

Lifecycle Prices and Production

Mark Aguiar
Federal Reserve Bank of Boston

Erik Hurst
University of Chicago and NBER

July, 2005

Preliminary– Comments Welcome*

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 age, 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 5 to 10 percent. From this elasticity and observed shopping intensity, we impute the opportunity cost of time for the shopper 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 two. 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. It also highlights the danger of interpreting lifecycle expenditure without acknowledging the changing demands on time and the available margins of substituting time for money.

* mark.aguiar@bos.frb.org, erik.hurst@gsb.uchicago.edu. We are particularly grateful to Jean-Pierre Dube for assistance with the Homescan database. We would also like to thank seminar participants at the Boston Fed, University of Chicago, Columbia University, MIT, Northwestern, and Rochester. The views are not necessarily those of the Federal Reserve Bank of Boston or the Federal Reserve System. Hurst thanks the University of Chicago Graduate School of Business for financial support.

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 can – and do – use to substitute time for money. In this paper, we use a novel dataset to document that households who shop more intensively pay lower prices for *identical* goods. We merge this data with a new dataset on time use to estimate parameters of the shopping and home production technologies that 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.

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 optimally substitute time for expenditures 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 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 data base

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), Brown and Goolsbee (2002), and McKenzie and Schargrodsky (2004). Rupert, Rogerson and Wright (1995, 2000), McGrattan, Rogerson, and Wright (1997), Campbell and Ludvigson (2001), and Aguiar and Hurst (forthcoming) provide recent empirical evidence on the importance of home production.

records the actual price paid by the household at the level of the UPC bar code. The data 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 also includes detailed information about the shopping trip, we can infer the household's shopping intensity.

We find that the price paid for a particular bar code is an increasing function of income. Specifically, households with annual income over \$70,000 on average pay 5 percent more for an identical good (UPC code) 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 to households with no children, all else equal.

One of our most striking results is that prices paid by households are humped-shaped over the lifecycle. Households in their early 40s pay, on average, between 6 percent and 8 percent more for identical goods than either households in their early 20s or households in their late 60s. Households in their 40s face the highest market opportunity cost of time (highest wages) as well as facing the highest non-market demands on time (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 age 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, we are able to estimate a "shopping function" that maps time and quantity purchased into price. We find that holding constant the quantity of goods purchased, households who shop more

frequently pay lower prices. Specifically, all else equal, households who increase their shopping frequency by 100 percent will pay prices that are 5 to 10 percent less on average. Likewise, holding shopping frequency constant, households who 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 lifecycle, 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 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 married households with only one worker. We estimate an elasticity of substitution between time and goods in home production of close to two.

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 lifecycle in a manner distinct from household expenditures. Specifically, the ratio of implied consumption to expenditures declines as households enter middle age and then rises rapidly through retirement. The lifecycle profile of this ratio reflects the changing cost of time as

households age and highlights the danger of inferring the lifecycle profile of consumption directly from expenditures. Specifically, expenditures are a poor proxy for actual consumption when the value of the household's time is changing.

Finally, we incorporate the fact that households can shop for bargains and undertake home production into an otherwise standard model of lifecycle consumption. We find that our simple model augmented with home production and shopping can quantitatively match the data along a variety of dimensions. In particular, our model generates a humped shaped profile in household expenditure over the lifecycle of similar magnitude as 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.

The lifecycle profile of expenditures has been well documented.³ These studies typically find that expenditure is “hump” shaped even controlling for household size. Heckman (1974) interprets the hump shape in expenditure over the lifecycle as being evidence that household utility is nonseparable in consumption and leisure. When household 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 that are associated with changing household structure. Our data, and the accompanying model, provide a microfounded story of how the ability to home produce implies a non-separability between *expenditure* and leisure even when utility is separable over *consumption* and leisure.

The remainder of the paper is structured as follows. In the next section, we describe our price data from the ACNielsen Homescan and explore how prices for identical goods vary with demographics and shopping intensity. In Section 3, we introduce our time use data and estimate a

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).

home production function. In Section 4, we discuss how consumption is distinct from expenditure over the lifecycle. In Section 5, we introduce a quantitative lifecycle model and discuss how well the theory explains the documented empirical behavior. Section 6 concludes.

2. Prices Paid Over the Lifecycle

2.1. Data

Our price data is from ACNielsen Homescan Panel. The Homescan data is designed to capture all consumer grocery package 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 1994 PSID (see Table A1).

Respondents in the the Homescan survey remain in the survey for upwards of twenty seven months. The survey is implemented at the household level and contains detailed demographics, which are updated annually. Specifically, we know the households age, sex, race, family composition, education, employment status, and income. The latter two categories 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 households shopping area (including grocery stores, convenience stores, specialty stores, super centers, and price clubs). If the households shops at a store outside their shopping area, the household can manually enter in the store information. ACNielsen maintains a database of

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

⁵ All packaged goods have a unique UPC code printed on their packaging. The codes are very specific. An liter bottle of Pepsi, a six pack of Pepsi cans, and a twelve pack of Pepsi cans all have distinct UPC codes.

current prices for all stores within the metropolitan area. Given the store and date information, ACNielsen can link each product scanned by the household to the actual price it was selling for 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, we have 2,100 separate households and over 950,000 transactions. For our analysis, we focus on households where 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 during 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 middle age households (richer cohorts) and older households (poorer cohorts). Likewise, it could cause us to understate the increase in expenditure between young and middle aged household. However, this should not be as important an issue for the normalized variables we focus on, such as the ratio of consumption to expenditure. In Appendix A we discuss and quantify a number of additional potential shortcomings of the data set. These include failure to scan in all purchases, attrition over time, and store and chain fixed-effects.

2.2. Prices Paid and the Opportunity Cost of Time

⁶ Households may pay lower than the stated store price if they use coupons or avail themselves to in store discounts. This information is manually entered by the households. Given that this information is likely fraught with large amounts of measurement error, we do not use it when computing our price indices. We have redone 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 coupon use by older households (see Cronovich et al (1997)).

⁷ The Homescan database records up to three ages for each household: male head, female head, and primary shopper. The former two are categorical variables while the latter takes on all 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 (i.e., 25-29), the majority of married households report the same age category for both heads. As a result, it makes little difference to our analysis whether we use the shoppers age, the male head’s age or the female head’s age.

Standard economics suggests that, all else 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 superstores which may involve longer commutes and check-out lines rather than shop at convenience stores, or clip coupons and mail in rebates.⁸

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 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 expenditures 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 Walmart offer prices between 5 percent and 55 percent less than 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 divided by 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 one:

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

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 distributions 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 one as reflecting a household that pays on average higher prices, and vice versa for an index less than one. It is important that the price premium is not reflecting higher quality. Given our index, this is not the case. The price differentials are for the *identical* goods as measured by UPC codes.⁹

Using our price index, we can revisit whether prices paid for the same goods vary across households with different costs of time. One measure for the opportunity cost of time is the market wage. In the Homescan data, we do not have wages; we only have categorical measures of household income. Using this data, we aggregate up to four income categories: income < \$30,000, income between \$30,000 and \$50,000, income between \$50,000 and \$70,000, and income >\$70,000. In Table 1 Column 1, we report mean price index for households within the four income categories. The results are striking. Households who earn less than \$30,000 a year, on average, pay 5 percent lower prices than households who earn over \$70,000 (p-value of difference < 0.01).¹⁰ Households who earn between \$30,000 and \$50,000 pay, on average, 3 percent lower prices than households who earn over \$70,000 (p-value of difference < 0.01). The difference in prices paid between households who earn less than \$30,000 and households who 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 for 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.¹¹

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

⁹ 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 in practice is not substantial – they share a correlation coefficient of 0.8.

