purchasing a larger basket of goods. We will explore how this influences price in the regressions reported in Section 2.4. It should be noted that a larger consumption basket increases the returns to shopping at the same time as higher income and larger household size raise the cost of shopping. The model of Section 5 will allow us to see how a household optimally weighs these considerations. Empirically, as Table 1 indicates, the costs of shopping dominate to the extent that richer and larger households pay higher prices.

We also explored whether married families in which both adults work at least 30 hours per week in the market differ from those in which only one spouse works in the market. Perhaps surprisingly, we find little difference in mean price paid. However, the absence of a differential may reflect endogeneity of market labor. For example, households that face greater time demands within the home may opt to have only one spouse work in the market, while those that do not face such heavy demands may have both spouses supply market labor. In this way, the opportunity cost of time may be uncorrelated with the labor status of a spouse. Moreover, there may be an income effect that reduces a spouse's willingness to supply market labor and also reduces the intensity of shopping. Note that this implication of labor supply endogeneity does not extend directly to retirement or unemployment. In those cases, withdrawal from market labor is due to such forces as a decline in wages, institutional features of pensions, or involuntary layoffs, and should predict a drop in the opportunity cost of time. Among married-couple families, we find that households in which neither spouse works more

than 30 hours per week in the market pay, on average, 2 percent less for goods (p-value=0.04) than married-couple households in which at least one spouse works full-time. Among all households, the difference is 1 percent, but the difference is not statistically significant (p-value=0.41).

Given that the timing of children and household wages both have a lifecycle component, we would expect our price index to vary with the age of the household head. Using the 2000 census, we find that the number of children in a married-couple household is greatest when the head is around age forty (see Figure A1). As seen in Figure A2, wages of both males and females, conditional on working, peak around the age range of 45-to-50.¹⁴ To the extent that labor force participation is declining in late middle age, the observed wages overstate the average cost of time for households in their fifties and sixties. Nevertheless, both the data on number of children present and the data on market wage suggest that the opportunity cost of time is greatest in middle age.

In Figure 1, we show the lifecycle profile of our price index for all households and for married-couple households. Consistent with our premise, households with heads in their middle forties pay the highest prices. Specifically, unconditional on marital status, households with heads aged 45-to-49 pay 7 percent higher prices than households with heads aged 25-to-29 (p-value <0.01) and 4 percent higher prices than households with heads aged 65 and older (p-value <0.01). Conditional on marriage, households with heads aged 40-to-44 pay 8 percent higher prices than households with heads aged 25-to-29 (p-value <0.01) and 6 percent more than those with heads 65 and older (p-value <0.01).¹⁵

One concern is that households may not be paying lower prices solely because of increased shopping intensity, but rather are experiencing lower utility from consumption. Consider two consumers who prefer Pepsi. The first always buys Pepsi, but the second selects Coke or Pepsi depending on which is cheaper. The second consumer will pay a lower price on average for the same Pepsi product. To control for this, we construct two additional measures of

¹⁴ The wage data come from the 1993-to-1995 cross-sections of the PSID.

¹⁵ We redid the analyses in Table 1 and Figure 1, using only the prices and purchases of milk (as opposed to the entire basket of purchases). Milk was the most common product category purchased in the data set. Using only UPC codes within the milk product category, the same conclusions can be drawn. Specifically, middle-aged, rich, and large households pay the highest price for milk.

goods purchased. The first is the number of “product categories” a household purchases per month, where a product category is a broad class such as milk, beer, orange juice, etc. The second is the number of individual UPCs, or “varieties,” a consumer purchases. Distinct varieties include a six-pack of Pepsi, a twelve-pack of Pepsi, a six-pack of Diet Coke, etc.¹⁶ Conditional on the number of product categories, the number of varieties captures the propensity of a household to substitute brands or sizes. As documented below, for a given shopping frequency, more goods imply higher prices (due to dilution of shopping time), and more varieties, for a given set of goods, imply lower prices (due to the propensity to switch brands or items). All the patterns documented in Table 1 and Figure 1 are robust to the addition of these controls.

Finally, in Figure 2 we plot price *dispersion* over the life cycle. We define dispersion in two ways. “Within-household” price dispersion tracks the change in price for the same good and the same household over time, using the panel dimension of the Homescan data. For each household and each year, we compute the standard deviation of log price, good-by-good. We then average these standard deviations across all goods purchased by the household (equally weighted). For the “between-household” price dispersion, we use the cross-sectional dimension of the Homescan data. To create this measure, we segment shoppers into our eight age ranges. For each UPC and each month, we calculate the standard deviation of log prices across households in the same age category. The measure of dispersion averages all the standard deviations of log price across all good-month cells within the age category.

Both series are plotted against the age of the household head in Figure 2, Panel A. To ensure that the observed effect is not due to a changing basket of goods over the life cycle, Figure 2, Panel B, breaks out milk (a single category purchased by almost every household) and performs the same analysis on UPCs within this category. The “within” dispersion peaks in middle age and declines in retirement, dropping by roughly 20-to-40 percent (or 5 percentage points) from peak to trough. The “between” dispersion drops by one-third to one-half between

¹⁶ We also replaced the number of UPCs as our measure of varieties with the number of “brands” (Coke, Pepsi, Miller, etc.) within a product category a household purchases per month. This counts Coke and Pepsi as different varieties, for example, but does not distinguish between six-packs and twelve-packs of Coke. The results were similar, but typically came with a larger standard error on the variety coefficient.

middle age and retirement. This pattern is easily interpreted in a search-theoretic framework and is consistent with the first moment of prices discussed above. Busy, middle-aged shoppers purchase goods at whatever price prevails on the date they shop, sometimes finding sales, but often paying high prices. Retirees, on the other hand, take the time to find the lowest price available. The resulting distribution of prices for retirees should, therefore, have a lower mean and be compressed relative to the distribution for middle-aged shoppers.

The patterns of mean price and price dispersion documented above suggest that shoppers behave in a manner consistent with basic search-theoretic intuition. Of course, these unconditional plots do not hold “all else equal.” Moreover, it is not clear that the observed patterns are quantitatively consistent with optimization. Whether household shopping is optimal, conditional on the equilibrium price dispersion and lifecycle time and consumption demands, is a question we answer within the framework of a quantitative model in Section 5.

2.3. Shopping Over the Life Cycle

Corresponding to the premise that the opportunity cost of time varies over the life cycle, whether due to the wage profile or to alternative demands on time, we would expect the time spent shopping to vary as well. However, the marginal benefit of additional shopping depends on the quantity purchased as well as on the price dispersion, and this makes shopping more valuable in middle age, when families are largest.

To examine time spent in shopping over the life cycle, we use data from the recently released 2003 American Time Use Survey (ATUS), conducted by the U.S. Bureau of Labor Statistics (BLS). Participants in ATUS are drawn from the exiting sample (that is, the final outgoing rotation group) of the Current Population Survey (CPS). Roughly 1,700 individuals complete the survey each month, yielding an annual sample of over 20,000 individuals. Only one individual per household is sampled. Respondents in the sample, via a telephone survey, complete a detailed time diary of their previous day. The BLS staff then aggregate the survey responses into time-use categories.¹⁷

¹⁷ See <http://www.bls.gov/tus> for a detailed description of the ATUS survey methodology and coding system.

We form two measures of time spent shopping. First, we use time spent only on shopping for groceries. Second, we use the total time spent shopping for all household items. As with the Homescan data, we restrict the sample to include only individuals between the ages of 25 and 75. In Table 2, we report the time spent shopping for all households (Panel A) and married-couple households (Panel B) over the life cycle.¹⁸

Both peak grocery and total shopping time occur for households with heads in their early forties and for households with heads 65 and older. Households with heads in their mid-forties have the largest family sizes and, as a result, have the greatest shopping needs. Households in their post-retirement years have the lowest opportunity cost of time, and, therefore, shop most intensively for a given basket of goods. Households with young heads shop relatively little because they buy relatively few goods and have work and education demands on their time. The ratio of grocery shopping to total shopping is fairly constant at 25 to 30 percent over the life cycle.

2.4. Estimation of the Price Function

We can undertake a more formal analysis of price paid by estimating a price function that maps shopping frequency and quantity purchased into the price paid. The estimated elasticities will be used in the lifecycle model of consumption outlined in Section 5. Formally, we wish to estimate the function:

$$p = p(s, Q), \tag{2.7}$$

where p is our price index [as defined in (2.5)], s is the amount of time shopping, and Q is the amount of goods purchased. Our hypotheses are $\partial p / \partial s < 0$; $\partial^2 p / \partial s^2 > 0$; $\partial p / \partial Q > 0$. In other words, holding Q constant, households that shop more will reduce the price they pay, but these returns to shopping diminish as shopping increases. Likewise, holding shopping time constant, households that purchase more goods pay higher prices.

¹⁸ Unfortunately, the BLS did not have each spouse within the same household fill out a time diary. We construct synthetic married-couple households by summing over married men and women, based on the husband's age. Given that each age group contains a fairly large cross-section and that the BLS randomly selects which spouse is recorded within a household, we feel that the bias from this approach is minimized.

The Homescan data set allows us to calculate the number of shopping trips undertaken by the household. Unfortunately, it does not report the time spent per trip. We therefore use trips per month as our measure of s . Below, we discuss how the omission of trip length may bias our estimates. Our benchmark regressions take Q to be purchases evaluated at the mean prices [as defined in equation (2.4)]. We also explore alternatives, such as the total number of product categories purchased and the variety of goods purchased.

Given that we have no strong priors regarding functional form, we estimate a number of specifications. The results are broadly consistent across all the specifications. To begin, we estimate the following two specifications:

$$\ln(p_m^j) = \alpha_0 + \alpha_1 \ln(s_m^j) + \alpha_2 \ln(Q_m^j) + \varepsilon_m^j \quad (2.8)$$

and

$$p_m^j = \beta_0 + \beta_1 s_m^j + \beta_2 (s_m^j)^2 + \sum_{k=1}^5 \beta_{k+2} (Q_m^j)^k + u_m^j. \quad (2.9)$$

The first specification, (2.8), assumes that price is log linear in shopping frequency and quantity. Specification (2.9) assumes that price is a second-order polynomial in shopping time and a fifth-order polynomial in quantity.¹⁹ Columns 1 and 2 of Table 3 report the estimates of (2.8) and (2.9), respectively. We estimate α_1 to be -0.08 (p-value < 0.01). A similar elasticity of -0.12 is obtained from (2.9), evaluated at the sample average. In other words, the data from Homescan indicate that a doubling of the shopping frequency reduces prices paid by roughly 8 to 12 percent, conditional on the quantity purchased. Moreover, the positive coefficient on the second-order term in shopping frequency from (2.9) indicates diminishing returns to search. We also find that the quantity purchased has a statistically significant impact on price in both specifications, with elasticities of 0.07 and 0.11, respectively. That is, for a given shopping frequency, the more purchases a shopper makes, the higher the price paid for the average good.

In columns 3 and 4 of Table 3, we explore other specifications of the price function. The results are stable across different specifications. Column 3 re-estimates (2.8) including controls for the number of product categories and the number of varieties purchased, as defined in

¹⁹ We experimented with polynomials of various lengths. Increasing the polynomial in shopping time beyond a second order and quantity beyond a fifth order had little effect on our results.

Section 2.2. We find the coefficient on the number of product categories, while not reported, to be significantly positive (the “dilution” of shopping time effect), while the coefficient on varieties (given the number of categories) is significantly negative (the brand/item “switching” effect). These additional controls do not dramatically change the elasticities reported in the first two specifications.

One concern with our benchmark specifications is that quantity purchased may be a function of price (the “demand” equation). This issue is less clearly a problem in our analysis than it may first appear. First, we are not tracking purchases by a household as the price varies over time. Rather, in a particular period, we are looking across households that all face the same distribution of prices. In this sense, our “supply curve” is fixed. Second, our price index measures how much one pays for a given UPC good relative to what the average person pays. The fact that one can buy in bulk to reduce the price is not relevant here. The bulk good is treated as a different UPC-coded good. Nevertheless, for completeness, we instrument for log quantity, using income. We also include dummies for household size and composition as additional regressors. Household size and composition may affect shopping efficiency and is correlated with our baseline regressors. Our identification assumption for our instrument is that income plays little direct role in shopping efficiency, once we control for changing family structure and shopping frequency. The results, reported in the final column of Table 3, suggest that the elasticity with respect to shopping frequency is unaffected and the elasticity with respect to quantity is slightly larger when we instrument and control for household structure.

One additional concern with our estimation is that we use shopping frequency rather than shopping time. This distinction is immaterial if time per trip is constant across households. However, the ATUS data suggest that time per trip is not constant over the life cycle. In fact, frequency and time per trip are negatively correlated over the life cycle. In a univariate regression, this would bias our estimated elasticity with respect to time toward zero. However, we cannot make claims regarding the direction of bias in the multivariate regressions. If we employ the time diaries’ average time per trip for each age group as an additional regressor, we find no significant direct impact or changes in the estimated elasticities. However, the need to use age averages reduces the amount of informative variation across individual households.

The results from (2.8) and (2.9) provide us with an empirical relationship between shopping intensity and prices paid. As we show below, this relationship will allow us to estimate the household's implied opportunity cost of time.

3. Estimation of the Home Production Function

At any point in time, an optimizing household will choose the least-cost method of acquiring consumption goods. In this section, we use this fundamental premise to leverage our price data into an estimator of a home production function. For example, a household can save on its food bill both by shopping more intensively and by purchasing raw ingredients and making a meal from scratch as opposed to buying pre-made (or take-out) food. On the margin, households should be indifferent between allocating another unit of time to shopping or to home production.

Time spent on home production varies systematically over the life cycle. Using the ATUS data, we define two measures of home production. The first is total time spent on food production (comprising meal preparation and meal clean-up). The second is total home production (comprising food production, indoor cleaning and chores, clothes care, outdoor maintenance, and lawn care). As seen in Table 2, the pattern of time spent in both food production and total home production over the life cycle mimics the pattern of shopping time. In particular, home production time peaks for households with heads in their early forties, and then again for households with heads who are 65 or older. As with shopping time, households with heads in their early forties have the greatest home production needs (that is, the largest family sizes) and households with heads 65 or older have the lowest opportunity cost of time. Moreover, over the life cycle, the ratio of time spent in food production to time spent in total home production is roughly constant, at approximately 28 percent.

