Consumption, Expenditure and Home Production Over the Life Cycle

Mark Aguiar
Federal Reserve Bank of Boston

Erik Hurst
University of Chicago

January, 2005

Preliminary and incomplete – Do not quote*

Abstract

We show that behind the well-known hump-shaped lifecycle profile of expenditure lies a pronounced hump in the prices paid for identical goods over the lifecycle. We combine this fact with detailed time use data on shopping and home production to estimate a model that replicates the lifecycle profile of expenditure but implies a pattern of consumption consistent with the standard permanent income model. A key innovation of this paper is the use of detailed scanner data of actual prices paid. We find strong evidence that there is large heterogeneity in prices paid across households for identical consumption goods in the same metro area at any given point in time. We show that the heterogeneity in prices paid corresponds directly with the household's opportunity cost of time. Moreover, the cross-sectional dispersion of purchase prices for identical goods peaks in middle age, consistent with the hypothesis that shopping intensity is lowest when the cost of time is highest. Such differences in prices can explain a hump shaped profile in expenditure without relying on households being myopic, liquidity constrained, or impatient. We use our scanner data, coupled with time use data, to estimate both shopping and home production functions. The data suggest that a doubling of shopping frequency lowers a good's price by 15%. From this elasticity and observed shopping intensity, we can impute an opportunity cost of time for the shopper. We find that the price of time peaks in middle age at roughly 35% higher than retirees. Moreover, we estimate an elasticity of substitution between time and goods in home production of close to two. We use these estimated functions to augment a standard PIH consumption model which replicates the hump shaped profile in expenditure, the hump shaped profile in prices paid, the large decline in expenditure during retirement, and observed patterns of time use over the life cycle.

---

* mark.aguiar@bos.frb.org, erik.hurst@gsb.uchicago.edu.

We are particularly grateful to Jean-Pierre Dube for assistance with the Homescan database.

I. Introduction

For more than five decades, empirical work has equated measures of household expenditure with measures of household consumption. Implicit in equating consumption with expenditure is the assumption that all households pay the same price for identical consumption goods. As recognized at least since Stigler (1961), there may be a large dispersion in prices at a given point in time for a particular commodity than can be exploited through search. Additionally, as noted in Becker (1965), household consumption is the output of combining market expenditures with time spent on home production. To the extent possible, as the relative price of time falls, households will substitute time for money spent on market goods by undertaking home production and by shopping more intensively. This margin of substitution makes expenditure a poor proxy for actual household consumption.

With this in mind, we take a new look at the pattern of lifecycle consumption. Specifically, we address the well-documented fact that expenditure is hump-shaped over the lifecycle. This hump is present even once we control for changing family composition. We show that such a hump in expenditures is consistent with the standard lifecycle model/permanent income hypothesis (PIH), including perfect foresight, separability between consumption and leisure in utility, patient consumers, and frictionless borrowing and lending. The key to the exercise is documenting empirically how time substitutes for expenditure through shopping and home production. To estimate key parameters of the shopping and home production functions, we use data from ACNielson's Homescan survey. This survey collects grocery scanner data at the level of the household. Two key features of this data make it preferable to retail establishment-based scanner data. First, the Homescan data, because it is implemented at the household level, is associated with detailed demographics about the purchaser and their family. Second, the Homescan data records purchases across a large variety of retail establishments which the shopper visits. Each purchase in the data base records the actual price paid by the household at the level of the UPC bar code. Because the data includes detailed information about the shopping trip, we can infer the household's shopping intensity.

An important contribution of this paper is that we document that prices paid for identical goods varies with the household's opportunity cost of time. 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 6% more for a given UPC code than households earning less than $30,000 (p-value of difference < 0.01). 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 imposed with having children, households with more children have higher opportunity costs of time than households with no children, all else equal.

One of our most striking results is that prices paid are humped-shaped over the lifecycle. Households in their early 40s pay, on average, between 6% and 8% higher prices 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 (i.e., highest wages) as well as facing the highest non-market opportunity cost of time (i.e., most children). Given this, it is not surprising that one observes a humped-shaped profile in expenditure over the lifecycle. Even if the actual consumption basket for an individual were held constant, the differences in prices paid over the lifecycle will generate a hump-shaped pattern in expenditure.

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 intensity by 100% will pay prices that are 15% less on average. Likewise, holding shopping frequency constant, households who purchase more goods pay higher prices.