¹⁰ 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 one. A similar caveat holds throughout.

¹¹ There is mixed evidence that prices are higher in poor neighborhoods (see 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 (<\$5,000) do pay slightly higher prices on average than those closer to \$30,000. However, the small number of extreme low income households makes it difficult to precisely characterize this potential non-monotonicity.

higher prices than households with smaller families. Specifically, households with only one household member pay 10 percent less for an item compared to households with family size greater than or equal to 5 (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 both income categories and family size categories, both sets of regressors enter significantly (results not reported).

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 that high income and large household size raises the cost of shopping. The model of Section 5 will allow us to see how a household optimally weighs these considerations. Empirically, Figures 2 and 3 indicate that the costs of shopping dominate to the extent that richer and larger households pay higher prices.

We have also explored whether married families in which both adults work at least 30 hours 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 that market labor is endogenous. For example, households which face greater time demands within the home may opt to have only one spouse work in the market while those which do not face such heavy demands have both spouses supply 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 which reduces a spouse's willingness to supply market labor that also reduces the intensity of shopping. Note that this implication of endogenous labor supply 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. For married 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 households in which at least one spouse works full time. For all households, the difference is 1 percent, but not statistically significant (p-value=0.41).

Given that both the arrival of children and household wages have a lifecycle component, we would expect our price index to vary with age. Using the 2000 census, we find that children in a married household peak when the head is in his or her early 40s (see Figure A1). As seen in Figure A2, wages of both males and females, conditional on working, peak around age 45-50. The wage data come from the 1993-1995 cross sections of the PSID. 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 50s and 60s. Nevertheless, both the profiles of children and market wage suggest 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 households. Consistent with our premise, households in their middle 40s pay the highest prices. Specifically, unconditional on marital status households aged 45-49 pay 6 percent higher prices than households aged 25-29 and 4 percent higher prices than households aged 65+. Conditional on marriage, households aged 40-44 pay 8 percent higher prices than households aged 25-29 and 6 percent more than those older than 65.¹²

One concern is that households may not be paying lower prices only 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 on sale. The second consumer will pay a lower price on average for the

¹² We redid the analysis in figures 1-3 and 6 using only the price index for milk (as opposed to the entire basket of purchases). Milk was the most common good purchased in the data set. Using only milk, the same conclusions could be drawn. Specifically, middle aged, rich, and large households paid the highest price for milk.

same Pepsi product. To control for this, we construct two additional measures of goods purchased. The first is the number of “goods” a household purchases per month, where good is a broad product category such as milk, beer, orange juice, etc. The second is the number of individual UPC codes, 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.¹³ Conditional on the number of goods, the number of varieties captures the propensity for a household to substitute brands or sizes. As documented below, for a given shopping frequency, more goods implies higher prices (due to dilution of shopping time) and more varieties condition on goods implies lower prices (due to the propensity to switch brands or items). All the patterns documented in Figures 1-3 and 6 are robust to the addition of these controls.

Finally, in Figure 2 we plot price *dispersion* over the lifecycle. 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 8 age ranges. For each UPC code 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 deviation 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 lifecycle, Figure 2 Panel B breaks out milk (a single category which almost every household purchases) and

¹³ We also replaced the number of UPC codes 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, but not six-packs vs. twelve-packs of Coke. The results were similar, but typically with a larger standard error on the variety coefficient.

performs the same analysis on UPC codes 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 a third to one half between middle age and retirement. This pattern is easily interpreted in a search theoretic framework and 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 retiree should therefore have a lower mean and be compressed relative to middle aged shoppers.

The patterns of mean price and price dispersion documented above suggest 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 Lifecycle

Corresponding to the premise that the opportunity cost of time varies over the lifecycle, whether due to the wage profile or 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 the price dispersion, which makes shopping more valuable in middle age when families are largest.

To examine time spent in shopping over the lifecycle, 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 of the Current Population Survey (CPS). Roughly 1,800 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 aggregates the survey responses into time use categories.¹⁴

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 all individuals between the age of 25 and 75. In Figure 3, we plot the time spent shopping for all households (Panel A) and married households (Panel B) over the lifecycle.¹⁵

Peak grocery and total shopping times occurs for households in their early 40s and for households older than 65. Households in their mid 40s 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 more intensively for a given basket of goods. Young households shop relatively little because they buy relatively few goods and have work and education demands on their time. Notice, the ratio of grocery shopping to total shopping is fairly constant at 25 percent over the lifecycle.

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:

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

¹⁴ See <http://www.bls.gov/tus> for a detailed description of the ATUS survey methodology and coding system.

¹⁵ Unfortunately, the BLS did not have each spouse within the same household fill out a time diary. We construct synthetic married households by summing over married men and women. 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.

where p is our price index, 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 who shop more will reduce their price. The returns to shopping diminish as shopping increases. Likewise, holding shopping time constant, households who purchase more goods pay higher prices.

The Homescan data 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 shopping intensity. We discuss below 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 goods 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:

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

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 price is log linear in shopping frequency and quantity. Specification (2.9) assumes price is a second order polynomial in shopping time and a fifth order polynomial in quantity.¹⁶ Columns 1 and 2 of Table 2 reports the estimates of (2.8) and (2.9), respectively. We estimate α_1 to be -0.08 (standard error = 0.01). A similar elasticity of -0.12 is obtained from (2.9), evaluated at the sample average. Moreover, the positive coefficient on the second-order term in shopping frequency from (2.9) indicates diminishing returns to search. We

¹⁶ 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.

also find that the quantity purchased has a statistically significant impact on price in both specifications, with an elasticity of 0.07 and 0.11, respectively. That is, for a given shopping frequency, the more purchases a shopper makes the higher the price of the average good.

In columns 3 and 4 of Table 2, 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 goods purchased and the number of varieties purchased, as defined in Section 2.2. While not reported, we find the coefficient on the number of goods is significantly positive (the “dilution” of shopping time effect), while the coefficient on varieties (given the number of goods) 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 may first appear. First, we are not tracking purchases by a household as the price varies over time. Rather, we are looking across households in a particular period who 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 code relative to what the average person pays. The fact that you 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. The results, reported in the final column of Table 2, suggest that the elasticity with respect to shopping frequency is unaffected and the elasticity with respect to quantity 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 lifecycle. In fact, frequency and time per trip are negatively correlated over the lifecycle. 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. We have merged in the time diaries' average time per trip for each age group as an additional regressor and found 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.

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 their food bill by both shopping more intensively and by purchasing raw ingredients and making their 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 rather than to home production.