To formalize the home production and shopping decisions, consider a household at time t that wishes to consume C units of a consumption good. Following Becker (1965), consumption goods are commodities produced by combining time and market goods via a home production function. Specifically, the household's cost minimization problem is:

$$\begin{aligned} \min_{\{s, Q, h\}} p(s, Q)Q + \mu(s + h) \\ \text{s.t. } f(h, Q) = C \end{aligned} \quad (3.1)$$

where s is the amount of time spent shopping, Q is the quantity of market goods purchased at price p , h is the amount of time devoted to home production, and μ is the price of time. In Section 5, we embed this cost minimization process in a lifecycle model in which the price of time is determined by the marginal utility of leisure of the shopper. The cost minimization problem does not depend on whether the goods in question are separable in utility with other consumption goods or leisure. However, we need to assume that our price function and home production function for food are adequately captured by our data set. That is, different goods and uses of time enter separably in *production*.²⁰

Letting μ_c denote the multiplier on the constraint, the first-order conditions are

$$\begin{aligned} -\frac{\partial p}{\partial s} Q &= \mu \\ \frac{\partial p}{\partial Q} Q + p &= \mu_c \frac{\partial f}{\partial Q} \\ \mu &= \mu_c \frac{\partial f}{\partial h} \end{aligned} \quad (3.2)$$

The first implication of (3.2) is that we can use our shopping data to estimate the shadow value of time (μ). Note, while $\frac{\partial \ln(p)}{\partial \ln(s)}$ is constant across households, assuming the log-log functional

form of (2.8), $\frac{\partial p}{\partial s} = \frac{\partial \ln(p)}{\partial \ln(s)} * \frac{p}{s}$ is household specific. Given Q , p , s (all defined above)

and $\frac{\partial \ln(p)}{\partial \ln(s)}$ [estimated from (2.8)], we can compute $-\frac{\partial p}{\partial s} Q$ for each household in our

Homescan data. In Figure 3, we plot the lifecycle profile of μ , by averaging $-\frac{\partial p}{\partial s} Q$ over all

²⁰ For the elasticity of substitution of the home production function, we need only assume weak separability. In particular, we need only assume that the ratio of marginal products does not vary with other goods or uses of time. However, when we compute the level of output of the home production function (below), we are making a stronger separability assumption.

households within a given age range and then expressing the series as differences from the 25-to-29 age group. We can see that the opportunity cost of time for the shopper is hump-shaped over the life cycle. It is also evident that the hump differs from that of wages for either males or females (Figure A2). Specifically, the shopper's cost of time rises faster than wages in the early thirties, but then is relatively flatter through middle age, before declining sharply. The wedge between the cost of time and wages should not be surprising. Not all shoppers are able to adjust labor supply at the margin. Indeed, the sharp increase in the shopper's cost of time in the early thirties may be driven by the arrival of children rather than by labor market forces. Moreover, reported wages are conditional on working and are, therefore, not directly informative regarding the unemployed or those out of the labor force. This highlights the benefit of the price data set in calculating the value of time for different types of households.

The first-order conditions imply that the marginal rate of transformation (MRT) between time and market goods in shopping equals the MRT in home production:

$$\frac{\frac{\partial f}{\partial h}}{\frac{\partial f}{\partial Q}} = \frac{\frac{\partial p}{\partial s} Q}{\frac{\partial p}{\partial Q} Q + p}. \quad (3.3)$$

Notice that once we specify a home production function, this first-order condition, together with our estimates of $\frac{\partial s}{\partial p}$, $\frac{\partial s}{\partial Q}$, p , Q , and h , will allow us to estimate the parameters of the home production function.

To see why the availability of the price data is crucial to estimating the home production function, consider the case in which we do not observe prices (or in which we assume that every household faces the same price). Estimation would rely on the fact that the MRT between time and goods in home production equals the relative price of time and goods [that is, we assume that prices are fixed and use the last two conditions in (3.2)]. The price of time would have to be inferred either from wages or leisure. The former is problematic, because many households have a single earner and the wage of the sole earner is not necessarily the opportunity cost of time of the home producer. Even with two-earner families, it is not clear that workers have the ability to vary labor supply smoothly at the margin. Imputing the cost of time from leisure

requires the measurement of leisure (usually taken as a residual) and knowledge of preferences over leisure—both questionable undertakings.

Our approach requires only that the opportunity cost of time for the shopper equal the opportunity cost of time for the home producer, a much more plausible assumption. Moreover, it strikes us as reasonable that households can adjust smoothly between the shopping and home production margins.

We restrict our home production function to have a constant elasticity of substitution between time and market goods:

$$f(h, Q) = (\psi_h h^\rho + \psi_Q Q^\rho)^{\frac{\alpha}{\rho}}, \quad (3.4)$$

where the elasticity of substitution is given by $\sigma = \frac{1}{1-\rho}$. We allow the function to be homogeneous of arbitrary degree α , although we will not be able to identify this parameter. Given (3.4), the MRT between time and goods from the home production function [left-hand side of (3.3)] is:

$$\frac{\psi_h}{\psi_Q} \left(\frac{h}{Q} \right)^{\rho-1}. \quad (3.5)$$

Substituting (3.5) into (3.3) and taking logs on both sides (and rearranging), we have:

$$\ln(h/Q) = \sigma \ln\left(\frac{\psi_h}{\psi_Q}\right) - \sigma \ln\left(-\frac{\partial p}{\partial s} / \left(\frac{\partial p}{\partial Q} Q + p\right)\right). \quad (3.6)$$

We construct the empirical counterpart of (3.6) by fitting the MRT in shopping from our price data, using the coefficients reported in Table 3, Column 1. Specifically, we use the estimated elasticities, together with observations on p and Q , to compute the last term on the right-hand side of (3.6).

Constructing the left-hand side is more difficult. Unfortunately, our price data do not contain time spent in home production. To get around this issue, we merge data from Homescan with data from ATUS, by creating demographic cells in both data sets, using age, sex, marital status, and education. Specifically, we use eight age ranges (those displayed in Figures 1 and 2), four education categories (less than high school, high school, some college, and

college or more), two marital status categories, and two sex categories. The demographic variables are those reported for the household head. Adjusting for the fact that not all combinations are represented, we have 92 separate cells. For each cell in the ATUS data set, we calculate the sample average of time spent in food production and total home production and merge this into the Homescan data set.

We combine this estimate of time spent in home production with the household's Q to obtain the left-hand side of (3.6). Note that while we have variation at the household level for p and Q , the measure of time use varies only according to our demographic cells. We therefore collapse each cell and run a "between-effects" regression. Averaging over a number of households in each demographic group should reduce the errors-in-variables inherent in our data. The averaging will also correct for idiosyncratic "productivity" shocks that are uncorrelated with demographics. Note that we are imposing the constraint that all demographic cells face the same production functions. This may be problematic to the extent that the quantity of "home capital" may vary across cells. However, the Homescan database contains dummy variables for presence of home durables (microwave, dishwasher, garbage disposal, etc.). Inclusion of these dummy variables does not alter the results. Therefore, we report the specifications without these controls, given the desire to preserve degrees of freedom.

Estimating equation (3.6) using information from the 92 demographic cells yields an estimate of $\sigma = 1.2$, with a standard error of 0.1.²¹ We perform the same analysis using the broader measure of time spent in home production (all housework, not just food preparation) as our measure of h , and estimate an elasticity of 1.3 with a standard error of 0.1. These estimates are reported in Columns 1 and 2 of Table 4, respectively. In both cases, we can reject the hypothesis of an elasticity of 1 at standard confidence levels. The fact that σ exceeds 1 has important implications for the impact of home production in many macroeconomic models [for example, see Benhabib, Rogerson, and Wright (1991) and Aguiar and Hurst (2005)].

One concern with the estimates in Columns 1 and 2 is that some of the demographic cells have few households (the minimum number of observations per cell is six). This may

²¹ Given the fact that we are using a generated regressor, we bootstrap all standard errors for Table 4, clustering on households and including the first-stage estimation of the shopping function in each repetition.

result in significant measurement error. To correct for this, we run a between-effects regression, using the eight age groups as our cells. We find an elasticity of 2.5 for food production and 2.7 for total home production, both with standard errors of 0.2. The estimates are reported in Columns 3 and 4 of Table 4. These cells are much larger, with a minimum observation count per cell of 2,449. The larger estimates may be indicative of attenuation bias in the specification of Columns 1 and 2.

For comparison, Rupert, Rogerson, and Wright (1995) report an elasticity of substitution between home and market goods, a measure roughly comparable to our elasticity, of 1.8 for single, employed women.²² This number is in line with our estimates. Moreover, restricting our sample to include only single women produces an estimated elasticity of 1.5. Rupert et al.'s parameter estimates for other demographic groups are generally imprecisely estimated (or implausible). This highlights the difficulty of relying on wages to value time for complex family structures and underscores the value of the price data. It is interesting that our estimates and theirs coincide for employed single women, a demographic for which wage is most plausibly the relevant price at the margin. Several studies have used equilibrium models and aggregate data to back out an elasticity of substitution for home production that is close to our estimates, which are based on micro data. For example, McGrattan, Rogerson, and Wright (1997) estimate an elasticity of 1.3 and Chang and Schorfheide (2003) estimate an elasticity of 2.3.

One concern with (3.6) is that Q is present in both the left-hand and (inversely) the right-hand sides of the regression. To the extent that Q is mismeasured, this may artificially imply a negative correlation and bias our estimate of σ upward. To check whether this is an issue, we run:

$$\ln(h) = \sigma \ln\left(\frac{\psi_h}{\psi_Q}\right) - \sigma \ln\left(-\frac{\partial p}{\partial(s^* + s)} \bigg/ \left(\frac{\partial p}{\partial Q} Q + p\right)\right) + \ln(Q). \quad (3.7)$$

The estimate of σ in this case is 2.5, with a standard error of 0.5, a result roughly the same as that found above. This specification also allows a test of whether the coefficient on $\ln(Q)$ is 1

²² Specifically, the interpretation of the elasticity of Rupert et al. (1995) is the same as ours if their home good is a linear product of time input, and if market work and home work are perfect substitutes over leisure in the utility function. This parameterization is consistent with their estimates.

(essentially a test of homotheticity). The estimated coefficient on $\ln(Q)$ is 1.0 with a standard error of 0.3. These results are reported in Column 5 of Table 4.

4. Lifecycle Consumption versus Lifecycle Expenditure

With a parameterized home production function, we can compare how lifecycle expenditure (an input into the home production function) compares with lifecycle consumption (the output of the home production function). To do this, we fit (3.4) over the life cycle (using the parameters from Column 3 of Table 4). Going from the ratios (the MRT) to levels requires us to assume a value for returns to scale, which we take to be 1. It is also the case that we can only estimate the ratio $\left(\frac{\psi_h}{\psi_Q}\right)$, so we set the denominator equal to 1. This assumption involves only a scaling of consumption and does not play a role in the analysis, once we normalize by young households.

The path of lifecycle consumption for the household is plotted in Figure 4. As before, we plot log deviations from households with heads aged 25-to-29. Household consumption has a “twin peaks” shape. Consumption rises rapidly early in the adult lifecycle, peaking around 40, then declines until late middle age before rising through retirement. Household consumption’s peak for households with heads in the late thirties is 26 percent higher than that for households headed by adults 25-to-29 (p-value <0.01), 9 percent higher than that of households with heads in their early fifties (p-value 0.08), and 3 percent higher than that of households with heads in their sixties (p-value 0.50). Again, given that family size is largest for households with heads around 40, it is not surprising that household consumption is largest in middle age.

To control for changing family size over the life cycle, Figure 5 plots the ratio of household consumption to household expenditure. Note that any proportional scaling factor due to changing household size is accounted for by the ratio. This figure highlights the finding that households use different ratios of time and market goods in consumption over the life cycle. We see that the ratio is at its lowest in middle age, when the price of time is highest. Moreover, the ratio increases dramatically in retirement. This occurs simultaneously with the

well-documented decline in expenditure during retirement.²³ As discussed in the next section in the context of a model, this results from two margins of substitution. First, as time is relatively cheap during retirement, households substitute away from market expenditures and toward time in producing consumption goods, lowering the denominator of Figure 5. Second, the total cost of consumption (inclusive of time) is relatively low in retirement. Households, therefore, have an incentive to delay consumption until retirement, raising the numerator.

5. A Lifecycle Model

This paper has documented a number of empirical facts that shed light on how households allocate their time to reduce expenditure over the life cycle. In this section, we embed the considerations raised by the price and time-use data into an otherwise-standard lifecycle model. In this fashion, we can view the empirical regularities in a single, coherent framework. We also demonstrate that the primary features of the data are consistent with the augmented lifecycle model, particularly for behavior observed from early middle age through old age.

Consider a household comprised of two adults, indexed by $i=1,2$. Where no index is used, this implies we have summed across adults to report a household-level variable. We denote the age of the household by a single index t , which runs from zero, when the household is formed, through T , when the adult members of the household die. At age t , the household also includes $n_t(\tau)$ children of age τ . Let n_t denote the age- t household's vector of number of children aged τ . We take the arrival of children as exogenous. There is no uncertainty in the model.

5.1. Preferences

Agents have preferences over consumption and leisure and seek to maximize total discounted utility over the life cycle:

²³ See Bernheim, Skinner, and Weinberg (2001) and Aguiar and Hurst (forthcoming).

$$\sum_{t=0}^T \beta^t U(C_t, \tilde{C}_t, l_{1,t}, l_{2,t}; n_t), \quad (5.1)$$

where C is household food consumption, \tilde{C} is household consumption of other goods, l_i , $i=1,2$, is individual i 's leisure, and β is the intertemporal discount factor.