Any standard optimizing model implies that the shopper equates the marginal value of shopping for lower prices to the opportunity cost of time. In this way, we can 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 less so than would be inferred 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.

Another important benefit of the price data is we can estimate 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 time use data on home production, we can use this first order condition 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 of the household head. This allows us to calculate a price of time for retirees and married households with one worker. We estimate an elasticity of substitution between time and goods in home production close to two.

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 match the data well along a variety of dimensions. In particular, our model generates a humped shaped profile in household expenditure as well as expenditure per adult equivalent over the lifecycle of similar magnitude as the data. Additionally, our model matches time use data on shopping and home production over the lifecycle as well as the lifecycle profile of prices paid. Lastly, we use data from food diaries which tracks actual food intake, to measure actual consumption over the lifecycle. We find that in contrast to expenditure, actual consumption does not display a hump over the lifecycle.

The lifecycle profile of expenditures has been well documented (Heckman (1974), Attanasio et al (1994), Banks et al (1994), Gourinchas and Parker (2002)). In figure 1, we plot food expenditure over the lifecycle using data from the Panel Study of Income Dynamics (PSID). The panel nature of the PSID allows us to control for idiosyncratic household characteristics (such as permanent income and tastes) through fixed effects. However, similar profiles are present in cross-sectional databases with more complete expenditure data such as the CEX (e.g. see Attanasio et. al. (1994) and Gourinchas and Parker (2002)). Unconditional on household size, households with heads aged in their late 30s/early 40s spend thirty percent more on food than households with heads in their mid 20s and nearly fifty percent more than households with heads in their mid 60s. Given that expenditure is measured as a family variable, and the fact that family size changes over the lifecycle, it is not surprising that expenditure is hump shaped over the lifecycle. However, in Figure 1, we also plot the lifecycle profile of food expenditure controlling for changes in family composition over the lifecycle.[1] To adjust for family structure, we include dummy variables for the number of children aged 0-5, the number of children aged 6-12, the number of female children aged 13- 17, the number of male children aged 13 – 17, the marital status of the head, the sex of the head, and the number of other adults in the household. As seen in Figure 1, household expenditure accounting for family composition is still hump shaped over the lifecycle (the dotted line). This result is consistent with the findings of Gourinchas and Parker (2002), Attanasio et al. (1994), and Banks et al. (1994).

The hump shaped expenditure profile over the lifecycle has often been interpreted as a rejection of the standard permanent income hypothesis (PIH) model. Carroll and Summers (1991) argue that the hump reflects Keynsian myopia or rule of thumb behavior in which agents consume a portion of current income. Gourinchas and Parker (2002) suggest that the hump shape

[1] In the Appendix, we discuss this PSID sample and specifically how we adjust for family structure.

in lifecycle expenditure is evidence of household precautionary savings. In their model, impatient and risk averse households are faced with labor income risk and imperfect credit markets. While young, households accumulate a buffer stock of savings to insure themselves against negative income shocks. This causes consumption to track income for young households. When the target buffer stock is achieved, household impatience dominates generating a declining consumption profile later in life. Gourinchas and Parker interpret the actual decline in expenditure for households later in life as being evidence that households are relatively impatient. Attanasio et al. (1994) also conclude that the hump shape expenditure profile over the lifecycle is likely the result of precautionary behavior. 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 consumption. Lastly, Attanasio et al. (1994) and Blundell, Browning, and Meghir (1994) attribute a portion of the lifecycle profile of expenditure to changing preferences over the lifecycle.

In this paper, we do not provide evidence against precautionary savings, non-separability of preferences, or changing preferences over the lifecycle. We do, however, conclude that the humped shaped profile in expenditure is not evidence, per se, that any of these other theories are true. Even in a world with patient consumers, separable preferences, and no preference shocks, expenditure will be humped shaped over the lifecycle if household shopping and home production responds to the opportunity cost of time. Moreover, the theory we purport matches the data along other dimensions, particularly for prices paid and time use.