Time spent on home production varies systematically over the lifecycle. Using the ATUS data, we define two measures of home production. The first is total time spent on food production (which includes preparing meals and meal clean up). The second is total home production (which includes food production, plus indoor cleaning and chores, clothes care, outdoor maintenance, and lawn care). As seen in Figure 3, time spent in both food production and total home production over the lifecycle mimics that of shopping time. In particular, home production time peaks for households in their early 40s and then again for households who are older than 65. As with shopping time, households in their early 40s have the greatest home production needs (i.e., largest family sizes) and households older than 1965 have the lowest

opportunity cost of time. Moreover, over the lifecycle, 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 in a lifecycle model where 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 is adequately captured by our data set. That is, different goods and uses of time enter separably in *production*.¹⁷

Letting μ_C denote the multiplier on the constraint, the first order conditions are

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

¹⁷ For the elasticity of substitution of the home production function, we need only assume weak separability. In particular, 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.

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 $-\frac{\partial p}{\partial s}Q$ each household in our Homescan data. In Figure 4, we plot the lifecycle profile of μ by averaging $-\frac{\partial p}{\partial s}Q$ over all households within a given age range and then expressing the series as differences from the age 25-29 group. We can see that the opportunity cost of time for the shopper is humped shape over the lifecycle. It is also evident that the hump early in the lifecycle is not as severe as that of wages for either males or females (Figure A2). This should not be surprising given that not all shoppers are on the margin of labor supply at the household head's wage, either because they face a different wage or that adjusting labor supply at the margin is not feasible. 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 dataset in calculated 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 this first order condition together with our estimates of $\frac{\partial s}{\partial p}$, $\frac{\partial s}{\partial Q}$, p , Q , and h , allows us to estimate 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 where we do not observe prices (or assumed every household faced 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, assume 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 smoothly vary labor supply 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 only requires that the opportunity cost of time for the shopper equals the opportunity cost of time for the home producer, a much more plausible assumption. Moreover, it strikes us as reasonable that households can smoothly adjust 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 homogenous 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 2 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 does not contain data on time spent in home production. To get around this issue, we merge together data from Homescan and ATUS. To do this, we create demographic cells in both data sets using age, sex, marital status, and education. Specifically, we use 8 age ranges (those displayed in figures 5-9), 4 education categories (less than high school, high school, some college, and college or more), 2 marital status categories, and 2 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. Estimating equation (3.6) in this manner using information from the 92 demographic cells yields an estimate of $\sigma = 1.2$, with a standard error of 0.1. 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 3, respectively.

One concern with the estimates in Columns 1 and 2 is that some of the demographic cells have few households (the minimum observations per cell is 6). This may result in significant measurement error. To correct for this, we run a between effects regression using the 8 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 3, respectively. These cells are much larger, with a minimum observation 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, which is roughly comparable to our elasticity, for single, employed women of 1.8.¹⁸ This number is in line with our estimates (moreover, restricting our sample to single women produces an estimated elasticity of 1.5). Their 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 coincide for employed single women, a demographic for which wage is most plausibly the relevant price at the margin.

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:

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

¹⁸ 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 market work and home work are perfect substitutes in the utility over leisure. This parameterization is consistent with their estimates.

The estimate of σ in this case is 2.5 with a standard error of 0.4, an elasticity roughly the same as that found above. This specification also allows a test of whether the coefficient on $\ln(Q)$ is one (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 3.

4. Lifecycle Consumption versus Lifecycle Expenditure

With a parameterized home production function, we can compare how lifecycle expenditure (the 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 lifecycle. Going from the ratios (the MRT) to levels requires us to assume a value for returns to scale, which we take to be one. 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 one. This assumption involves only a scaling of consumption and does not play a role in the analysis (given that we only display the ratio over the lifecycle relative to young households).

The path of lifecycle consumption for the household is plotted in Figure 5. As before, we plot log deviations from households aged 25-29. Household consumption has a “twin peaks” shape. Consumption rises rapidly early in the adult lifecycle, peaking around 40, declining until late middle age and then rising through retirement. Household consumption’s peak in the late 30s is 26 percent higher than adults 25-29 (p-value <0.01), 11 percent higher than those in their early 50s (p-value 0.04), and 5 percent higher than those in their 60s (p-value 0.28). Again, given 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 lifecycle, Figure 6 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 that households

use different ratios of time and market goods in consumption over the lifecycle. 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.¹⁹ 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 6. 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 lifecycle. 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 vector of children in a household of age t . We take the arrival of children as exogenous. There is no uncertainty in the model.

5.1. Preferences

¹⁹ See Bernheim, Skinner and Weinberg (2001) and Aguiar and Hurst (forthcoming).

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

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

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

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.²⁰ 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.²¹ In this fashion, we can highlight the distinction between separability between 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 non-separably. We feel the distinction is useful to understand the microfoundations behind reduced-form non-separability. Specifically, period utility is given by

²⁰ 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.

²¹ Keep in mind 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 pertained to home production and shopping.

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

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 $N_t = \left(n_A + \sum_{\tau} \alpha_{\tau} n_t(\tau) \right)^{\eta}$. n_A is the number of adults in the household and, as noted before, $n_t(\tau)$ is the number of children aged τ . The parameter α_{τ} is the relevant weight in consumption of a child of age τ to an adult. This specification of adult equivalencies has been 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 parameter in the adult equivalency schedule, and the elasticity of inter-temporal substitution (*EIS*) in consumption (discussed below). We select the normalization that the home production function is constant returns to scale and adjust the other two parameters accordingly.

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

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

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

adequately model labor supply over the lifecycle, we would need to account for the fact that workers in their late 20s and 30s 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 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. Consider the case in which the market price of goods is constant (that is, no shopping function). As the cost of time ($v'(l)$) increases, all else 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 *CRS* home production function, this implies the ratio Q/h increases (i.e., agent will substitute goods for time). To satisfy (5.6), $u'(c)/N$ must decrease. All else equal (including family size), this implies 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 use food diaries to show that retirees eat better than non-retirees 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,

$$Q^{-\gamma} g\left(\frac{Q}{h}\right) = \lambda p \quad (5.9)$$

where $g\left(\frac{Q}{h}\right) = \left(\bar{h}\left(\frac{Q}{h}\right)^{-\rho} + 1\right)^{\frac{-\gamma+1-\rho}{\rho}}$. Note that $g' > 0$ if and only if $\gamma > 1 - \rho$. Or, in other

words, g is an increasing function of the price of time 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 $p^{-\gamma}$,

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

Therefore, as the price of time increases, expenditures increase 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. 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

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

where $\xi_Q = \frac{p_Q Q}{p}$, the elasticity of price with respect to Q (which we assume here to be a constant). Sufficient conditions for expenditures to increase with the cost of time are then $\gamma > 1/\sigma$ (as before) and $\gamma \geq 1$.²² Given that empirical estimates of γ tend to be greater than one, this additional condition will typically be satisfied.