Consumption is the product of the home production function discussed and estimated in Section 3. This function was estimated for a subset of goods, Q —namely, food items captured by Homescan. We assume that utility is separable in food and other goods. This allows us to model in partial isolation decisions regarding food expenditures and time spent shopping for and preparing food. The purchase of other goods and time spent shopping for other goods enter only through the budget constraints. To account for other goods, let φ_Q denote the fraction of total expenditures captured by our Homescan goods. Similarly, let φ_S and φ_H denote the fraction of total shopping and home production time, respectively, accounted for by food. We assume that these shares are invariant to the level of expenditures and the amount of time spent in home production.²⁴ Total expenditures in terms of time and money are then constant multiples of expenditures on food.

We further assume that utility is separable between consumption and leisure.²⁵ In this fashion, we can highlight the distinction between separability of consumption and leisure in the primitive utility function and the ability to substitute time for money through shopping and home production. The combination induces a reduced form in which time and expenditures enter nonseparably. We feel the distinction is useful to understand the micro foundations behind reduced-form nonseparability. Specifically, period utility is given by

$$U(C, \tilde{C}, l_1, l_2; n) = u(C; n) + \sum_i v(l_i) + u(\tilde{C}; n). \quad (5.2)$$

The family composition vector, n , enters as a taste shifter. We implement this by defining “per capita” consumption $c_t \equiv \frac{C_t}{N_t}$, where $N_t = \left(n_A + \sum_{\tau} \alpha_{\tau} n_{\tau}(\tau) \right)^{\eta}$. n_A is the number of

²⁴ This assumption implicitly assumes similar elasticities of substitution between time and market inputs across goods. While this condition is unlikely to hold precisely in practice, we impose it as a tractable approximation.

²⁵ Keep in mind that these separability assumptions pertain to the model. For the empirical analysis, we made no assumptions regarding separability in preferences. The separability assumption required for the empirical results pertains to home production and shopping.

adults in the household and, as noted before, $n_t(\tau)$ is the number of children aged τ . The parameter α_τ is the relative weight in consumption of a child of age τ to an adult. This specification of adult equivalencies was suggested by Banks and Johnson (1994). The parameter η captures returns to scale in household consumption. Given the functional form assumptions, we have an extra degree of freedom in setting the parameters governing returns to scale in home production, returns to scale in the adult equivalency schedule, and the elasticity of intertemporal substitution (*EIS*) in consumption (discussed below). We select the normalization that the home production function exhibits constant returns to scale, and we adjust the other two parameters accordingly.

We follow standard practice and select iso-elastic utility functions for leisure and consumption. Specifically,

$$\begin{aligned} u(C; n) &= \frac{c^{1-\gamma}}{1-\gamma} \\ v(l_i) &= \theta \frac{l_i^{1-\nu}}{1-\nu} \end{aligned} \tag{5.3}$$

The parameter γ is the *EIS* for consumption, and ν is the corresponding elasticity of leisure. The parameter θ governs the relative weight of leisure in utility.

5.2. Budget Sets

Each adult in the household allocates his or her time over a number of tasks. To simplify the analysis, we treat claims on time due to market work, children, sleep, etc., as exogenously determined. Treating labor supply decisions as exogenous is a simplification. However, to model labor supply adequately over the life cycle, we would need to account for the fact that workers in their late twenties and thirties are acquiring skills and experience on the job that will be reflected in future wages. This consideration would be necessary to help explain why wages are fairly symmetric around middle age, but hours are asymmetric (younger workers put in more hours than workers near retirement with similar wages). Moreover, it is not evident that workers are able to adjust market hours freely at the margin. As our focus is on shopping and home production—margins that can more plausibly be adjusted freely—we endow adult i in a

household of age t a total of $H_{i,t}$ units of time that can be allocated to shopping, home production, and leisure. The remaining time is exogenously committed to market work, childcare, sleep, personal care, etc. The budget constraint for time is, therefore, given by

$$\frac{s_t^i}{\varphi_s} + \frac{h_t^i}{\varphi_H} + \ell_t^i \leq H_i(t), \quad i = 1, 2, \quad (5.4)$$

where we have scaled up time spent shopping for and home producing food to account for the corresponding time devoted to other goods.

The household has access to borrowing and lending at the interest rate r . Given that labor income is exogenous, we can collapse the budget constraint into

$$\sum_{t=0}^T (1+r)^{-t} p_t(s, Q) Q_t \leq \varphi_Q A, \quad (5.5)$$

where A is the net present value of labor income plus any initial assets, φ_Q is the share of expenditures on food, and we have made explicit the assumption that price is given by the shopping function $p=p(s, Q)$, estimated using (2.8).

5.3. First-Order Conditions and the Lifecycle Profile of Expenditure

The household's problem is to maximize (5.1) subject to the budget constraint (5.5), the time constraint, the home production and shopping technologies, and non-negativity constraints on all choice variables. The first-order conditions associated with the problem are

$$\beta^t (1+r)^t u'(c) N^{-1} f_Q = \lambda (p + p_Q Q), \quad (5.6)$$

$$u'(c) N^{-1} f_h = v'(l_i) \text{ if } h_i > 0, \quad (5.7)$$

$$\beta^t (1+r)^t v'(l_i) = -\lambda p_s Q \text{ if } s_i > 0. \quad (5.8)$$

One question that arises is: How does the change in the opportunity cost of time influence the lifecycle profile of expenditures for a given level of lifetime resources? Consider the case in which the market price of goods is constant (that is, there is no shopping function). As the cost of time [$v'(l)$] increases, all else being equal, the ratio of f_Q to f_h decreases [this can be seen from the ratio of (5.6) to (5.7)]. For a constant-returns-to-scale (CRS) home production function, this implies that the ratio Q/h increases (that is, the agent substitutes goods for time).

To satisfy (5.6), $u'(c)/N$ must decrease. All else being equal (including family size), this implies that household consumption is greatest when the cost of time is lowest. That is, adjusted for family size and impatience, consumption is highest in retirement. This is consistent with evidence documented in Aguiar and Hurst (forthcoming), which uses food diaries to show that retirees eat better than nonretirees along a number of dimensions.

But what about expenditures? Using the iso-elastic functional forms for $u(c)$ and $f(Q,h)$ and continuing to assume price is fixed, we have,

$$Q^{-\gamma} \left(\bar{h} \left(\frac{Q}{h} \right)^{-\rho} + 1 \right)^{\frac{-\gamma+1-\rho}{\rho}} = \lambda p. \quad (5.9)$$

Note that $\left(\bar{h} \left(\frac{Q}{h} \right)^{-\rho} + 1 \right)^{\frac{-\gamma+1-\rho}{\rho}}$ is increasing in Q/h if and only if $\gamma > 1 - \rho$. Or, in other words, it

is increasing in Q/h if and only if the *EIS* of consumption, $(1/\gamma)$, is less than the intra-temporal elasticity of substitution between time and goods in home production, $\sigma = 1/(1-\rho)$. To get expenditures, multiply through by p^γ ,

$$(pQ)^{-\gamma} \left(\bar{h} \left(\frac{Q}{h} \right)^{-\rho} + 1 \right)^{\frac{-\gamma+1-\rho}{\rho}} = \lambda p^{1-\gamma}. \quad (5.10)$$

Therefore, holding p and λ constant, expenditures increase with the price of time if $1/\gamma < \sigma$. The intuition behind this relationship is the following. An increase in the price of time provides an incentive to purchase more market goods and less time as inputs into home production for a given level of consumption. However, the total cost of consumption is relatively high when time is scarce, providing an incentive to reduce consumption (and market goods as an input) in those periods. Which effect dominates depends on the relative elasticities of substitution. See Ghez and Becker (1975) for a related discussion.

In our framework, we allow the price of goods to vary with the cost of time, as well. However, this does not imply a dramatically different interpretation of the response of expenditures to the price of time. Specifically, we replace (5.9) with

$$(pQ)^{-\gamma} \left(\bar{h} \left(\frac{Q}{h} \right)^{-\rho} + 1 \right)^{\frac{-\gamma+1-\rho}{\rho}} = \lambda (\xi_Q + 1) p^{1-\gamma}, \quad (5.11)$$

where $p = \phi s^{\xi_s} Q^{\xi_Q}$. Sufficient conditions for expenditures to increase with the cost of time are then $\gamma > 1/\sigma$ (as before), $\xi_Q > -\xi_s$, and $\gamma \geq 1$.²⁶ The same conditions are also sufficient for price to increase with the cost of time, holding constant lifetime resources. Our estimates of the shopping and home production functions (plus the many studies that estimate $\gamma \geq 1$) suggest these conditions hold empirically.

In summary, a relatively low *EIS* for consumption and the changing cost of time directly imply a “hump” in expenditures over the life cycle. This is driven solely by the opportunity cost of time and the ability to substitute time for expenditure in shopping and home production. In a sense, there is a similarity to Heckman’s (1974) explanation of the lifecycle profile of expenditures as stemming from a nonseparability between consumption and leisure in utility. However, in our present framework, the nonseparability is between market expenditures and time in home production and shopping and is, therefore, more directly tied to the analysis of Becker (1965) and Ghez and Becker (1975). We should also reiterate that the “hump” in expenditures does not necessarily reflect a “hump” in consumption (see Figure 5), and that it is perfectly consistent with rational, patient agents with access to complete markets.

5.4. Results

The parameters used to calibrate the model are reported in Table 5 and discussed in detail in Appendix B. We calibrate to married households and consider the empirical counterpart to the age of the household to be that of the male head. The simulated lifecycle profile implied by the model is displayed in Figures 7 through 11, along with the corresponding data. The well known lifecycle “hump” in expenditures is present for food expenditures from the Homescan data [X from (2.1)]. As Figure 6 indicates, the model tracks these data closely over the life cycle. However, the model overpredicts expenditures early in the life cycle by a few percentage points, and, therefore, the model’s hump is slightly shallower than the data’s. While

²⁶ Proof is omitted but available from the authors upon request.

the model's additional expenditure is fairly small, it does suggest that there may be room for borrowing constraints or habit formation that suppresses expenditure early in the life cycle. Note, however, that the model captures the dramatic decline in expenditures in later middle age and retirement. This decline is a combination of declining family size and the falling opportunity cost of time. As noted above, the fact that agents are more inclined to substitute time for goods within a period than to substitute consumption across periods implies that expenditures track the price of time.

To see that the profile of expenditures is a rather poor guide to lifecycle consumption, consider Figure 7. In this figure, we plot the ratio of household consumption to household expenditure as predicted by the model. For comparison, we include the consumption-to-expenditure ratio as found in the data, evaluated using the calibrated home production parameters.²⁷ For the model, consumption is calculated using the calibrated home production function plus the predicted inputs of market goods and time. In both the data and the model, consumption is relatively low in middle age and high late in life. This reflects both the intra-temporal and intertemporal margins of substitution. Time is at a premium in middle age, and agents will substitute toward market goods along any consumption isoquant, lowering the ratio of consumption to expenditure. Moreover, consumption is cheap when time is cheap, and agents will accordingly substitute away from consuming in middle age and towards consuming in retirement. The fact that household consumption rises relative to expenditure after middle age requires a careful interpretation of the familiar empirical regularity that expenditure declines dramatically after middle age. Naively extrapolating this series into a decline in consumption overlooks the dramatic shift in the allocation of time away from the market and towards home production that occurs simultaneously.

As documented in Section 2, the empirical hump in expenditures is associated with a hump in prices. The model yields this prediction as well, as shown in Figure 8. As with expenditure, the model's hump is somewhat shallower early in the life cycle. All else being equal, the fact that households with young heads in the model consume more than they do in

²⁷ The "data" series in Figure 7 differ slightly from those in Figure 5, as the former calibrate the weight on time in the home production function (ψ_h) as 0.9, while the latter use 1.1. See Appendix B for details.

the data is reflected in the pattern for prices (given that $\xi_Q > 0$). The decline in prices as household heads age toward retirement is again nicely demonstrated in the model.

Behind this pattern of prices is varying time spent shopping (Figure 9). Specifically, the model predicts that households with middle-aged and elderly heads will shop intensively. The behavior of the former is due to larger family size; the behavior of the latter reflects a lower cost of time. This “twin” peak in shopping time translates into the single peak in prices, as the additional time spent shopping in middle age is more than offset by the need to purchase a larger basket of goods. As discussed in the Appendix, we calibrate the relative weight on leisure in utility (θ) to match the amount of time spent shopping by households with middle-aged heads. Therefore, the fact that, in the model, households with heads aged 40-to-44 spend as much time shopping as their empirical counterparts is a product of calibration. However, the shape of the lifecycle profile is not determined solely by this parameter. The model also captures the rough features of the lifecycle profile of home production observed in the data (Figure 10), although at a higher level. In short, the estimated elasticities for the shopping and home production functions, when fed into the model, yield lifecycle profiles for shopping and home production that match the empirical patterns quite closely.

5.5 Sensitivity Analysis

To shed light on how predicted behavior in the model changes with parameters, we explore three alternative parameterizations. The results for expenditure, the ratio of consumption to expenditure, and prices are depicted in Figure 11.

The first alternative (parameterization “a” in Column 3 of Table 5) lowers the elasticities in the home production and shopping functions by roughly one-half. Specifically, the elasticity of substitution between time and goods in home production is lowered from 2.5 to 1.1. Similarly, the elasticities of price with respect to shopping time and goods purchased are lowered to -0.05 and 0.10, from -0.11 and 0.21, respectively. From Figure 11, we see that the lifecycle profile of expenditure remains roughly the same as in the benchmark. As would be expected, the largest departures occur for the ratio of consumption to expenditure and prices over the life cycle. In particular, the relative increase in consumption during retirement is muted

because of the lower elasticity of substitution in home production. Similarly, the lifecycle path of prices is flatter than in the benchmark.