Our paper builds on the idea that consumption requires both time and commodities as inputs, a point stressed by Becker (1965).[2] In the Beckerian model, utility may be separable

[2] See also Ghez and Becker (1975). Becker's insight was revived and extended by, among others, Benhabib, Rogerson, and Wright (1991), Greenwood and Hercovitz (1991), Rios-Rull (1993), and Baxter and Jermann (1999). Rupert, Rogerson and Wright (1995, 2000) and McGrattan, Rogerson, and Wright (1997) provide empirical evidence documenting the importance of home production.

over *consumption* and time, but the need for home production implies a non-separability between market *expenditures* and time. Our estimated shopping and home production functions may be viewed as providing a testable, micro-founded structure to a reduced-form non-separability between expenditures and time that rests on utility that is separable in consumption and leisure.

There are two recent papers that use shopping and/or home production to reinterpret existing consumption puzzles. Aguiar and Hurst (2004) shows that a theory of shopping and home production can explain the retirement consumption puzzle. The retirement consumption puzzle states that at the incidence of retirement household non-durable consumption falls by roughly 20% for the average household (Bernheim, Skinner, and Weinberg (2001); Banks, Blundell, and Tanner (2000); Haider and Stephens (2004)). Many authors have interpreted the decline in expenditure as being evidence that household fail to plan sufficiently for retirement. Aguiar and Hurst (2004) shows that time spent on shopping and home production go up in magnitudes that are consistent with the decline in expenditure. Most importantly, they use detailed food diaries from the Department of Agriculture to show that despite the decline in expenditure, actual consumption remains fairly constant as households enter retirement. In other words, the entire decline in expenditure at retirement is due to a decline in prices paid (either through shopping or home production) as opposed to representing a decline in actual consumption. Baxter and Jermann (xxxx) document that a standard permanent income model augmented with home production can yield excess sensitivity in consumption.

The remainder of the paper is set up as follows. In the next section, we describe our price data from the ACNielson Homescan. We document that the price paid for identical goods varies over the lifecycle in ways which are consistent with changes in the opportunity cost of time. Also in that section, we estimate a price function which links prices paid with shopping intensity. In section III, we introduce our time use data and show how shopping time and home production times vary over the lifecycle. We outline our model in section IV. In section V, we discuss our estimation/calibration and show the results of our model simulations. In the final section, we conclude.

II. Prices Paid Over the Lifecycle

A. Data

Our price data is from ACNielsen Homescan. The Homescan data is designed to capture all consumer package goods purchased by the household at a wide variety of retail outlets.[3] We use the Homescan database for Denver covering the period January 1991 through April 1993. The survey is designed to be representative of the metropolitan statistical area (MSA). The Homescan survey is implemented at the household level and contains detailed demographics. 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.[4] The Homescan data base tracks purchases of a given household across multiple retail outlets. In terms of demographics and coverage of multiple outlets, the Homescan database is superior to retail establishment-based scanner data for lifecycle analysis.

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 the purchased goods. The shopper also enters the age and sex of the primary shopper for the trip.[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. ACNielson maintains a database of current prices for all stores within the metropolitan area. Given the UPC code, store and date, ACNeilson provides a purchase price (before coupons and discounts) for each

[3] While we only have access to the packaged goods component of the survey, there are additional supplements which focus on fresh foods and a separate supplement that focuses on consumer durables.
[4] For those not working, we do not know if they are unemployed, retired, disabled, or a home maker.
[5] While the age and sex of the primary shopper is trip specific, we use the transaction frequency –weighted average age (for each family in each year) as our measure of shopper age. In practice, based on age and sex it appears the same person does a vast majority of the shopping within a household. On average, the price paid within a household is invariant to who does the shopping for that particular trip.

transaction. (The information regarding coupons and other discounts entered by the shopper are for the entire shopping trip, as opposed pertaining to any particular good. xx)

Households in our Homescan data base participate for up to twenty-eight months. To the extent that households participate over time, we have a panel of household purchases. The panel is subject to two caveats. First, the demographic variables are only updated annually. Second, there is some attrition in the sample over time. However, the median household is in the sample for 27 out of 28 months.

While the actual store of purchase is not included in our Homescan database, we do know the type of store (convenience, grocery, price clubs, superstores, etc.). For our 1993-1995 Denver Supplement of 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 71. This restriction leaves us with just over 2,000 separate households. Descriptive statistics for our sample are shown in appendix table A1.