In summary, a relatively low *EIS* for consumption and the changing cost of time directly implies a “hump” in expenditures over the lifecycle. 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 (1974)’s explanation of the lifecycle profile of expenditures as stemming from a non-separability between consumption and leisure in utility. However, in our present framework, the non-separability 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

²² To derive this, take logs of both sides of (5.11) and differentiate. Note that the incentive to shop falls as the cost of time increases. All else equal, this raises the price paid.

expenditures does not necessarily reflect a “hump” in consumption (see Figure 6) and 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 4 and discussed in detail in the appendix. We calibrate to married households and consider the empirical counterpart to age of the household to be that of the male head. The simulated lifecycle profile implied by the model is displayed in Figures 13 through 17, 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)). Figure 7 indicates the model tracks this data closely over the lifecycle. However, the model over-predicts expenditures early in the lifecycle by a few percentage points and therefore the model’s hump is slightly shallower than the data’s. While the model’s additional expenditure is fairly small, it does suggest there may be room for borrowing constraints or habit formation that suppresses expenditure early in the lifecycle. 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 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 8. In this figure we plot the ratio of household consumption to that of household expenditure for the model. For comparison, we include the consumption to expenditure ratio found in the data (from Figure 6). 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 inter-temporal 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 consumption in retirement. The fact that household consumption is rising 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 9. As with expenditure, the model's hump is somewhat shallower early in the lifecycle. The fact that young households in the model consume more than young households in the data generates, all else equal, the same pattern for prices (given that $\xi_Q > 0$). The decline in prices as households age toward retirement is again nicely reflected in the model.

Behind this pattern of prices is varying time spent shopping (Figure 10). Specifically, the model predicts that middle-aged and elderly households will shop intensively. The former is due to larger family size, the latter reflects a lower cost of time. This "twin" peak in shopping time is translated 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.

The model also captures the lifecycle profile of home production observed in the data (Figure 11). The increase during retirement is somewhat larger for home production than for shopping in the data. This feature is reflected in the model as well. That is, 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 12.

The first alternative (parameterization “*a*” in Column 3 of Table 4) 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 2.3 to 1.1. Similarly the elasticity of price with respect to shopping time and goods purchased is lowered to -0.05 and 0.10 from -0.10 and 0.20, respectively. From Figure 12, we see that the lifecycle profile of expenditure remains roughly the same as that of the benchmark. As would be expected, the largest departures occur for the ratio of consumption to expenditure and prices over the lifecycle. In particular, the relative increase in consumption during retirement is muted due to the lower elasticity of substitution in home production. This brings the model’s increase in this ratio over late middle age closer to the data (compare to Figure 8). However, the lifecycle path of prices is noticeably flatter than the benchmark, increasing the deviation from the data.

Parameterization “*b*” alters the adult equivalence scales. Specifically, all children, regardless of age, are considered 0.5 adults in consumption. Moreover, we set the returns to scale parameter to one. From Figure 12 Panel A, we see that this raises expenditure early in the lifecycle. This is a direct result of the increased relative weight on infants and toddlers, the number of which peaks in the late 20s (see Figure A1). The increased purchases early in the lifecycle lead to slightly higher prices paid, as well. Although not depicted, households do shop more intensively early in the lifecycle, but not enough to offset the price effect of the larger quantity of market goods purchased. The increased expenditures early in the lifecycle 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 lifecycle (Panel B).

The last parameterization, “ c ”, raises the inter-temporal 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 12, where expenditures in retirement are noticeably higher than under the other parameterizations. The ratio of consumption to expenditure during retirement remains high relative to middle age, but is slightly lower than under 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 does not appear to be overly sensitive to lowering the home production elasticity close to one. Adult equivalency scales that place a higher weight on teenagers than toddlers help 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 inter-temporal elasticity of substitution. That is, the decline in expenditures late in the lifecycle 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 consistent with standard economic principles.

This paper has brought some new data and insights regarding these margins of substitution. However, the data have some limitations. The scanner data consists of a subset of

grocery items. We cannot access whether similar patterns hold for other goods. The time use data does suggest that households shop and engage in home production for non-food 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 is cross-sectional in nature and therefore we must be cognizant that some of our lifecycle results may be confounded with cohort effects. Cohort effects are likely to be less of an issue for normalized variables, such as the ratio as consumption to expenditure and the dispersion of prices within a household over time.

There is a growing interest in the role of non-market 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 been primarily on lifecycle consumption, we feel 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 Dataset

This appendix discusses and quantifies a number of potential concerns related to the Homescan Dataset . One issue is coverage. The average monthly expenditure for packaged goods for 1993-1994 in our sample 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 covers a little more than half of total grocery expenditures reported in the PSID. The difference between the Homescan data and the PSID likely comes from two sources. First, the Homescan data does not include meat, fresh foods or vegetables. Moreover, as discussed below, it may be the case that households fail to scan in all grocery items in the Homescan database.

Second, there is some attrition in 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 who quickly withdraw from the survey.²³ However, we can directly observe attrition on the intensive margin. On average, a household reports 1 percent less expenditures in the first quarter of 1994 compared with the same household during first quarter of 1993, and 5 percent less 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, the omissions are random within a household). However, it may influence such items as total expenditures and frequency of shopping. For each of our analyses, we have compared estimates using only the first quarter of the sample with those obtained using the sample from the last quarter and did not uncover substantial differences.

²³ Within the dataset, 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.

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 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 0.44.²⁴ 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 constant across demographic groups. However, the initial level of under-reporting appears to be correlated with education, but not age. To assess this, we compare expenditures in Homescan with the PSID. Specifically, we create cells in the PSID by age of head (using the 8 categories), education (less than high school, high school, some college, and college or more), and year (1993, 1994, 1995). For each cell, we calculate the average expenditure on food at home reported in the PSID and merge these values into the Homescan dataset. We then construct the ratio between Homescan households and their corresponding PSID cells. This gap shows no correlation with the age of the household head (p-value of F test is 0.56). However, the gap is correlated with education. For example, reported expenditure in Homescan for households with a college education or better is on average 42 percent than that reported in the PSID. The comparable fraction for high school graduates is 55 percent (p-value of difference <0.01). This suggests that higher educated households are less likely to scan all purchases. 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, due to 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 are treating 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. We do have store and chain identifiers in the data set. We find that households do

²⁴ As discussed below, we analyze 8 age ranges: 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-64, and 65+.

shop at multiple chains within a year. Moreover, we have adjusted our price index for differences across chains in mean prices for each good and found no substantial change in the results. Specifically, for a good i purchased in month m at chain k , we calculate the average price of good i sold in that chain over the relevant quarter (we average over a quarter rather than a month to ensure that a reasonable number of purchases constitute the average). We then calculate the cost of a basket purchased within the month if each good were purchased at the relevant chain's average price. This quantity is used in place of (2.4). We found no substantial differences in the patterns described in the text using this alternative index. oduce our price index.