Parameterization “*b*” alters the adult equivalence scales. Specifically, all children, regardless of age, are considered to be equivalent to 0.5 adult in consumption. Moreover, we set the returns-to-scale parameter to 1. From Figure 11, Panel A, we see that this raises expenditure early in the life cycle. This is a direct result of the increased relative weight on infants and toddlers, whose numbers peak in households with heads in their late twenties (see Figure A1). The increased purchases early in the life cycle lead to slightly higher prices paid, as well. Although not depicted, households do shop more intensively early in the life cycle, but not sufficiently to offset the price effect of the larger quantity of market goods purchased. The increased expenditures early in the life cycle necessarily lower expenditures later in life (as financial resources are fixed). Households compensate later in life by increasing time spent on home production (not depicted), raising the ratio of consumption to expenditure later in the life cycle (Panel B).

The last parameterization, “*c*”, raises the intertemporal elasticities of substitution for consumption and leisure. Specifically, both elasticities are set to 0.67, or $\gamma = \nu = 1.5$. As discussed above, the *EIS* of consumption plays a role in the lifecycle profile of expenditures. As this elasticity increases, households are more willing to delay consumption until retirement, when time is cheap. This can be seen in Panel A of Figure 11, which shows that expenditures in retirement are noticeably higher than under the other parameterizations. The ratio of consumption to expenditure during retirement remains high relative to its value in middle age, but it is slightly lower than in the benchmark. The increased level of expenditures leads to higher prices during retirement relative to the benchmark.

In summary, the alternative parameterizations indicate that large price elasticities with respect to time and quantity are useful to generate the sharp hump in prices seen in the data. However, the consumption and expenditure series do not appear to be overly sensitive to lowering the home production elasticity close to 1. Adult equivalency scales that place a higher weight on teenagers than toddlers help to generate the relatively low expenditure early in the adult lifecycle (and correspondingly higher expenditure in middle age). Finally, the sharp

decline in expenditure during retirement suggests a low intertemporal elasticity of substitution. That is, the decline in expenditures late in the life cycle is inconsistent with a strong willingness to delay consumption until it is cheapest (that is, retirement).

6. Conclusion

This paper has estimated the elasticity between time and money due to shopping and home production. We find that households can and do alter the relationship between expenditures and consumption by varying time inputs. Moreover, they do so in a way that is consistent with standard economic principles.

This paper has contributed some new data and insights regarding these margins of substitution. However, the data have some limitations. The scanner data set consists of a subset of grocery items. We cannot state whether similar patterns hold for other goods. The time-use data do suggest that households shop and engage in home production for nonfood goods. Nevertheless, the results in this paper should be considered only suggestive of how households exploit time in the consumption of goods other than food. Moreover, the data are cross-sectional in nature, and, therefore, we must be cognizant that some of our lifecycle results may be confounded with cohort effects. However, cohort effects are less likely to be an issue for normalized variables, such as the ratio of consumption to expenditure and the dispersion of prices within a household over time, than for non-normalized variables.

There is a growing interest in the role of nonmarket activities and the allocation of work between the market and the household. The insights of household production have already proved fruitful in explaining phenomena as disparate as baby booms and business cycles. While our focus has primarily been on lifecycle consumption, we feel that the data and analysis presented in this paper support the broader emphasis on how time is spent outside of market labor.

Appendix A: Issues Related to the Homescan Data Set

This appendix discusses and quantifies a number of potential concerns related to the Homescan data set. First, there is a potential issue with the extent to which households actually scan in the products they purchase. Within our Homescan data, the average monthly expenditure for packaged goods scanned is \$176 per month, expressed in current dollars. The comparable figure for “food at home” reported in the 1993 and 1994 waves of the Panel Study of Income Dynamics (PSID) is \$323 per month. This implies that the Homescan data set covers a little more than half of total grocery expenditures reported in the PSID. The difference between the Homescan data set and the PSID data set likely comes from two sources. First, the Homescan data set does not include meat, fresh foods, or vegetables. Moreover, as discussed below, it may be the case that households fail to scan all grocery items into the Homescan database.

Second, there is a potential issue with attrition from the Homescan sample over time. A direct assessment of the magnitude of attrition on the extensive margin is complicated by the fact that ACNielsen drops data from households that quickly withdraw from the survey.²⁸ However, we can observe attrition directly on the intensive margin. On average, a household reported 1 percent lower expenditure in the first quarter of 1994 than it did during the first quarter of 1993, and 5 percent lower expenditure in the first quarter of 1995 compared with the same quarter in 1994. The failure to record all transactions is not crucial to many of the facts regarding price dispersion documented in this paper, as long as the transactions a household does record are representative of that household’s purchases (that is, as long as the omissions are random within a household). However, the failure to record all transactions may influence such items as total expenditures and frequency of shopping. For each of our analyses, we compared estimates using only the first quarter of the sample with those obtained using the sample from the last quarter, and we did not uncover substantial differences.

More importantly, the decline in household expenditure over the sample does not appear to vary with such demographics as age and education, suggesting that attrition is not

²⁸ Within the data set, roughly two-thirds of the households are present for at least 16 months of the survey, and over half remain for the entire 27 months.

highly correlated with our key controls. In a regression of the month-to-month decline in expenditure on age and time dummies, the p-value of the test that all (seven) age dummies are zero is 0.44.²⁹ In a regression of the month-to-month decline in expenditure on education and time dummies, the p-value of the test that all (five) education categories are zero is 0.78.

Taken together, these results indicate that the rate of attrition is fairly constant across demographic groups. However, the initial level of under-reporting (potential issue 1) appears to be correlated with education, but not with age. To assess this, we compared expenditures in Homescan with those in the PSID. Specifically, we created cells in the PSID by age of head (using the eight categories), education (less than high school, high school, some college, and college or more), and year (1993, 1994, 1995). For each cell, we calculated the average expenditure on food at home reported in the PSID, and we merged these values into the Homescan data set. We then constructed the ratio between Homescan households and their corresponding PSID cells. This gap shows no correlation with the age of the household head (the p-value of the F test is 0.56). However, the gap is correlated with education. For example, reported expenditure in Homescan for households with a college-educated head is, on average, 42 percent of that reported in the PSID. The comparable share for high school graduates is 55 percent (p-value of difference <0.01). This suggests that households with more-highly educated heads are less likely than others to scan all purchases (or that they buy more goods outside the scope of Homescan, such as meat and produce). Again, for the main analysis, as long as the scanned items are representative of the household's purchases, this will not generate a bias. However, because of these results, we do not sum the Homescan transactions to infer how shopping frequency and total expenditures vary with education or income.

Another issue is that we treat the Denver metropolitan area as a single market. It may be the case that there is extensive market segmentation due to income, age, and family composition. However, our data set includes specific store and chain identifiers. Within the data, 83.6 percent of purchases were made at grocery stores, 4.1 percent were purchased at discount stores, 3.1 percent at price clubs, 1.7 percent at convenience stores, and 1.5 percent at drug stores. The remaining purchases occurred at specialty stores, including liquor stores, gas

²⁹ As discussed in the text, we analyze eight age ranges: 25–29, 30–34, 35–39, 40–44, 45–49, 50–54, 55–64, and 65+.

stations, vending machines, pet stores, etc. We find that households do shop at multiple chains within a year.

Some may be concerned that the quality of a purchase may not be perfectly proxied by the UPC. For example, high-income individuals may shop at high-end grocery stores (like Whole Foods) because the store displays are nicer or because they have access to a wider variety of high-quality goods. The higher price for a specific UPC good at such a store may be the result of better store amenities. However, the nature of the data suggests that this is not an issue. Eighty-five percent of all products purchased at grocery stores were purchased within four grocery chains of similar quality: Albertsons, King Sooper, Safeway, and Cubes Food.

Nevertheless, to address this concern formally, we adjusted our price index for differences across chains in mean prices for each good, and we found no substantial change in the results. Specifically, for a good i purchased in month m at chain k , we calculated the average price of good i sold in that chain over the relevant quarter (we averaged over a quarter rather than a month to ensure that the average comprises a reasonable number of purchases). We then calculated the cost of a basket purchased within the month as if each good had been purchased at the relevant chain's average price. This quantity was used in place of (2.4). We found no substantial differences in the patterns described in the text using this alternative index. In other words, controlling for grocery-chain effects cannot explain the results presented in this paper.

Appendix B: Calibration

This appendix describes the details behind the calibrated parameters reported in Table 5 and used in Section 5. We take a household to consist of two adults and calibrate to married households in the data. We use the age of the male head as our empirical counterpart of household age, where the lifecycle begins at age 25 and ends at age 81. A time interval is taken to be a year, and we set the annual interest rate to 2 percent. The time preference parameter β is set equal to the inverse of 1 plus the interest rate, or 0.98. This implies that agents would like to maintain constant marginal utilities of consumption and leisure over the life cycle, all else being equal. The parameter γ is set to 5, or an *EIS* of 0.2, which is in line with many empirical estimates [see, for example, Browning and Lusardi (1996) and Browning, Hansen, and Heckman (1999)].

We have less empirical guidance on the elasticity of leisure, ν . As our model is a partial equilibrium model, the appropriate elasticities are those observed in micro data. However, the vast majority of studies estimate the elasticity of *market labor*. This would not pose a major obstacle if labor varied one-for-one with leisure. However, if leisure is considered to be a good providing utility (as it is in the model) rather than the complement of market labor, then we must consider how nonmarket time is allocated to such activities as shopping and home production. Aguiar and Hurst (2005), document that the fairly stable level of market hours over the last 40 years masks dramatic changes in leisure. Perhaps the study that is closest in spirit to estimating the intertemporal elasticity of leisure is Heckman and McCurdy (1980). However, there again, leisure in the cross-section is assumed to increase minute-for-minute with declines in market work. These caveats aside, our reading of the labor literature implies a plausible estimate of ν to be 3. That is, a 1-percent increase in the price of time induces a 0.33 percent increase in leisure. In our ATUS sample, market hours for married men aged 25-to-55 is slightly higher than reported leisure hours, implying a Frisch labor elasticity evaluated at the mean of roughly 0.2.³⁰ For women, reported time spent in market labor is less than time spent in leisure, implying a labor supply elasticity closer to 0.4. We choose θ , the parameter governing the

³⁰ We define leisure as time spent in active recreation, socialization, entertainment, relaxing, and civic and religious activities. See Aguiar and Hurst (2005) for a more detailed analysis.

relative importance of leisure in utility, so that the time spent shopping by households with heads aged 40-to-44 in the model line up with the data.

The household consists of two adults plus children placed into three age groups, 0-to-5, 6-to-12, and 13-to-18, corresponding to $\tau=1,2,3$. The number of children in each age range is calibrated to the life cycle of married-couple households reported in the 2000 Census. The three series are plotted in Figure A1. The weights α_τ determine the consumption of children of various ages relative to that of adults. There is no single schedule of “adult equivalents” uniformly used in the literature. We should also point out that we are scaling consumption rather than expenditures, while many of the studies generating equivalence scales relate to expenditures. Given the scant guidance from the literature, we somewhat arbitrarily set the relative consumption weights to 0.1, 0.5, 1, for the three age ranges of children. We set the “returns-to-scale” parameter, η , equal to 0.9, which implies mild positive returns to scale to household size. We discuss the sensitivity of the results to these parameters in the robustness section.

We set the expenditure “share” parameter, ϕ_Q , to match the ratio of average expenditure in the Homescan database and total nondurable expenditures reported in the Consumer Expenditure Survey. The data from ATUS indicate that the mean time spent home producing food is roughly one-quarter to one-third of total time spent on housework. A similar ratio holds between shopping for food and total shopping time. The parameters ϕ_H and ϕ_S are both set to 0.3, accordingly. The present value of total lifetime resources, A , is calibrated to lifetime expenditures. Specifically, we scale the Homescan expenditures by ϕ_Q and then discount to age 25 using the age of the household head and an annual interest rate of 2 percent. This value is \$1.8 million dollars, expressed in terms of average prices for the Homescan period.

The elasticity of substitution in home production is set to 2.5, in line with the estimates reported in Table 4. The scale parameter, ψ_h , is calibrated so that the MRT between time and goods in shopping equals that in home production when evaluated at the empirical means of shopping time, home production time, and market goods purchased for households with heads aged 40-to-44. The resulting value is 0.9. An alternative is to use the estimate of the intercept of (3.6), which is 1.1. Note that the latter value equates the mean $\ln(\text{MRT})$ in shopping and home

production, which in general will differ from equating the MRTs evaluated at the sample means. For the shopping function, we assume a log-linear functional form in shopping time and goods. Guided by the estimates reported in Column 4 of Table 4, we set the elasticity with respect to shopping to -0.11 and with respect to goods to 0.21. Recall that the estimates in Table 4 used shopping frequency (trips per month) rather than time as the regressor. For the model, we assume that shopping time per trip is constant and adjust the intercept (in logs) of the price function so that $\ln(p)=0$, evaluated at the average frequency of trips and quantity purchased per month in the Homescan data.

The endowment of time for each adult is obtained from the ATUS. Specifically, for each age (of the male household head), we take the sum of time allocated to home production, shopping, and leisure for married men and women and average across households. The two series are plotted in Figure A3.