The average monthly expenditure for packaged goods for our sample was \$180/month (in 1996 dollars). The Panel Study of Income Dynamics PSID records total expenditure for all food purchased at any grocery establishment. This measure excludes food purchased at restaurants, fast food establishments, and cafeterias. Using a sample of household heads between the ages of 24 – 71 from the 1993 – 1995 PSID, we find that average monthly expenditure on all grocery goods is \$340 (in 1996 dollars). 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. Additionally, it may be the case that households fail to scan in all grocery items in the Homescan database. This latter issue may influence our sample means regarding frequency of shopping or coupon usage. However, under the reasonable assumption that households do not systematically select goods to scan based on price paid, it should not bias our estimate of the price paid per UPC code for each household.

[to be added: discussion of cross sections and cohort effects]

B. Prices Paid and the Opportunity Cost of Time

Basic economics suggests that 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 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.[6]

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* purchased by household *j* on shopping trip (date) *t*. Let $q_{i,t}^{j}$ represent the corresponding quantity purchased. Total expenditures during month *m* is simply

$$X_m^j = \sum_{i,t \in m} p_{i,t}^j q_{i,t}^j \tag{2.1}$$

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:

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

where

$$\overline{q}_{i,m} = \sum_{j,t \in m} q_{i,t}^j \tag{2.3}$$

[6] Recently, Hausman and Leibtag (2004), using Homescan data, document that stores like Walmart offer prices between 5% and 55% less than the same product in traditional grocery stores.

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 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,

$$Q_m^j = \sum_{i,t \in m} \overline{p}_{i,t} q_{i,t}^j \tag{2.4}$$

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 then normalize the index by dividing through the average price index across households within the month, ensuring that each month the index is centered around one:

$$p_m^j = \frac{\tilde{p}_m^j}{E(\tilde{p}_m^j)} \tag{2.5}$$

where

$$\tilde{p}_m^j = \frac{X_m^j}{Q_m^j} \tag{2.6}$$

and

$$E\left(\tilde{p}_m^j\right) = \frac{1}{J} \sum_j \tilde{p}_m^j \tag{2.7}$$

The price index shares the typical feature (as with Laspeyres and Paasche indices(sp?)) 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 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 case. In the usual case, 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 comparisions, such as the CPI. By revealed prefence, 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 interepret 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. The price differentials are for the *identical* goods.

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.

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 Figure 2, we show the log of the price index and the log of the price index including coupon use for households within the four income categories. The results are striking. Households who earn less than \$30,000 a year pay, on average, 6% lower prices than households who earn over \$70,000 (p-value of difference < 0.01). Households who earn between \$30,000 and \$50,000 pay, on average, 4% lower prices than households who earn over \$70,000 (p-value of difference < 0.01). The difference between these groups 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.[7] Of course, more than the price of time varies across these income categories. In particular, the higher income households are purchasing a larger basket of goods. We will explore how this influences price in the regressions reported in section II-D.

A second measure of the opportunity cost of time is alternative demands on time at home arising from children. In Figure 3, we see that households with larger families pay higher prices than households with smaller families. Specifically, households with only one household member pay 14% less for an item compared to households with family size greater than or equal to 5 (p-value < 0.01). Similarly, single females with no children pay 11% lower prices than married couples with children (p-value < 0.01), while single males without children pay 6% lower prices than married couples with children (p-value < 0.01). These latter patterns are documented in Figure 4. These differences holds conditional on income as well. When we regress the log price premium measure on both income categories and family size categories, both sets of regressors enter significantly.

Given that both the arrival of children and household wages have a lifecycle component, we would expect our price premium measure to vary over the lifecycle. Using a sample of PSID households (discussed in the appendix), we find that children in a married household peak when the head is in his or her early 40s. Correspondingly, the number of 6-12 year olds peaks within a married household when the household head is in his or her late 40s. As seen in Appendix Figure A1, wages of the male household head within a married couple peak around the age of 45. Both the profiles of children and market wage suggest the demands on time are greatest in middle age.

In Figure 5, we show the lifecycle profile of our log price premium for all households and for married households. Our age measure is the age of the "household head" within the household.[8] Consistent with our premise, households in their middle 40s pay the highest prices.

---

[7] There is some evidence that prices are higher in poor neighborhoods (cites xx). 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) pay 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 apparent non-monotonicity.

[8] For married households, we follow standard practice (as in the PSID) and use the age of the male.

Specifically, unconditional on marital status households aged 45-49 pay 6% higher prices than households aged 25-29 and 4% higher prices than households aged 65+. Conditional on marital status, households aged 40-44 pay 10% higher prices than households aged 25-29 and 6% more than those older than 65.