Appendix B: Calibration

This appendix describes the details behind the calibrated parameters reported in Table 4 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 one plus the interest rate, or 0.98. This implies agents would like to maintain constant marginal utilities of consumption and leisure over the lifecycle, all else 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, v . As our model is partial equilibrium, 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 a good providing utility (as it is in the model) rather than the complement of market labor, then we must consider how non-market time is allocated to things such 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 inter-temporal 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. With these caveats aside, our reading of the labor literature implies a plausible estimate of ν to be 3. That is, a one 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 25-55 is slightly more than reported leisure, implying a Frisch labor elasticity evaluated at the mean of roughly 0.2.²⁵ For women, reported market labor is less than leisure, implying a labor supply elasticity closer to 0.4. We choose θ , the parameter governing the relative importance of leisure, so that average hours of shopping and home producing from the model roughly line up with the data's sample mean.

The household consists of two adults plus children placed into three age groups, 0-5, 6-12, and 13-18, corresponding to $\tau=1,2,3$. The number of children in each age range is calibrated to the lifecycle of married households reported in the 2000 Census. The three series are plotted in Figure A1. The weights α_τ determine the relative consumption of children of various ages to 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, and many of the studies generating equivalence scales relate to expenditures. Given the little 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 $\eta = 0.9$, 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 non-durable expenditures reported in the CEX. The data from ATUS indicates that the mean time spent home producing food is roughly one third of total

²⁵ 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.

housework. Similarly, the ATUS suggests that shopping for food accounts for approximately 30 percent of total shopping time. The parameters φ_H and φ_S are both set to 1/3 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 2.3, in line with the estimates reported in Table 3. The scale parameter, ψ_h , is calibrated so that the MRT for home production equals the MRT for shopping 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 3, we set the elasticity with respect to shopping to -0.1 and with respect to goods to 0.2. Recall that the estimates in Table 3 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 Homescan.

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 (defined in footnote 25) 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*, January 1999, 17, 1, 22-35.

- Banks, James and Paul Johnson (1994), "Equivalence Scale Relativities Revisited," *The Economic Journal*, July 1994, 104, 883-890.
- Baxter, Marianne and Urban J. Jermann (1999), "Household Production and the Excess Sensitivity of Consumption to Current Income", *American Economic Review*, Vol. 89 (4), pp. 902-920.
- Becker, Gary (1965), "A Theory of the Allocation of Time", *Economic Journal* 75 (299), pp. 493-508.
- Benhabib, Jess, Richard Rogerson, and Randall Wright (1991), "Homework in Macroeconomics: Household Production and Aggregate Fluctuations", *Journal of Political Economy*, Vol 99 (6), pp. 1166-1187.
- 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*, June 2002, 110, 3, 481-507.
- Browning, Martin and Annamaria Lusardi (1996), "Household Savings: Micro Theories and Micro Facts", *Journal of Economic Literature*, Vol. 34 (4), pp. 1797-1855.
- Browning, Martin, Lars Peter Hansen, and James Heckman (1999), "Micro Data and General Equilibrium Models," in John Taylor and Michael Woodford, ed., *Handbook of Macroeconomics*, Vol 1A, Amsterdam, New York and Oxford: 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*, 2001, November, Volume 33, No. 4, pages 847-875.
- Cronovich, Ron, Rennae Daneshvary, and R. Keith Schwer (1997), "The Determinants of Coupon Usage," *Applied Economics*, 29, pp. 1631-41.
- Ghez, Gilbert and Gary Becker (1975). *The Allocation of Time and Goods over the Life Cycle*, Columbia University Press, New York.
- Gourinchas, Pierre-Olivier and Jonathan Parker (2002), "Consumption over the Life Cycle," *Econometrica*, January 2002, 70, 1, 47-89.
- Greenwood, Jeremy, Richard Rogerson, and Randall Wright (1995), "Household Production in Real Business Cycle Theory", in Thomas Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995.
- Hausman, Jerry and Ephraim Leibtag (2004), "CPI Bias from Supercenters: Does the BLS Know that Wal-Mart Exists?" NBER Working Paper 10712.

- Heckman, James, and Thomas Macurdy (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 Fernandex-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*, Vol 38 (2), pp. 267-290.
- McKenzie and Schargrotsky (2004), "Buying Less, But Shopping More: Changes in Consumption Patterns During A Crisis." mimeo.
- Rios-Rull, Jose-Victor (1993), "Working in the Market, Working at Home, and the Acquisition of Skills: A General-Equilibrium Approach", *American Economic Review*, Vol 83 (4), pp. 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*, Vol 46, pp. 557-579.
- Sorensen, Alan (2000), "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy*, August 2000, 108,4, 833-850.

Table 1: Average Price Paid by Demographic Groups

| Income Category | (1) Average <i>p</i> | (2) Household Size | Average <i>p</i> | (3) Household Composition | Average <i>p</i> |
|-----------------------|----------------------------|--------------------------|---------------------|-------------------------------------|---------------------|
| <\$30,000 | 0.98 (0.002) | 1 | 0.96 (0.003) | Married with Children | 1.03 (0.003) |
| \$30,000- \$50,000 | 1.00 (0.003) | 2 | 0.99 (0.003) | Unmarried Female w/ Children | 0.99 (0.006) |
| \$50,000- \$70,000 | 1.03 (0.004) | 3 | 1.01 (0.004) | Unmarried Male w/ Children | 1.03 (0.013) |
| >\$70,000 | 1.03 (0.005) | 4 | 1.04 (0.004) | Married w/o Children | 1.01 (0.003) |
| | | >4 | 1.06 (0.006) | Unmarried Female w/o Children | 0.96 (0.003) |
| | | | | Unmarried Male w/o Children | 0.99 (0.005) |

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. Standard errors reported in parentheses. For this table, standard errors assume observations are independent within households across time and across households.

Table 2: Average Price Paid as a Function of Shopping Frequency and Total Quantity

| | (1) | (2) | (3) | (4) |
|---|-----------------|---|---|--|
| Dependent Variable: | ln(p) | 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), #Goods ^a , #Varieties ^b | ln(Q)(IV) ^c , #Goods ^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. #Goods defined as log of number of good categories (milk, beer, etc) purchased in month.

b. # Varieties 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 (8 categories) and household composition (8 categories).

Table 2: 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 (0.4) |
| $\ln(Q)$ | | | | | 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 8 age categories. Age, sex, and education refer to household head. See text for definition of categories. Time spent on home production 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.

Table 3: 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 | | |
| 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 | 2.3 | 1.1 ^a |
| Scale Parameter for Home Production | ψ_h | 0.66 | |
| Shopping Technology: | | | |
| Elasticity of Price wrt Time | ξ_S | -0.1 | -0.05 ^a |
| Elasticity of Price wrt Q | ξ_Q | 0.2 | 0.1 ^a |
| Budget Share Parameters | | | |
| Homescan Food in total Expenditure | φ_Q | 0.04 | |
| Food Shopping in total Shopping | φ_S | 0.33 | |
| Food Production in total Housework | φ_H | 0.33 | |
| 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 5.6 and reported in Figure 18.


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 to the 1994 wave of PSID. Homescan sample restricted to households in which the head is at least 25 and the average age of the primary shopper is between 24 and 75. ATUS and PSID samples restricted to households in which the head is between 25 and 75. For married households, head refers to the male (to accord with PSID methodology). All demographics except employment is that of the household head.

a: Not all demographics are available for the full sample of households.