References

- Aguiar, Mark, and Erik Hurst. Forthcoming. "Consumption vs. Expenditure." *Journal of Political Economy*.
- Aguiar, Mark, and Erik Hurst. 2005. "Measuring Leisure: Evidence from Five Decades of Time Use Surveys." University of Chicago Mimeo.
- Attanasio, Orazio, James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Lifetime Consumption." *Journal of Business and Economic Statistics* 17(1) January: 22-35.
- Banks, James, and Paul Johnson. 1994. "Equivalence Scale Relativities Revisited." *Economic Journal* 104 July: 883-890.
- Baxter, Marianne, and Urban J. Jermann. 1999. "Household Production and the Excess Sensitivity of Consumption to Current Income." *American Economic Review* 89(4): 902-920.
- Becker, Gary. 1965. "A Theory of the Allocation of Time." *Economic Journal* 75(299): 493-508.
- Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy* 99(6): 1166-1187.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth among U.S. Households? " *American Economic Review* 91(4): 832-857.
- Blundell, Richard, Martin Browning, and Costas Meghir. 1994. "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." *Review of Economic Studies* 61:57-80.
- Brown, Jeff, and Austan Goolsbee. 2002. "Does the Internet Make Markets More Competitive: Evidence from the Life Insurance Industry." *Journal of Political Economy* 110(3) June: 481-507.
- Browning, Martin, and Annamaria Lusardi. 1996. "Household Savings: Micro Theories and Micro Facts." *Journal of Economic Literature* 34(4): 1797-1855.
- Browning, Martin, Lars Peter Hansen, and James Heckman. 1999. "Micro Data and General Equilibrium Models." In *Handbook of Macroeconomics*, Vol 1A , eds. John Taylor and Michael Woodford. New York: Elsevier Science, North-Holland.
- Campbell, John, and Sydney Ludvigson. 2001. "Elasticities of Substitution in Real Business Cycle Models With Home Production." *Journal of Money, Credit and Banking*. 33(4) November: 847-875.
- Chang, Yongsung, and Frank Schorfheide. 2003. "Labor Supply Shifts and Economic Fluctuations." *Journal of Monetary Economics* 50(8): 1751-1768.
- Cronovich, Ron, Renae Daneshvary, and R. Keith Schwer. 1997. "The Determinants of Coupon Usage." *Applied Economics* 29:1631-41.
- Ghez, Gilbert, and Gary Becker. 1975. *The Allocation of Time and Goods over the Life Cycle*. New York: Columbia University Press.
- Gourinchas, Pierre-Olivier, and Jonathan Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70(1) January: 47-89.

- Greenwood, Jeremy, Richard Rogerson, and Randall Wright. 1995. "Household Production in Real Business Cycle Theory." In *Frontiers of Business Cycle Research*, ed. Thomas Cooley. Princeton: Princeton University Press.
- Greenwood, Jeremy, Ananth Sesahdri, and Guillaume Vanderbrouke. 2005. "The Baby Boom and the Baby Bust: Some Macroeconomics for Population Economists." *American Economic Review* 95(1) March: 183-207.
- Hausman, Jerry, and Ephraim Leibtag. 2004. "CPI Bias from Supercenters: Does the BLS Know that Wal-Mart Exists?" NBER Working Paper 10712.
- Heckman, James. 1974. "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption over the Life Cycle." *American Economic Review* 64(1) march: 188-194.
- Heckman, James, and Thomas McCurdy. 1980. "A Life Cycle Model of Female Labor Supply." *Review of Economic Studies* 47:47-74.
- Kaufman, Phillip, James MacDonald, Steve Lutz, and David Smallwood. 1997. "Do the Poor Pay More for Food?" *USDA Agricultural Economic Report No 759*.
- Krueger, Dirk, and Jesus Fernandez-Villaverde. Forthcoming. "Consumption over the Life Cycle: Some Facts from CEX Data." *Review of Economics and Statistics*.
- McGrattan, Ellen, Richard Rogerson, and Randall Wright. 1997. "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy." *International Economic Review* 38(2): 267-290.
- McKenzie, David, and Ernesto Schargrotsky. 2004. "Buying Less, But Shopping More: Changes in Consumption Patterns During A Crisis." BREAD Working Paper 92.
- Rios-Rull, Jose-Victor. 1993. "Working in the Market, Working at Home, and the Acquisition of Skills: A General-Equilibrium Approach." *American Economic Review* 83(4): 893-907.
- Reid, Margaret. 1934. *The Economics of Household Production*. New York: J. Wiley and Sons.
- Rupert, Peter, Richard Rogerson, and Randall Wright. 1995. "Estimating Substitution Elasticities in Household Production Models." *Economic Theory* 6:179-193.
- Rupert, Peter, Richard Rogerson, and Randall Wright. 2000. "Homework in Labor Economics: Household Production and Intertemporal Substitution." *Journal of Monetary Economics* 46:557-579.
- Sorensen, Alan. 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *Journal of Political Economy* 108(4) August: 833-850.
- Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69(3) June: 213-225.

Table 1: Average Price Paid by Demographic Groups

Income Category	(1)	(2)		(3)	
	Average p	Household Size	Average p	Household Composition	Average p
<\$30,000	0.98 (0.01)	1	0.96 (0.01)	Married with Children	1.03 (0.01)
\$30,000- \$50,000	1.00 (0.01)	2	0.99 (0.01)	Unmarried Female w/ Children	0.99 (0.02)
\$50,000- \$70,000	1.03 (0.01)	3	1.01 (0.01)	Unmarried Male w/ Children	1.03 (0.04)
>\$70,000	1.03 (0.01)	4	1.04 (0.01)	Married w/o Children	1.01 (0.01)
		>4	1.06 (0.01)	Unmarried Female w/o Children	0.96 (0.01)
				Unmarried Male w/o Children	0.99 (0.01)

Notes: See text for details of construction of p . An observation is a household in a particular month. There were 2,060 households and 41,408 total observations. Robust standard errors (clustered on household) in parentheses.

Table 2: Time Use over the life cycle

A. Average Minutes per Day for All Households

	Shopping		Home Production	
	Grocery	All	Food	All
25–29	8.9	34.1	42.6	128.9
30–34	11.4	41.6	54.6	172.4
35–39	11.5	44.2	63.2	196.7
40–44	11.8	42.6	63.7	210.1
45–49	11.6	40.4	60.8	209.2
50–54	11.9	44.9	54.0	205.6
55–64	11.3	40.3	64.4	247.0
65+	14.9	50.1	75.8	270.1

B. Average Minutes per Day for Married Households

	Shopping		Home Production	
	Grocery	All	Food	All
25–29	14.6	49.9	68.8	190.1
30–34	15.0	52.7	73.1	226.7
35–39	14.1	54.2	79.4	242.3
40–44	14.9	53.1	80.8	264.7
45–49	14.3	50.3	75.4	256.9
50–54	14.0	54.0	64.9	245.1
55–64	12.4	44.9	75.8	289.6
65+	18.2	61.4	91.5	323.7

Notes: Data from American Time Use Survey 2003. In the case of shopping, “All” refers to shopping for all goods. In the case of home production, “All” refers to general household activities. Food production refers to food preparation and clean-up. Panel A is all households. Panel B is married households. In both panels, household time for married households is calculated by summing married men and women in the sample, using the age of the husband as reference. Age refers to age of household head.

Table 3: Average Price Paid as a Function of Shopping Frequency and Total Quantity

Dependent Variable:	(1)	(2)	(3)	(4)
	Eq 2.8 ln(p)	Eq 2.9 p	ln(p)	ln(p)
ln(shopping frequency)	-0.08 (0.01)		-0.05 (0.01)	-0.11 (0.02)
Shopping frequency		-0.02 (0.005)		
(Shopping frequency) ²		4x10 ⁻⁴ (3x10 ⁻⁴)		
<i>Elasticity with respect to shopping frequency:</i>	-0.08	-0.12	-0.05	-0.11
Additional Terms	ln(Q)	Q, ..., Q ⁵	ln(Q), #Prod ^a , #Varieties ^b	ln(Q)(IV) ^c , #Prod ^a , #Varieties ^b , Household Characteristics ^d
<i>Elasticity with respect to Q:</i>	0.07	0.11	0.10	0.21
N	41,408	41,408	41,408	41,408
R-squared	0.04	0.07	0.06	NA

Notes: An observation is a household in a particular month. There were 2,060 total households, restricted to households with heads at least 25 years of age. Robust standard errors clustered on household in parentheses. See text for details of specifications and definitions of p and Q . Elasticities are calculated at sample averages.

a. #Prod is defined as log of number of product categories (milk, beer, etc) purchased in month.

b. # Varieties is defined as log of number of individual UPC codes purchased in month.

c. ln(Q) is instrumented with household income.

d. Household characteristics are dummies for household size (eight categories) and household composition (eight categories).

Table 4: Estimated Elasticity of Home Production Function

	(1)	(2)	(3)	(4)	(5)
Dependent Variable:	$\ln(h/Q)$	$\ln(h/Q)$	$\ln(h/Q)$	$\ln(h/Q)$	$\ln(h)$
σ (elasticity of substitution in home production)	1.2 (0.1)	1.3 (0.1)	2.5 (0.2)	2.7 (0.2)	2.5 (0.5)
$\ln(Q)$					1.0 (0.3)
Source of Variation (Number of Groups)	Age*Sex* Marriage* Education (92)	Age*Sex* Marriage* Education (92)	Age (8)	Age (8)	Age (8)
Additional Controls	Sex and Marriage Dummies, Constant	Sex and Marriage Dummies, Constant	Constant	Constant	Constant
N	41,408	41,408	41,408	41,408	41,408
R-squared (Between)	0.82	0.83	0.96	0.98	0.92
Measure of h	Food Prep	All Housework	Food Prep	All Housework	Food Prep

Notes: Between-effects regression using Homescan demographic categories. The first two columns use 92 age*sex*marriage*education categories. The remaining columns use eight age categories. Age, sex, and education refer to household head. See text for definition of categories. Time spent on home production is from ATUS 2003. Q is index of quantity purchased defined in text. The elasticity of substitution between time and goods in home production is the (negative of) the coefficient on the MRT between time and goods in shopping. See text for details. Bootstrapped standard errors using 500 repetitions and clustered on households, where each repetition includes estimation of the right-hand regressors, are reported in parentheses.

Table 5: Model Calibration

		Benchmark	Robustness
Preferences:			
Discount Rate	β	0.98	
Inverse of EIS Consumption	γ	5	1.5 ^c
Inverse of EIS Leisure	ν	3	1.5 ^c
Relative Preference for Leisure	θ	6×10^{-5}	$8.6 \times 10^{-8a}, 6 \times 10^{-5b}, 0.7^c$
Adult Equivalences:			
Children <5 years	α_1	0.1	0.5 ^b
Children 6-12 years	α_2	0.5	0.5 ^b
Children >13	α_3	1.0	0.5 ^b
Returns to Scale	η	0.9	1.0 ^b
Home Production Technology:			
Elasticity of Substitution in Home Production	σ	2.5	1.1 ^a
Scale Parameter for Home Production	ψ_h	0.9	
Shopping Technology:			
Elasticity of Price wrt Time	ξ_s	-0.11	-0.05 ^a
Elasticity of Price wrt Q	ξ_Q	0.21	0.1 ^a
Budget Share Parameters			
Homescan Food in total Expenditure	ϕ_Q	0.04	
Food Shopping in total Shopping	ϕ_s	0.3	
Food Production in total Housework	ϕ_H	0.3	
Interest Rate	r	0.02	

Notes: Values of parameters used in model of Section 5. Superscripts *a*, *b*, and *c* refer to the three separate sensitivity analyses discussed in Section 5.5 and reported in Figure 11.

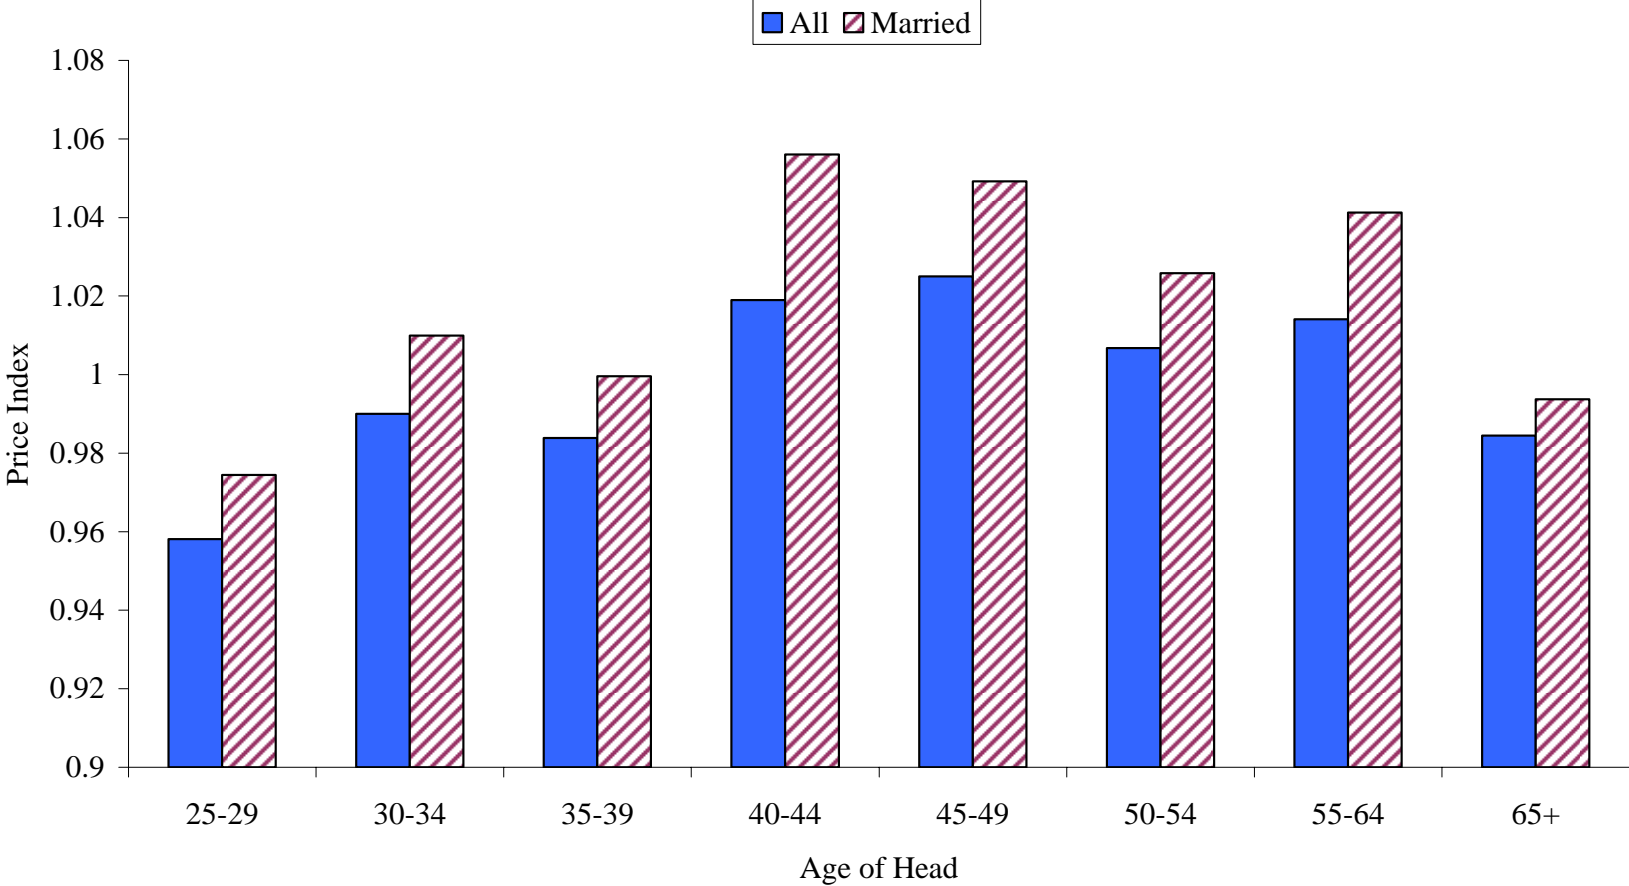
Table A1: Summary Demographics

	Homescan (1994)	ATUS (2003)	PSID (1994)
Number of Households ^a	1,607	16,816	6,508
Percent Married	55%	66%	55%
Percent with Children	35%	41%	38%
Percent Employed			
Male	80%	83%	78%
Female	68%	74%	63%
Percent High School or less	31%	44%	52%
Percent Age 25–39	33%	34%	36%
Percent Age 40–54	38%	37%	33%
Percent Age 55 and older	29%	29%	28%
Percent White	92%	77%	84%

Notes: Summary demographics for Homescan and ATUS samples, as well as a reference wave (1994) of the PSID. For this table, Homescan data restricted to 1994 for direct comparison with the 1994 wave of PSID. Homescan sample restricted to households in which the head was at least 25 and the average age of the primary shopper was between 24 and 75. ATUS and PSID samples restricted to households in which the head was between 25 and 75. For married-couple households, head refers to the male (to accord with PSID methodology). All demographic information except employment refers to the household head.

a: Not all demographic information is available for the full sample of households.