Finally, in figure 6 we plot price *dispersion* over the lifecycle. Specifically, for each UPC code and each month we split the observed purchase prices by age of shopper. For each subsample, we calculate the standard deviation of prices and then average over goods. This averaged is then plotted against the age of the household head in figure 6. To ensure that the observed effect is not due to a changing basket of goods over the lifecycle, figure 6 also breaks out a single good – milk – and performs the same analysis. In both cases, we see that dispersion is highest in middle age and lowest in retirement. This pattern is easily interpreted in our 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. We have performed the same analysis for dispersion on a given *day* and found a similar pattern.

C. 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.

In Figure 7, we plot the time spent shopping, engaging in home production, and childcare for all households (panel a) and married households (panel b) over the lifecycle. This data is from the American Time Use Survey (ATUS) conducted by the BLS. (See Appendix A for a discussion of this dataset). Peak shopping and home production occurs occurs in the early 40s and the mid 60s. Households in their mid 40s have the largest family sizes and, as a result, have the greatest shopping needs. They also devote the most time to childcare. Households in their post retirement years have the lowest opportunity cost of time and therefore shop more intensively for a given basket and rely more on home production. Young households shop relatively little because they buy relatively few goods and have work and education demands on their time.

D. 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:

$$p = p(s, Q) \tag{2.8}$$

where p is our price index, s is the amount of time shopping, and Q is the amount of goods purchased. Our hypothesis is $\partial p / \partial s < 0;\ \partial^2 p / \partial s^2 > 0;\ \partial p / \partial Q > 0$. In other words, holding Q constant, households who shop more can 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* as purchases evaluated at the mean prices (as defined in equation (2.4)). We will also explore alternatives such as the physical quantity of goods as well as 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 specification:

$$p_m^j = \alpha_0 + \alpha_1 \ln(s_m^j) + \sum_{k=1}^{5} \alpha_{k+1} \left(Q_m^j\right)^k + \varepsilon_m^j \tag{2.9}$$

where we normalize our price index (dependent variable) to have a sample mean of one. This function specifies that price is semi-log in shopping time and is a function of a 5$^{\text{th}}$ order polynomial in quantity. Column 1 of Table 2 shows the estimates of (2) using the monthly expenditure at average prices for our measure of Q. We estimate α_1 to be -0.16 (standard error = 0.01). Recall that we have normalize *p* to have a mean of one. This implies that a 100% increase in shopping frequency reduces the price paid by 16 percent. We also find that the quantity purchased has a statistically significant impact on price, with an elasticity evaluated at the average *Q* of 0.21. That is, for a given shopping frequency, the more purchases a shopper makes the higher the price of the average good.

In columns 2 through 5 of Table 2, we explore other specifications of the price function. The results are stable across different specifications. Column 2 includes a quadratic in shopping frequency which indicates price is decreasing and convex in shopping time. Columns 3 and 4 uses alternatives measures of quantity – total number of goods ($\sum_{i,t \in m} q_{i,t}^j$) and the number of different UPC varieties purchased in a month, respectively. Column 5 explores a log-log specification. In all cases price is decreasing and convex in shopping time and increasing in quantity purchased. We have also included higher order terms for both price and quantity and found them to be statistically insignificant with no change in the estimated elasticities.

(clarify) 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. Firstly, 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. Secondly, our price index measures how much you pay 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 age and household composition dummies. Not surprisingly given our discussion above, the results, reported in the final column of Table 1, are not sensitive to this issue.

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, as discussed in the appendix, alternative data sources suggest that time per trip is not constant over the lifecycle. In fact, the 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 included the 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.

III. An Augmented PIH/LC Model with Shopping and Home Production

In this section we introduce a standard lifecycle/permanent income (PIH/LC) model which is augmented to allow for shopping and home production. We model the household as a single decision making unit. The household consists of one to two adults and $\sum_{\tau=1}^{T_A} n(\tau)$ children, where $n(\tau)$ represents the number of children aged τ, and T_A is the age after which children leave the household and receive no further support (in time or money) from the parents. We treat the arrival of children as exogenous. We denote the age of the household by a single index t, which runs from one when the household is formed through T when the adult members of the household die. If there are two adults in the household, they both face the same T (i.e., we abstract from the issue of widows or widowers).[9]

All exogenous variables are deterministic functions of t. In particular, the number of children at each age and wages are functions of t.