Figure 1: Price Paid by Age



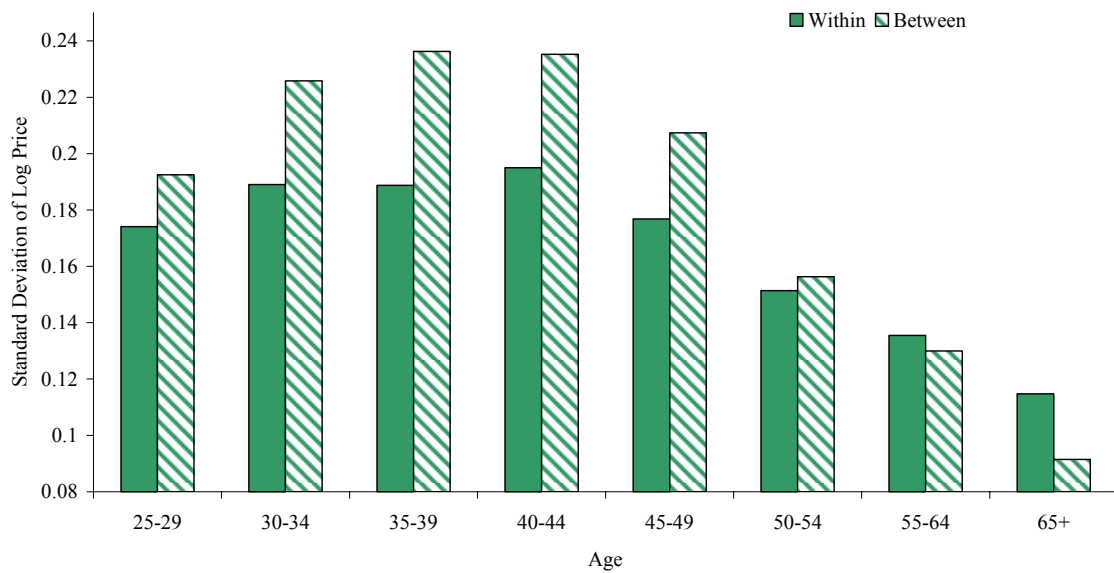
Note: Data from AC Nielsen Homescan. See text for details on construction of Price Index.

Figure 2: Price Dispersion

Panel A: All Goods



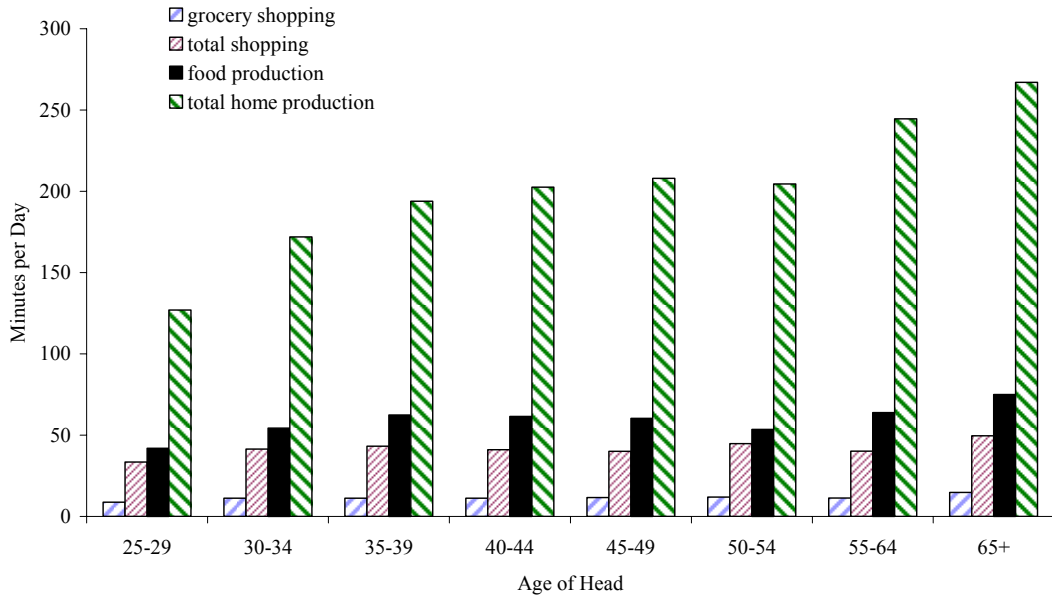
Panel B: Milk




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: Time Use over the Lifecycle

Panel A: All Households




Panel B: Married Households




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

Figure 4: Implied Empirical Opportunity Cost of Time




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

Figure 5: Implied Household Consumption



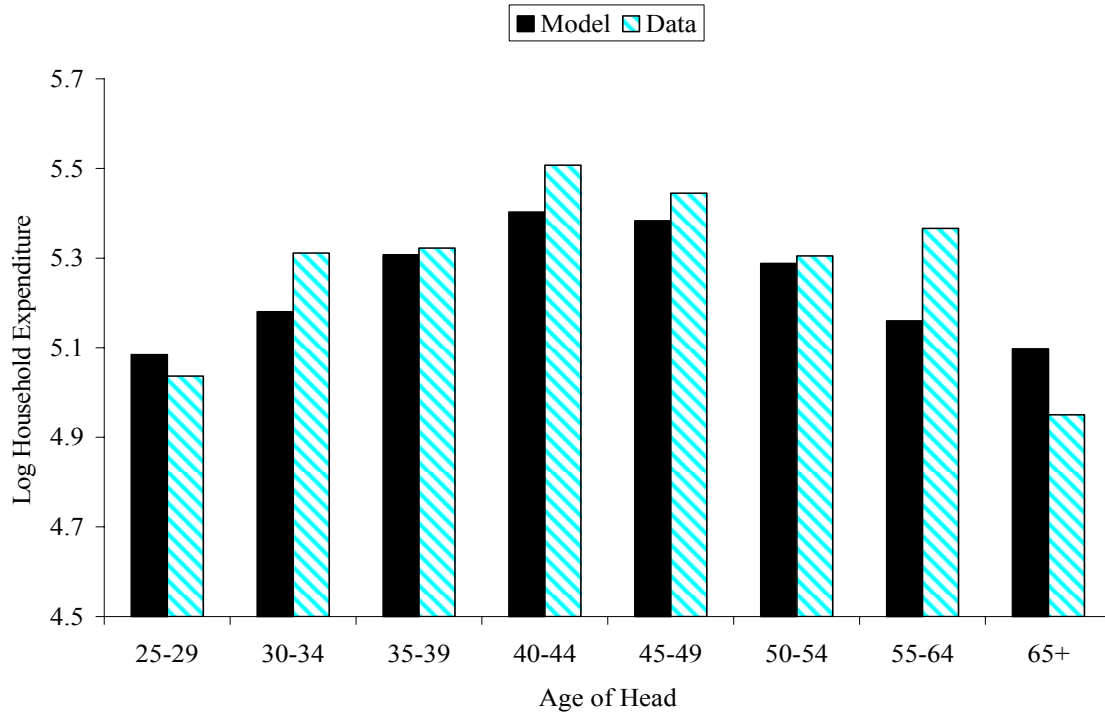
Note: Consumption calculated using parameterized home production function discussed and estimated in Section 3. Inputs of time and goods from ATUS and Homescan datasets, respectively. Figure depicts log deviations from households whose head is aged 25-29.