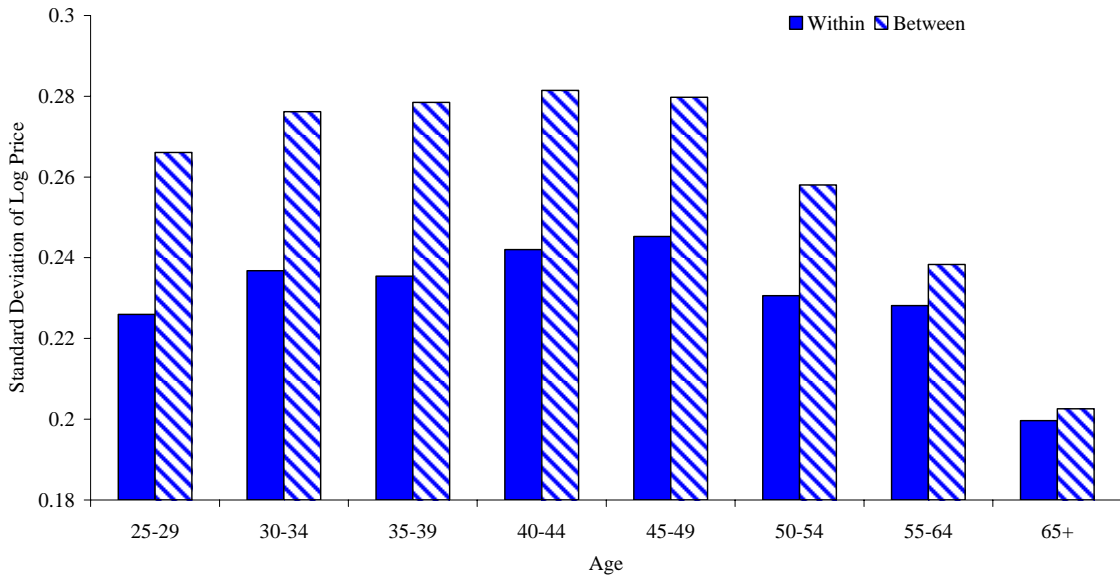
Figure 1: Price Paid, by Age of Household Head



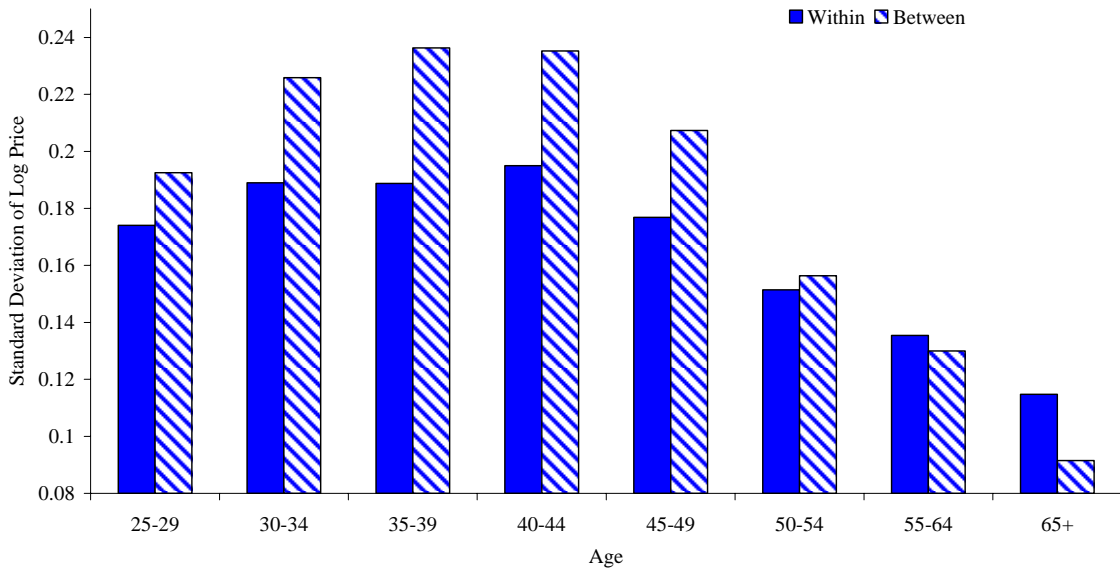
Note: Data from AC Nielsen Homescan. See text for details on construction of Price Index.

Figure 2: Price Dispersion

Panel A: All Goods

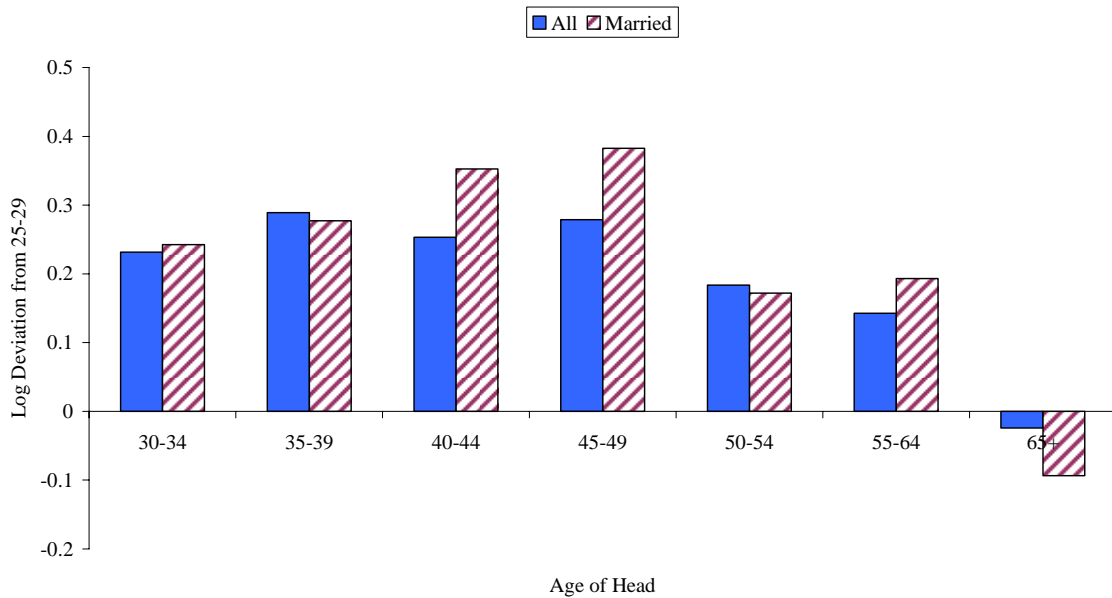


Panel B: Milk



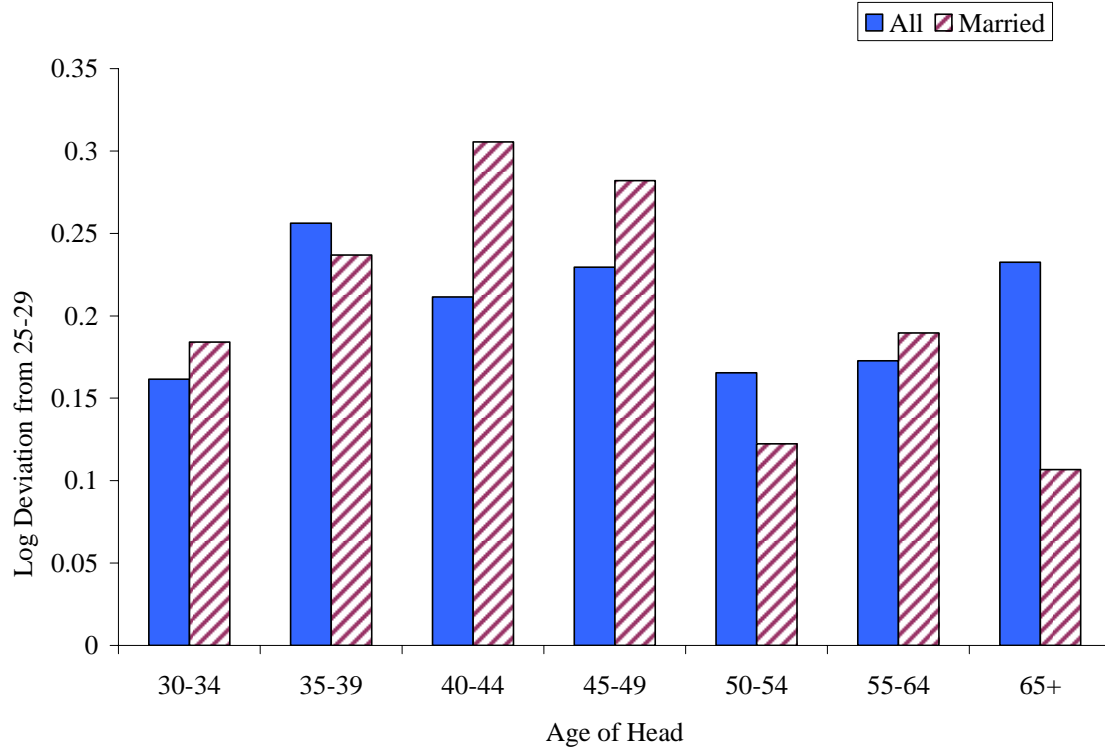
Note: Average standard deviation of log price by age group. Panel A uses all goods. Panel B uses only milk. “Within” is constructed by calculating the standard deviation of log price for each UPC code and household across shopping trips in each year. We then average over goods, years, and households within each age range. “Between” is constructed from the standard deviation of log price paid for each UPC code and month across households in an age group. We then average across items. Averages across goods and households are weighted by number of shopping trips.

Figure 3: Implied Empirical Opportunity Cost of Time



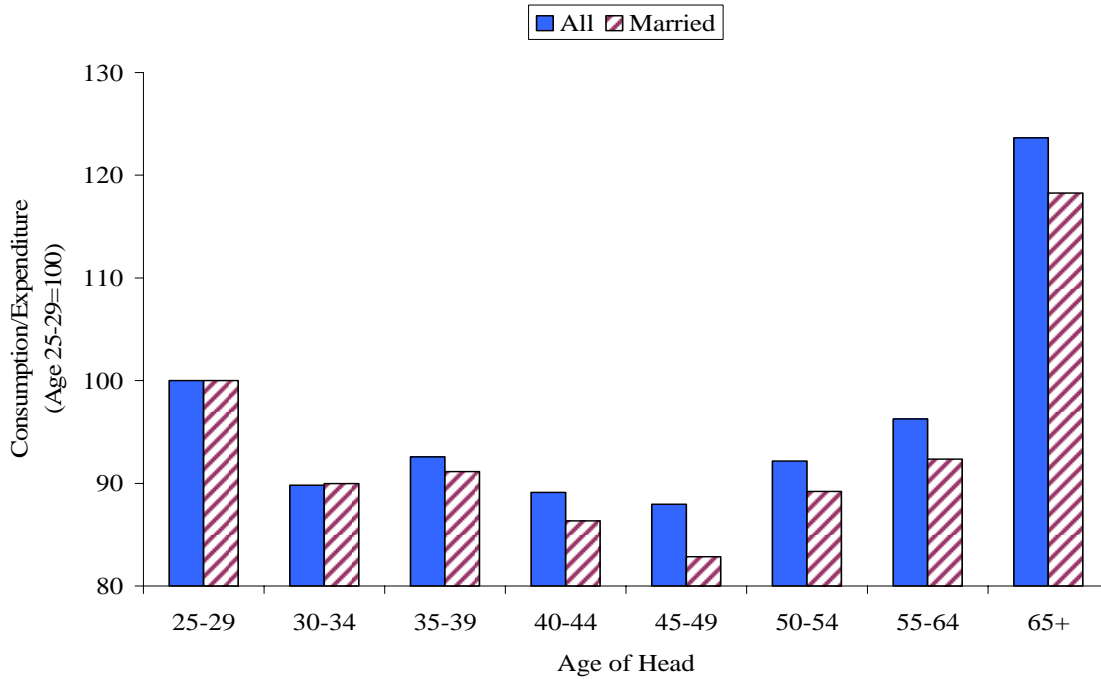
Note: The opportunity cost of time is calculated as the derivative of price with respect to shopping intensity times quantity purchased. See Section 3 for details. Figure depicts log deviations from households whose head is aged 25–29.