A. Budget Sets:

Each adult in the household is endowed each period with a unit of time, which we normalize to one. Time can be allocated to leisure (ℓ), labor in the formal market (L), shopping (s), home production (h), and taking care of children. The time spent taking care of a child aged τ is denoted as $\ell(\tau)$, where we echo the notation of leisure to reflect that children may obtain utility from their parent's attention.

The household faces up to two wage rates at a point in time. We denote these w_t* and w_t. The asterisk notation identifies the adult with the lower opportunity cost of time (which will also be the adult with the lower wage). In every household, at least one adult shops and (it will be shown below) the shopper has the lower opportunity cost of time. The asterisk therefore designates the "shopper". Note that the asterisk is not tied to one household member over time, but identifies the member with the lower cost of time at a particular point in time (or the lone

[9] We can also think of T as a more general dissolution of the household, including due to divorce. The analysis will be exactly the same as long as no household assets are carried forward beyond T.

adult in a single-head household). It may be the case that the other adult also shops and/or engages in home production. This will be the case if the two adults have the same opportunity cost of time. In this case, the model's symmetry makes the distinction immaterial.

In each period, the household purchases Q_t units of the market good at price p_t. As discussed above, the price paid is a function of the quantity purchased and the time spent shopping. That is

$$p_t = p(s_t^* + s_t, Q_t) \tag{4.1}$$

where the adults are perfect substitutes in their ability to shop.

The household has access to borrowing and lending at the interest rate r. Household assets at time t are denoted a_t, which can be negative in case the household is a net debtor. The evolution of assets is described by

$$a_{t+1} = (1+r)a_t + w_t^* L_t^* + w_t L_t - p_t Q_t \tag{4.2}$$

Iterating the difference equation for a given initial asset level a_0 and imposing the terminal condition $a_{T+1} \geq 0$, we have

$$\sum (1+r)^{-t} \left(w_t^* L_t^* + w_t L_t \right) + a_0 \geq \sum (1+r)^{-t} p_t Q_t \tag{4.3}$$

B. Preferences:

We denote household utility at time t by $U(c^*, c, c(\tau), \ell^*, \ell, \ell^*(\tau) + \ell(\tau), t)$, where c*, c, c(τ) refers to the consumption of the shopper, the other adult (if present), and the vector of children aged $\tau = 1, \ldots, T_A$, respectively. The terms ℓ*, ℓ,ℓ*(τ)+ℓ(τ) refer, respectively, to the leisure of the shopper, the leisure of the other adult, and the total amount time spent by both parents on the vector of children indexed by τ. We do not need to include the child's leisure as a separate argument as children spend all their time as leisure.

We assume that within the household utility, is additively separable across individuals, so we can write:

$$U = u(c^*, \ell^*) + u(c, \ell) + \sum_{\tau} \omega(c(\tau), \ell^*(\tau) + \ell(\tau), \tau) n(\tau) \tag{4.4}$$

Note that we have imposed symmetry of preferences across adults. Children's preferences are distinct and may vary with the child's age. If there is only one adult, the second term on the right is set to zero. We assume that utility functions satisfy the standard properties (strictly increasing, strictly concave, twice differentiable, plus the Inada conditions in consumption and leisure).

The household's problem can now be written as:

$$\max_{\left\{c^*, c, c(\tau), \ell^*, \ell, \ell^*(\tau), \ell(\tau), h^*, h, s^*, s, L^*, L, Q\right\}_{t=1,\ldots,T}} \sum_{t=1}^{T} \beta^t \quad U(c_t^*, c_t, c_t(\tau), \ell_t^*, \ell_t, \ell_t^*(\tau) + \ell_t(\tau), t) \tag{4.5}$$

The maximization is subject to the budget constraint **Error! Reference source not found.** and the following additional constraints.

The first constraint is the home production function:

$$C_t \equiv c_t^* + c_t + \sum_{\tau} c_t(\tau) n(\tau) = f(Q_t, h_t^* + h_t) \tag{4.6}$$

Home production combines market goods and time into consumption goods, as in Becker (1965). We assume that f is increasing, strictly concave, and satisfies the Inada conditions in each of its arguments.