Figure 6: 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 datasets, respectively. Expenditure from Homescan. Both series normalized to 100 for households whose head is aged 25-29.

Figure 7: Predicted Expenditure over Lifecycle




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

Figure 8: Consumption/Expenditure



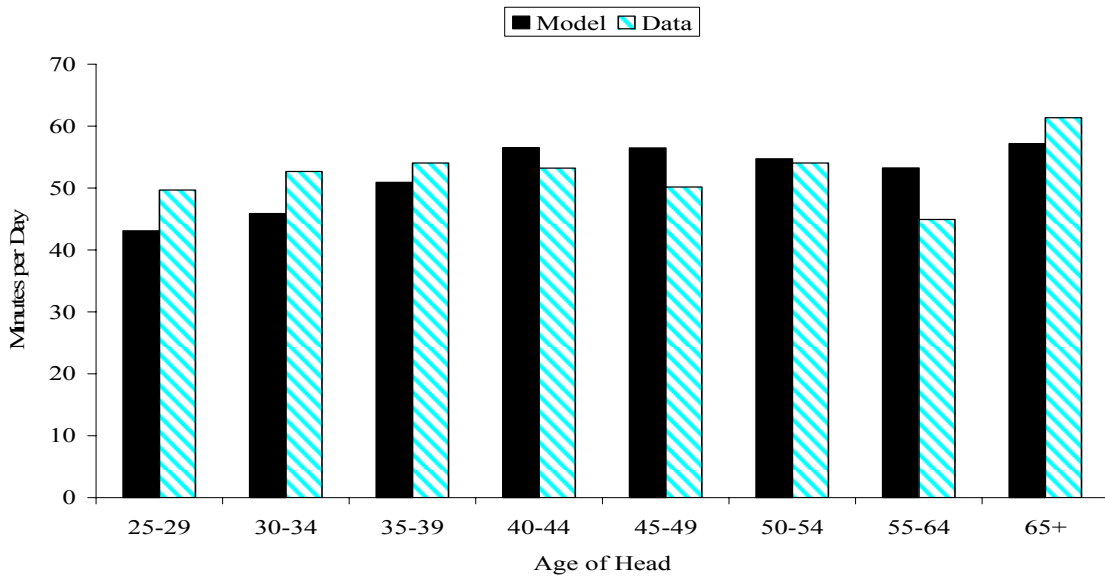
Note: Ratio of household consumption to expenditure. Age 25-29 normalized to 100 for both series. Consumption constructed using market goods and time spent in home production as inputs into production function. Data refers to married households in the AC Nielsen Homescan database with time use merged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

Figure 9: Predicted Price Paid over Lifecycle



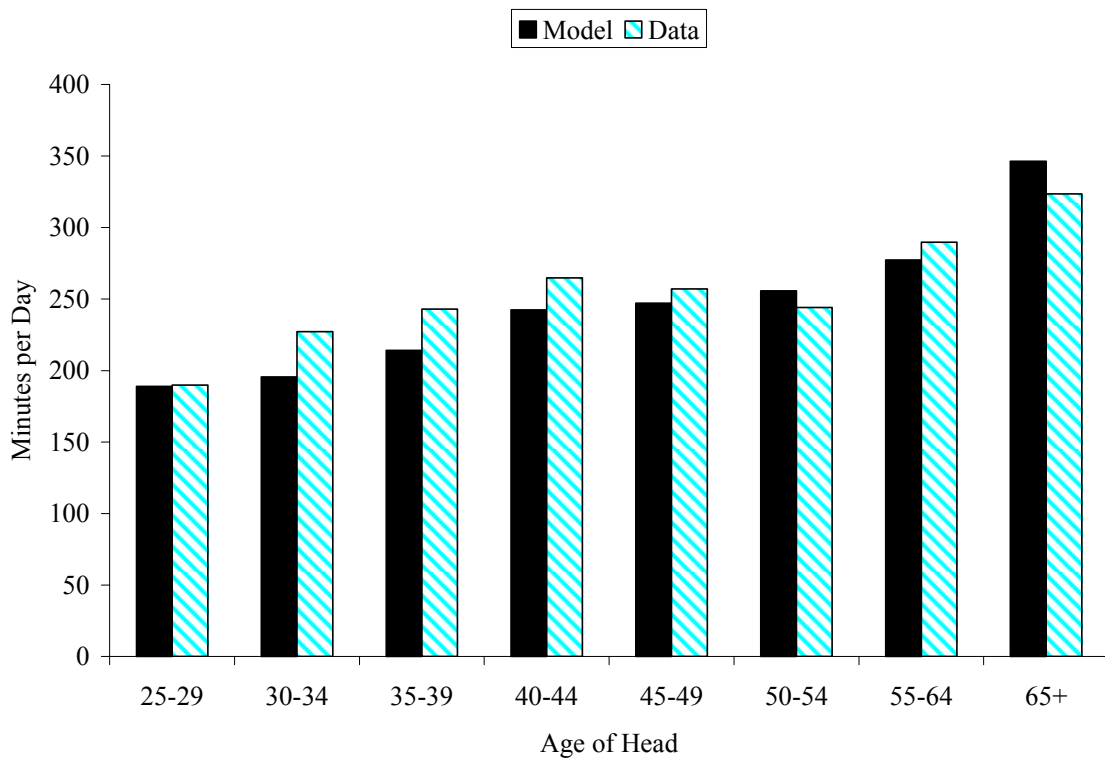
Note: See text for details. Data refers to married households in the AC Nielsen Homescan database. See text for details regarding price index.