Figure 4: Implied Household Consumption



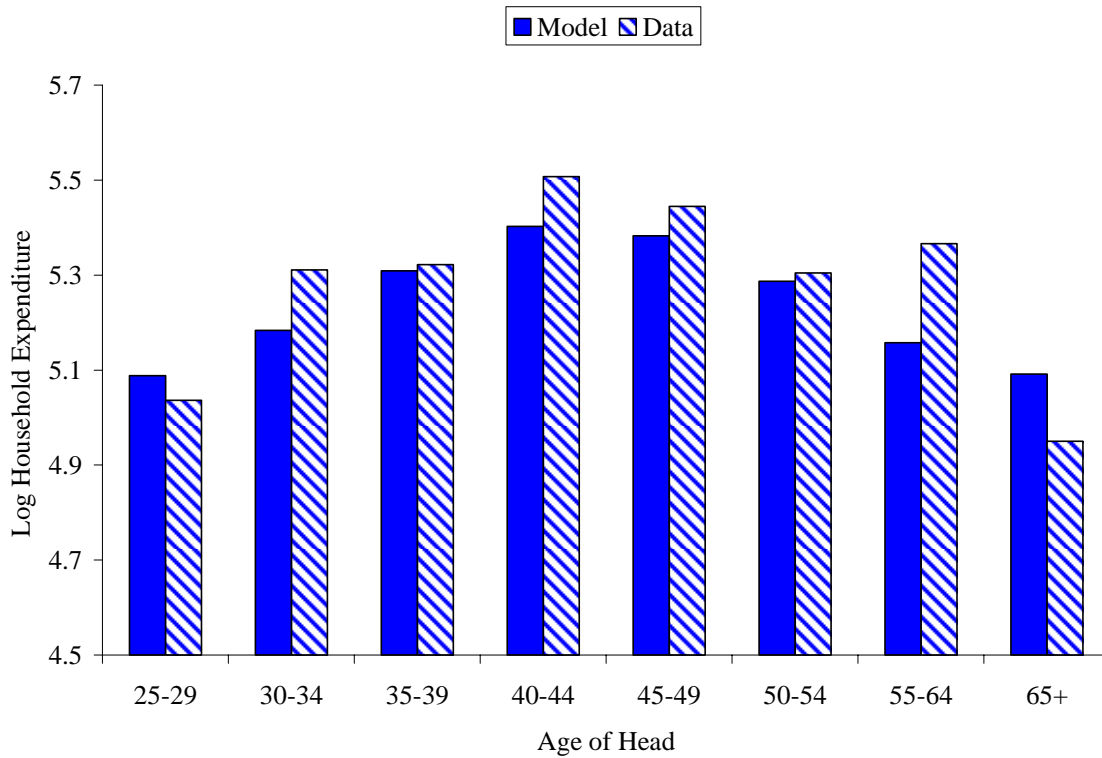
Note: Consumption calculated using parameterized home production function discussed and estimated in Section 3. Inputs of time and goods from ATUS and Homescan data sets, respectively. Figure depicts log deviations from households whose head is aged 25–29.

Figure 5: Consumption/Expenditure



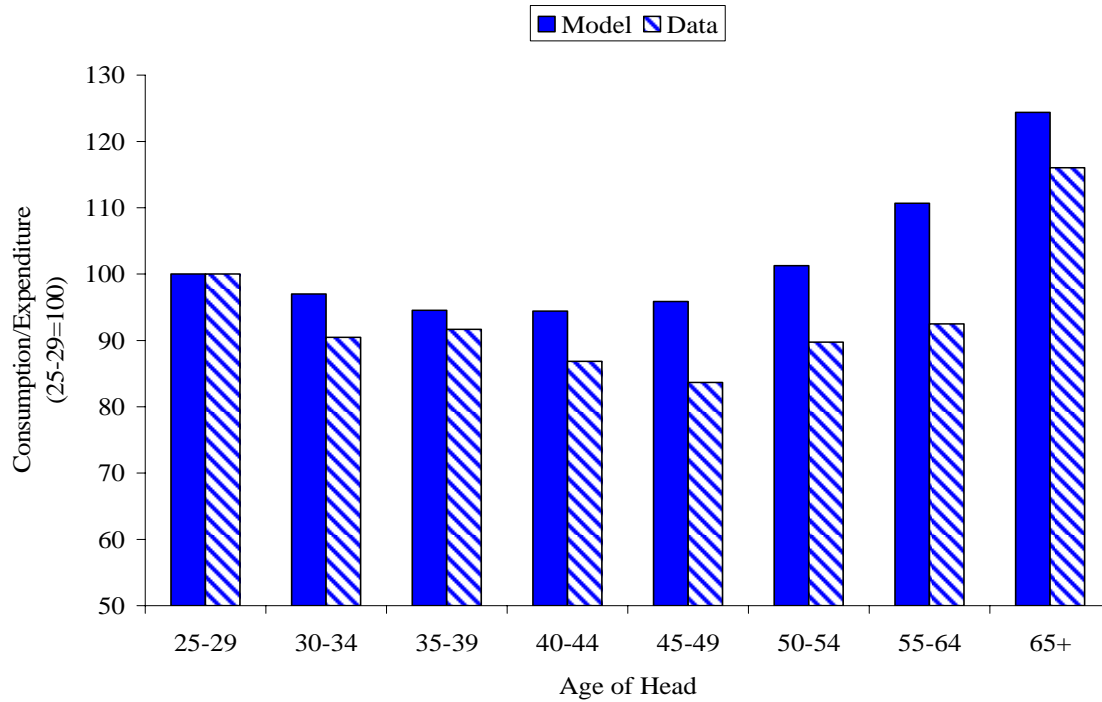
Note: Consumption calculated using parameterized home production function discussed and estimated in Section 3. Inputs of time and goods from ATUS and Homescan data sets, respectively. Expenditure from Homescan. Both series normalized to 100 for households whose head is aged 25–29.

Figure 6: Predicted Expenditure over the Life Cycle



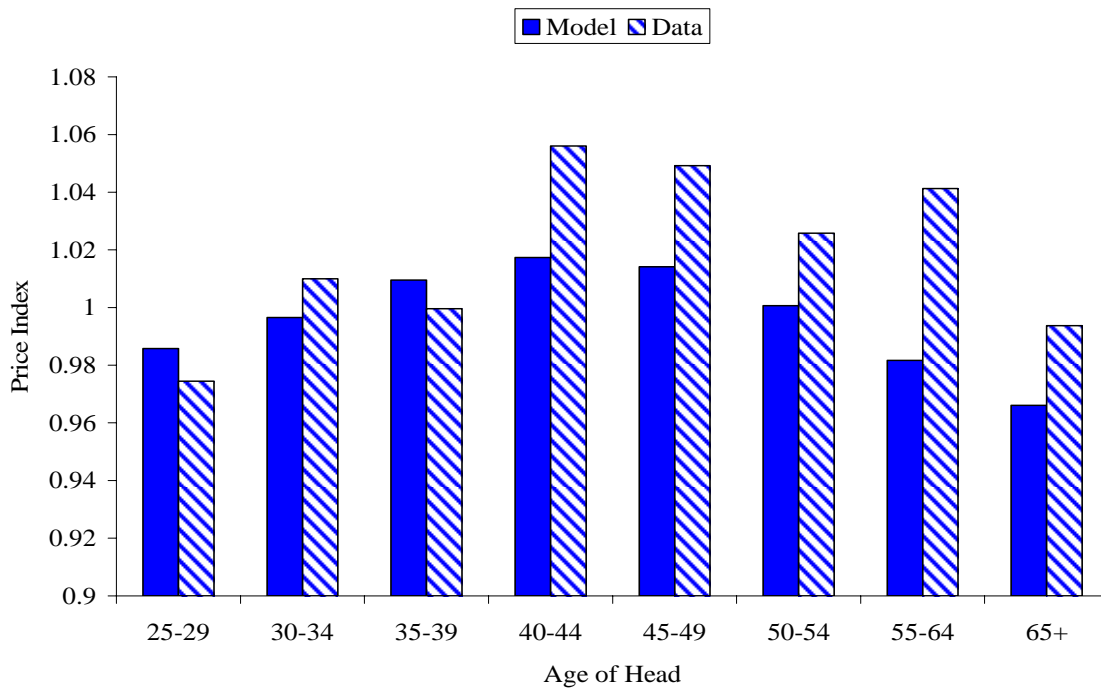
Note: Model's predictions. See text for details. Data refer to married households in Homescan.

Figure 7: Consumption/Expenditure



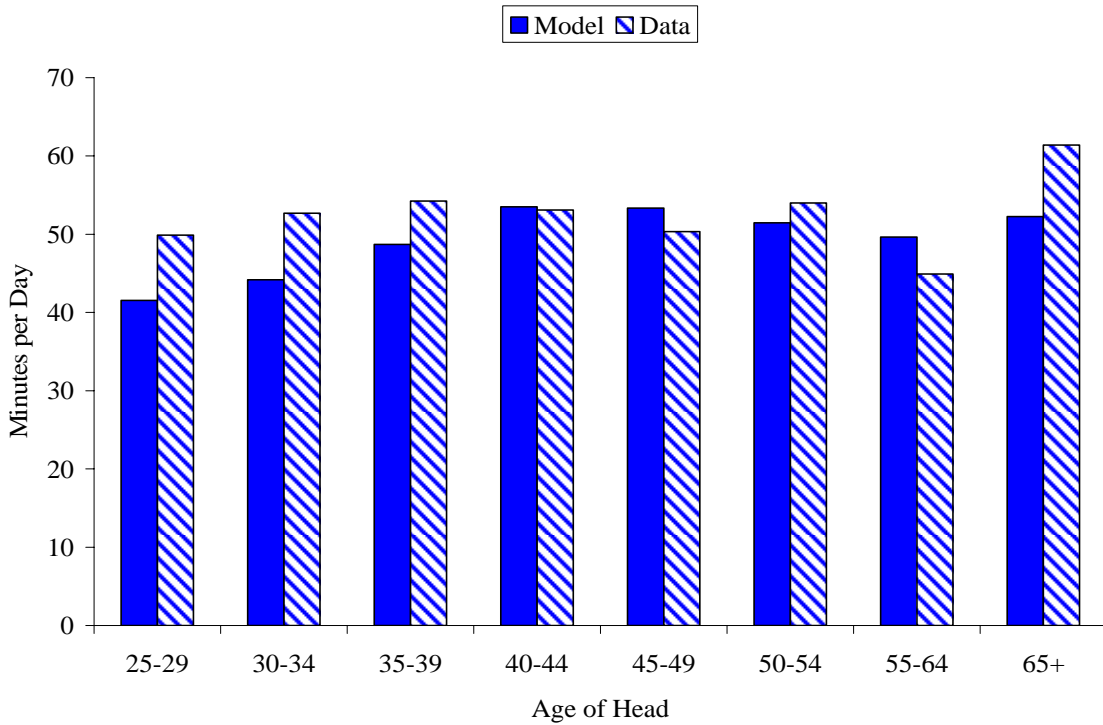
Note: Ratio of household consumption to expenditure. Age 25–29 normalized to 100 for both series. Consumption constructed using market goods and time spent in home production as inputs into production function. Data refer to married households in the AC Nielsen Homescan database with time use merged in from ATUS. See text for details.

Figure 8: Predicted Price Paid over the Life Cycle



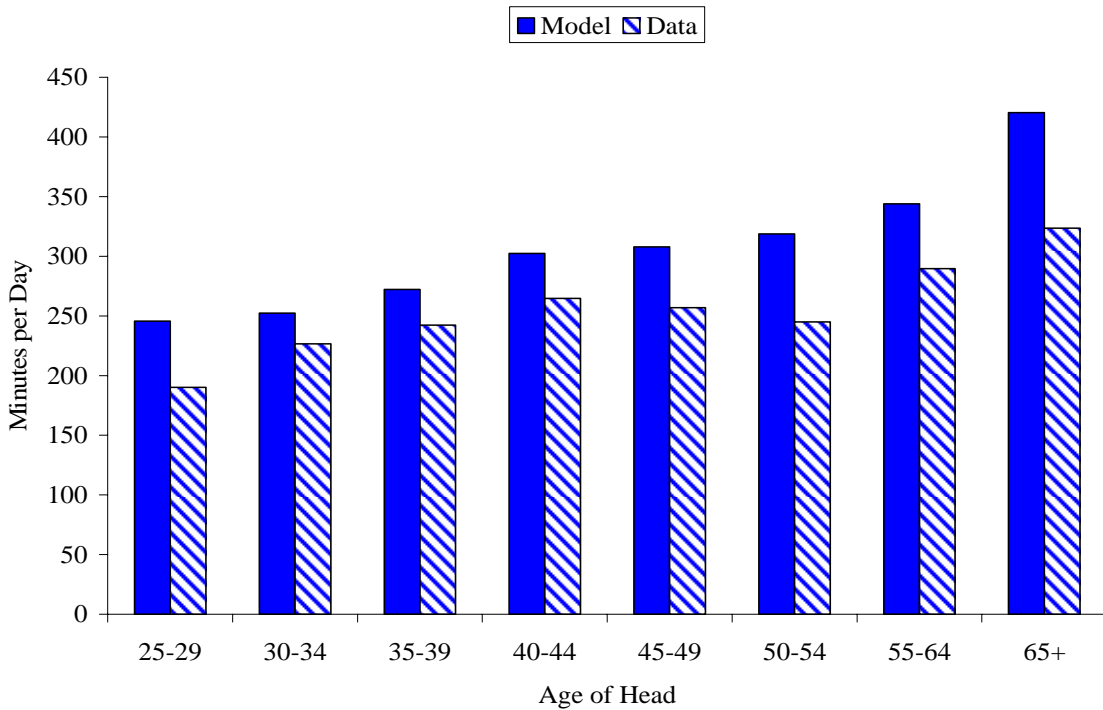
Note: See text for details. Data refer to married households in the AC Nielsen Homescan database. See text for details regarding price index.