The next two constraints are the time constraints on the shopper and the other adult:

$$L_t^* + \frac{k}{\eta} L_t^{*\eta} + s_t^* + h_t^* + \ell_t^* + \sum_{\tau} \ell_t^*(\tau) n_t(\tau) = 1 \tag{4.7}$$

$$L_t + \frac{k}{\eta} L_t^{\eta} + s_t + h_t + \ell_t + \sum_{\tau} \ell_t(\tau) n_t(\tau) = 1 \tag{4.8}$$

The second term in each line represents a convex cost of providing labor to the market, where $k \geq 0$ and $\eta \geq 1$ are parameters. This is designed to capture commuting costs and other costs associated with labor supply.[10]

The final set of constraints consists of the non-negativity constraints for time components and consumption. We assume Inada conditions on utility in regard to consumption and leisure, so the non-negativity constraints on these never bind. However, they are relevant for labor supply to the market and (individual) time spent on shopping, home production, and children.

We let λ denote the Lagrange multiplier on the lifetime budget constraint (4.3). The multiplier on constraint (4.6) at age t is given by $\beta^t \mu_{C,t}$. The multipliers on the time constraints (4.7) and (4.8) at age t are given by $\beta^t \mu^*_t$ and $\beta^t \mu_t$, respectively. We let $\beta^t v_t(x)$ denote the multiplier on the constraint that $x \geq 0$ at time t, where x represents the choice variables. We set up the problem so the usual Kuhn-Tucker conditions ensure that these multipliers are non-negative.

The first order conditions for the household's problem are as follows. The household's total utility from consumption is maximized by equating marginal utilities to each other and to the shadow cost of producing consumption goods:

$$\frac{\partial u}{\partial c_t^*} = \frac{\partial u}{\partial c_t} = \frac{\partial \omega}{\partial c_t(\tau)} \forall \tau = \mu_{C,t} \tag{4.9}$$

The marginal utilities of leisure of the adults are set equal to their respective shadow values of time. For the children, at least one of the parents provide time inputs (due to the Inada conditions). Given that the child treats each parent's time as perfect substitutes, both parents provide time to children only if their shadow values of time are equal (i.e. $\mu_t^* = \mu_t$). Otherwise the parent with the lower opportunity cost of time does all the parenting.

[10] One might think of such costs as fixed costs. However, the quadratic cost provides a more tractable framework. Moreover, as we are trying to match the behavior of the "average" household, we can follow much of the investment literature and model what appears as a fixed cost at the level of the individual agent as a convex cost for the representative agent.

$$\frac{\partial u}{\partial \ell_t^*} = \mu_t^* \tag{4.10}$$

$$\frac{\partial u}{\partial \ell} = \mu_t \tag{4.11}$$

$$\frac{\partial \omega}{\partial(\ell_t^*(\tau) + \ell_t(\tau))} = \mu_t^* - \nu_t(\ell^*(\tau)) = \mu_t - \nu_t(\ell(\tau)),\ \forall\,\tau \tag{4.12}$$

The marginal product of time spent in home production (times the shadow price of consumption) is set equal to the opportunity cost of time of the parent(s) engaging in the activity:

$$\mu_{C,t}\frac{\partial f}{\partial(h + h^*)} = \mu_t^* - \nu_t(h^*) = \mu_t - \nu_t(h) \tag{4.13}$$

The marginal value in dollars of shopping for goods is the change in price times the quantity purchased. This is converted into utility by λ, the shadow value of a dollar (at time zero). In turn, the marginal value of shopping is set equal to the opportunity cost of time.

$$-\frac{\partial p}{\partial(s_t^* + s_t)}Q_t\lambda = (1+r)^t\beta^t(\mu_t^* - \nu_t(s_t^*)) = (1+r)^t\beta^t(\mu_t - \nu_t(s_t)) \tag{4.14}$$

The marginal cost of a market good, including its potential affect on price, is set equal to the marginal product of Q in home production:

$$\lambda\left(\frac{\partial p}{\partial Q_t}Q_t + p_t\right) = (1+r)^t\beta^t\mu_{C,t}\frac{\partial f}{\partial Q_t} \tag{4.15}$$

Finally, the value of an additional unit of labor is set equal to the opportunity cost of time for each adult:

$$\lambda w_t^* = (1+r)^t \beta^t \left(\mu_t^* (1 + kL_t^{*\eta-1}) - \nu_t(L^*) \right) \tag{4.16}$$

$$\lambda w_t = (1+r)^t \beta^t \left(\mu_t (1 + kL_t^{\eta-1}) - \nu_t(L) \right) \tag{4.17}$$

The first order conditions imply many of the familiar properties of lifecycle behavior. However, we develop a few propositions to highlight aspects that are unique to models with home production and shopping.