Figure 10: Shopping



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


Figure 11: Home Production



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

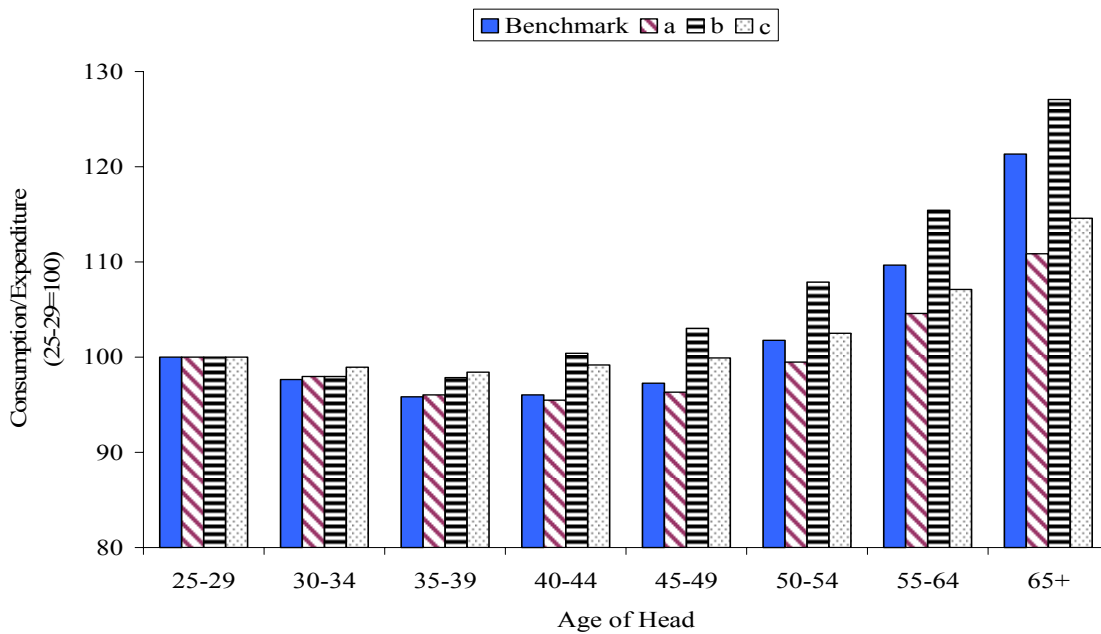
Figure 12: Robustness

Panel A: Expenditures



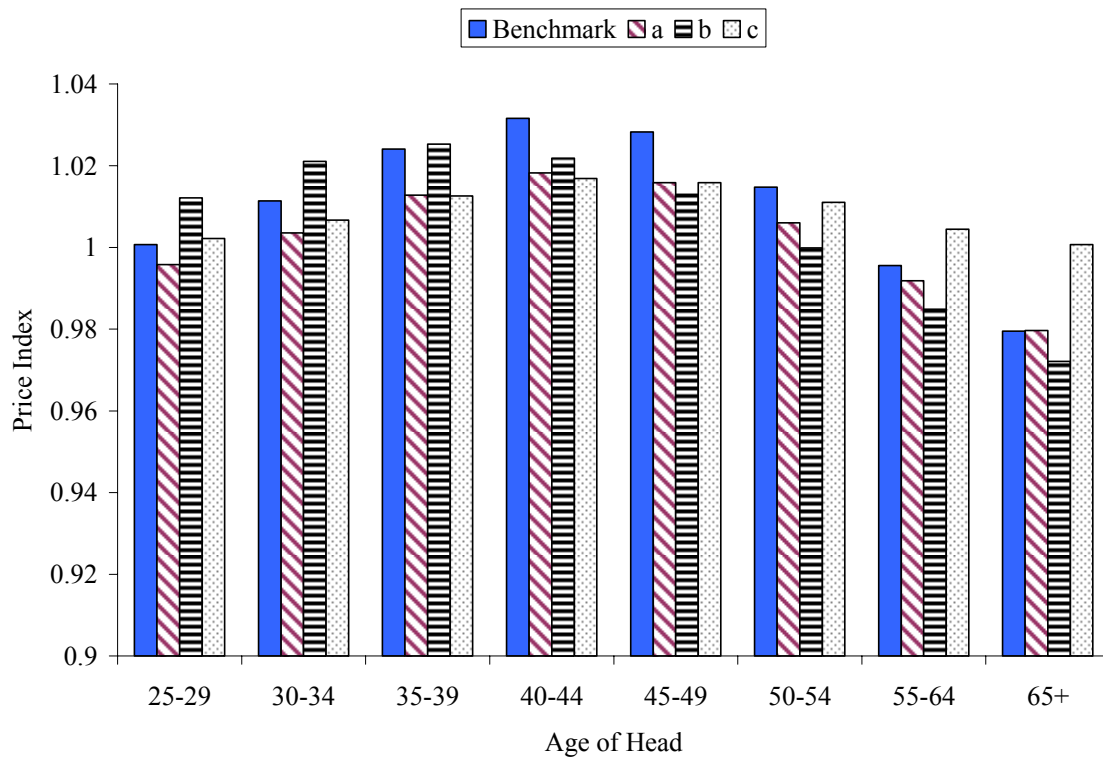
Note: Benchmark is same as "Model" in Figure 13. Robustness a,b,c, refer to predictions from alternative parameterizations of model reported in Column 3 of Table 3.

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.

Panel C: Prices



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 Lifecycle



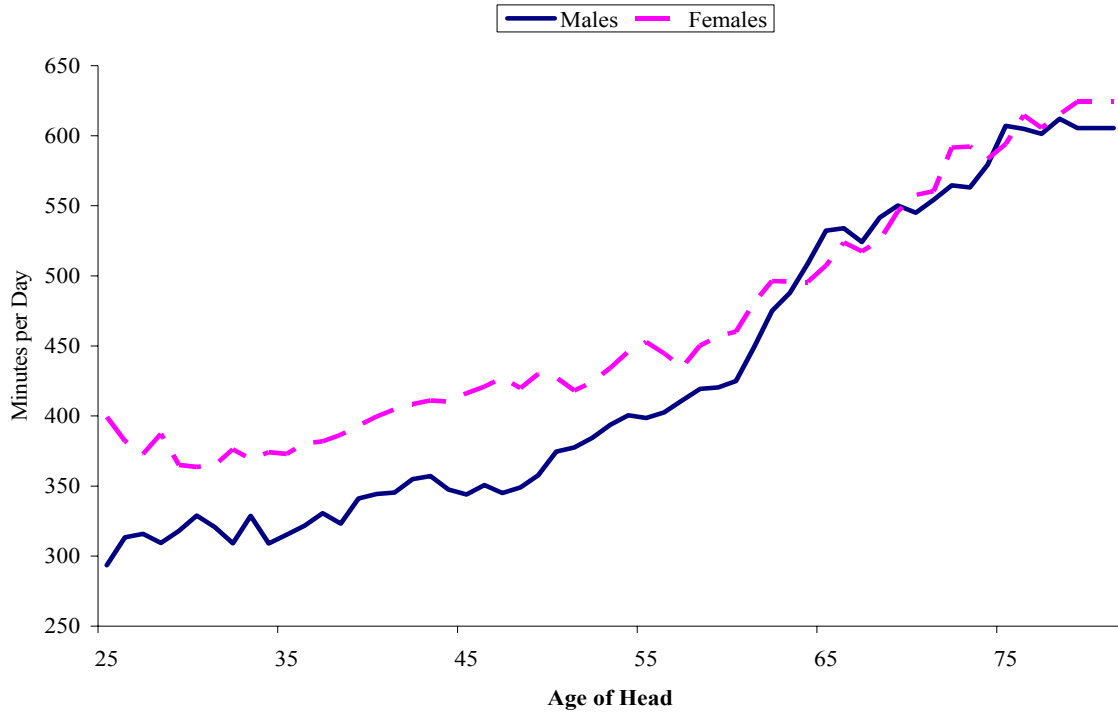
Note: Source: 2000 Census. Series represents 3-year moving average of number of children per household. Age refers to age of household head.

Figure A2: Lifecycle Wage Profile



Note: PSID wage series for men and women 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 to Homescan, wages are expressed in contemporaneous dollars. Wages are conditional on working.

Figure A3: Time Allocation over Lifecycle



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