Figure 9: Shopping



Note: Time spent shopping for all goods. Model's predictions refer to food shopping scaled up by $1/\phi_s$. Data refer to shopping for all goods reported by married households in the ATUS 2003 database.

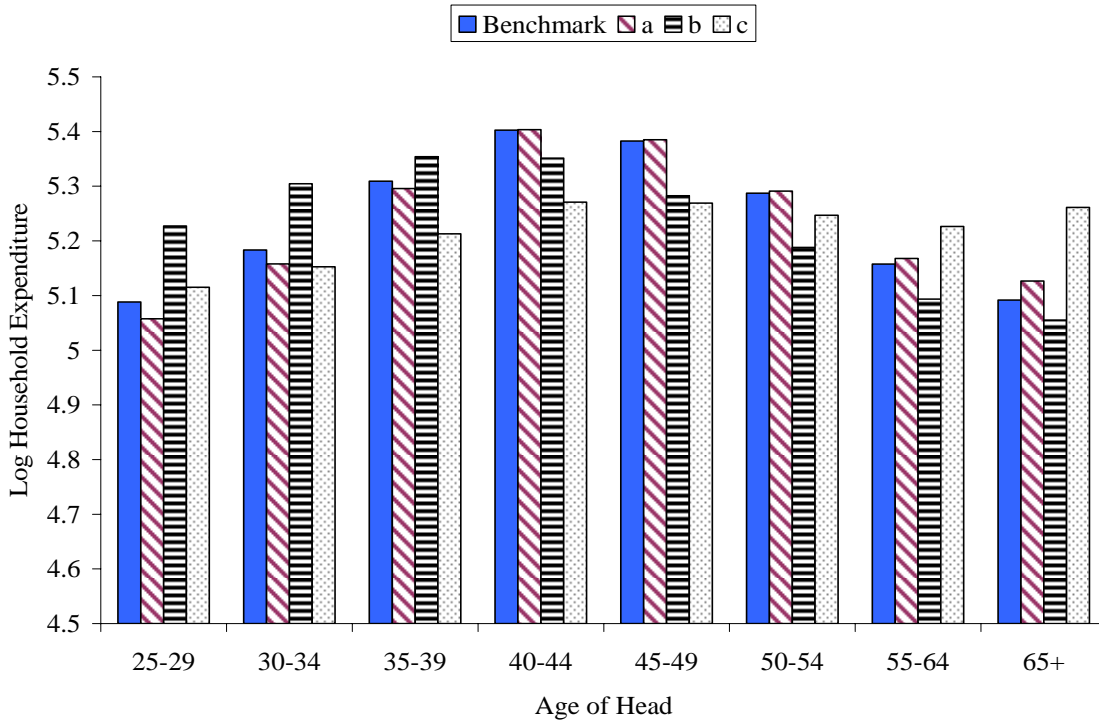
Figure 10: Home Production



Note: Time spent in home production. Model's predictions refer to food preparation scaled up by $1/\phi_H$. Data refer to all housework and home production reported by married households in the ATUS 2003 database.

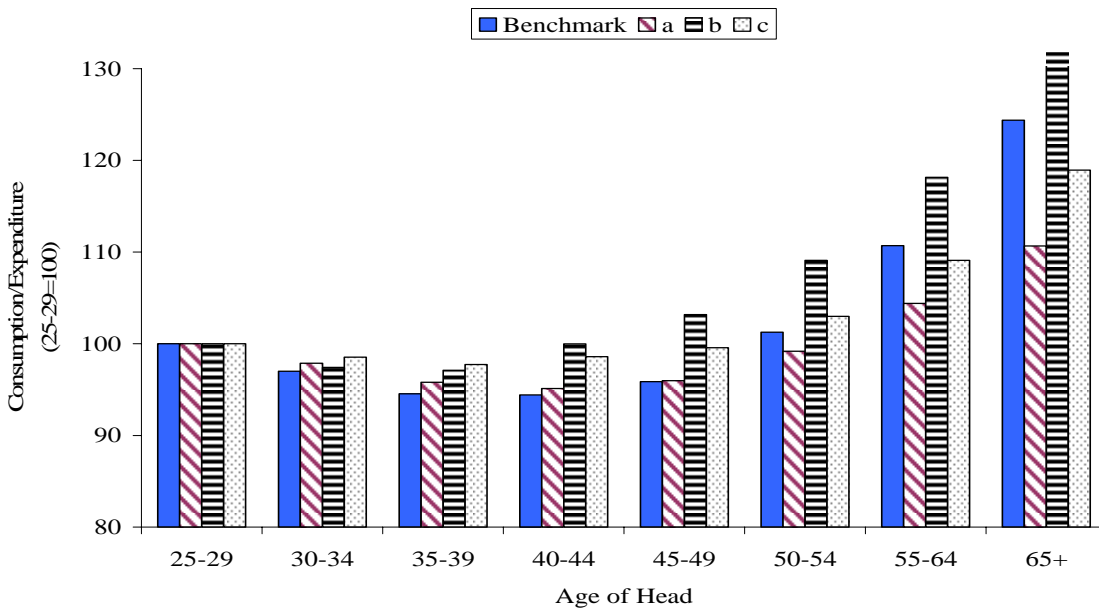
Figure 11: Robustness

Panel A: Expenditures

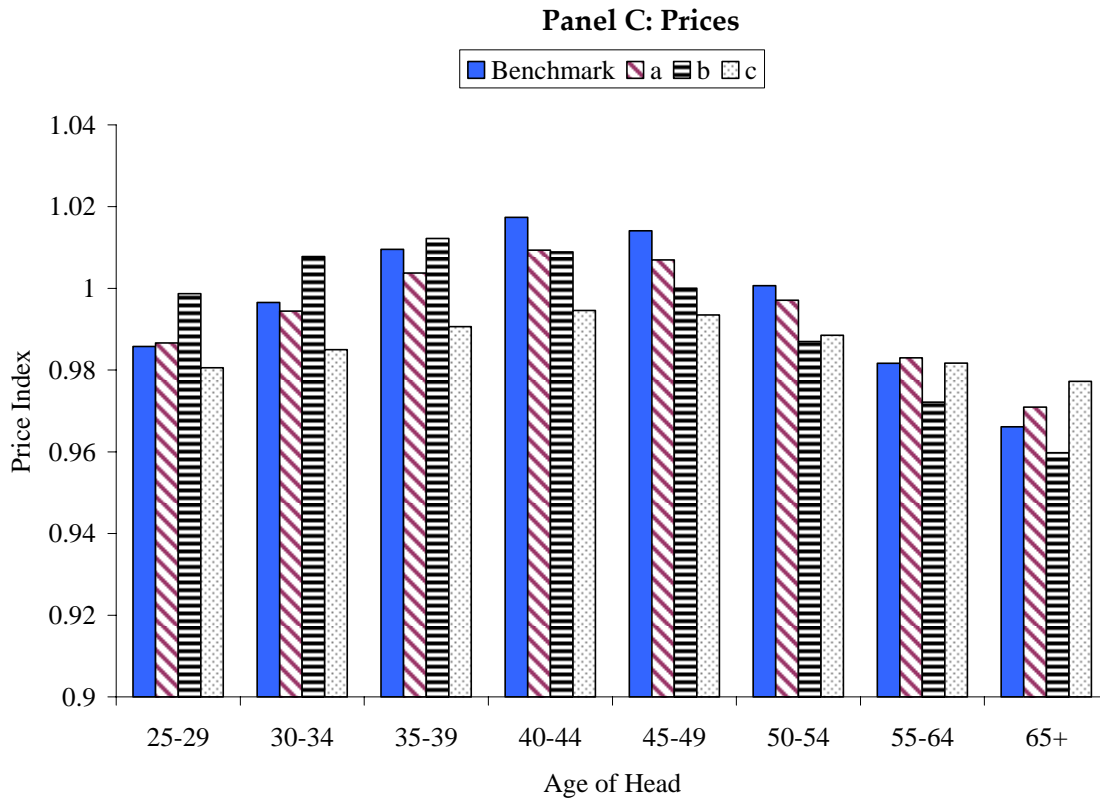


Note: Benchmark is same as “Model” in Figure 13. Robustness a,b,c, refer to predictions from alternative parameterizations of model reported in Column 3 of Table 3.

Panel B: Consumption/Expenditure

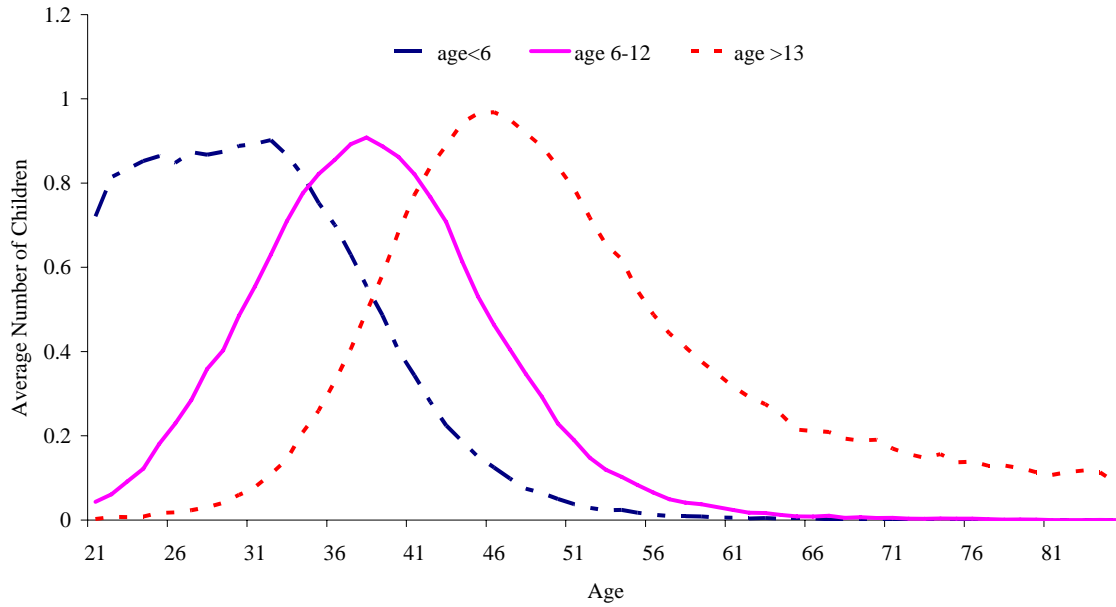


Note: Benchmark is same as “Model” in Figure 14. Robustness a,b,c, refer to predictions from alternative parameterizations of model reported in Column 3 of Table 3.



Note: Benchmark is same as “Model” in Figure 15. Robustness a,b,c, refer to predictions from alternative parameterizations of model reported in Column 3 of Table 3.

Figure A1: Number of Children over the Life Cycle



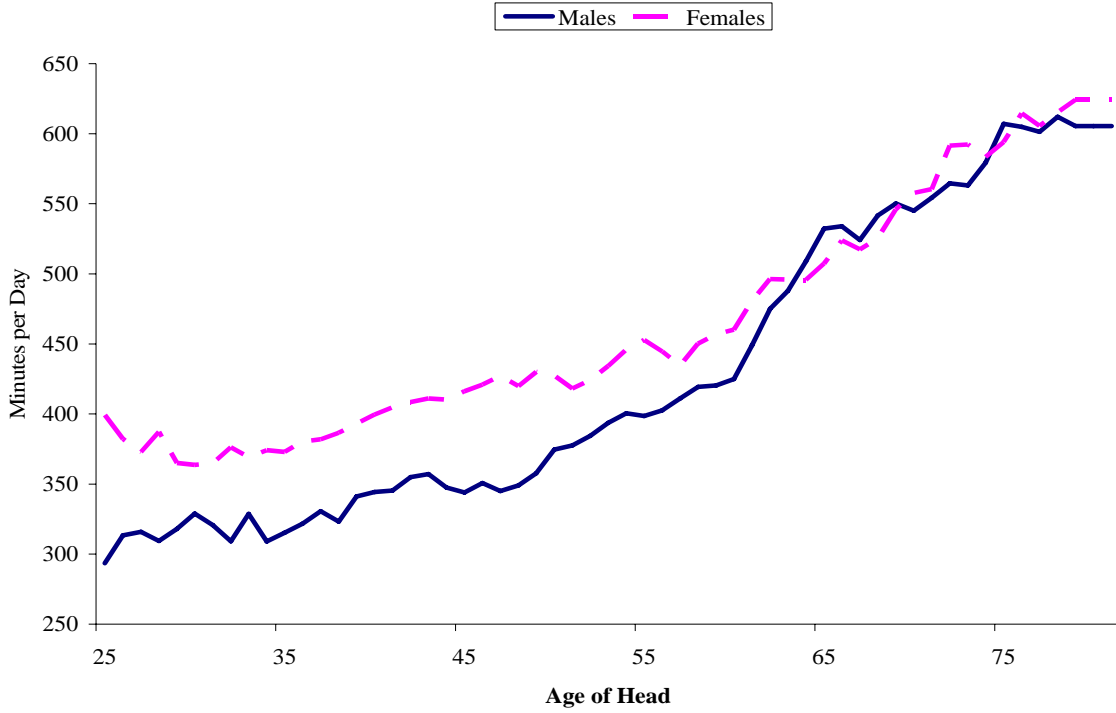
Note: Source: 2000 Census. Series represents 3-year moving average of number of children per household. Age refers to age of household head.

Figure A2: Lifecycle Wage Profile



Note: PSID wage series for men and women in households with head aged 25–74. Wages are those reported for 1993–1995 (asked of waves 1994–1996). Series expressed in log deviation from households with heads aged 25–29. For direct comparison with Homescan, wages are expressed in contemporaneous dollars. Wages are conditional on working.

Figure A3: Time Allocation over the Life Cycle



Source: ATUS. Series depict minutes per day allocated in total to home production, shopping, and leisure for households in ATUS.