Proposition 1: The person who engages in child care, shopping, or home production ("the shopper") has an opportunity cost of time that is less than or equal to the other adult's. If both adults engage in these activities, they have equal shadow values of time.

Proof: This follows directly from the fact that the adults are perfect substitutes in time spent on children, shopping, and home production. Out of two perfect substitutes, it is always optimal to choose the input with the lower price. This can be seen mathematically from (4.12), (4.13) or (4.14), and the fact that $\nu(\bullet) \geq 0$.

Proposition 2: If $\frac{\partial^2 u}{\partial \ell \partial c} \geq 0$*, the "shopper" is the low wage adult. That is,.* $\mu_t \geq \mu_t^*$ *implies* $w_t \geq w_t^*$.

Proof : If $\mu_t = \mu_t^*$, the distinction is immaterial and we can arbitrarily assign the asterisk to the lower wage. Similarly, if neither adult works, then concavity of the utility functions implies that the household should equate their marginal utilities of leisure to maximize total utility. This, plus (4.10) and (4.11), imply $\mu_t = \mu_t^*$. Now consider the case in which $\mu_t > \mu_t^*$ and at least someone in the household works in the market. Suppose (to generate a contradiction) that $\mu_t > \mu_t^*$ but $w_t < w_t^*$. In this case, it is never optimal for $L_t \geq L_t^*$. This follows directly from the fact that the two are perfect substitutes on every other time dimension except market labor. If $L_t \geq L_t^*$, it would be possible to keep all else equal but raise income by having the high wage worker work less in the market and the low wage worker work less at home/shopping. Therefore, $L_t < L_t^*$. This, plus $\mu_t > \mu_t^*$ implies the "shopper" does all the shopping/homework plus more market work and therefore $\ell_t > \ell_t^*$. If utility is separable in leisure, this implies $\mu_t < \mu_t^*$, a contradiction. If utility is nonseparable and $\frac{\partial^2 u}{\partial \ell \partial c} \geq 0$ then $\ell_t > \ell_t^*$ and (4.9) imply that $c_t > c_t^*$. However, given that utility is concave, it would improve total household utility with no change in expenditure or income to lower c_t and ℓ_t (by shopping more) and raise the corresponding values of c^*_t and ℓ^*_t.accordingly. Therefore, this cannot be an optimum.[11]

Proposition 3: Marginal utility of consumption follows an augmented Euler Equation:

$$\frac{\mu_{C,t+1}}{\mu_{C,t}} = (1+r)^{-1}\beta^{-1}\left(\left(\frac{\partial p}{\partial Q}Q_{t+1} + p_{t+1}\right)\Big/\left(\frac{\partial p}{\partial Q}Q_t + p_t\right)\right)\left(\frac{\partial f}{\partial Q}_t \Big/ \frac{\partial f}{\partial Q}_{t+1}\right) \tag{4.18}$$

Proof: This follows directly from (4.15).

In a standard model, $\partial p / \partial Q = 0$ and $\partial f / \partial Q = 1$, yielding the typical Euler Equation. In the current model, the inter-temporal relative price of a unit of consumption also depends on the relative price of an additional market good and the relative marginal product of an additional market good. If the marginal cost of Q in the future is relatively high or the marginal product of Q in the future is relatively low, households prefer to consume more today. This will be important in understanding the path of consumption as prices and quantities vary over the lifecycle.

---

[11] The statement of the proposition may not hold if $u_{c\ell}<0$. Suppose that $u = u(c+\ell)$ and $w<w^*$. Then the high wage adult could do more market labor as well as all the shopping, but still maintain the same marginal utilities of consumption and leisure as the other adult by consuming more. That is, the first order conditions are satisfied if $c + \ell = c^* + \ell^*$.

C. Estimation of the Home Production Function

An important implication of the first order conditions is that the shopper's opportunity cost of time over the life cycle can be inferred, up to a constant, from the amount of time spent shopping and the quantity purchased. That is, (4.14) plus the fact that at least one person shops implies:

$$\ln(\mu_t^*) = \ln\left(-\frac{\partial p}{\partial s} Q_t\right) + \ln(\lambda) \tag{4.19}$$

where we have assumed $\beta=1/(1+r)$. From our estimation of the price function, we can construct the first term on the right hand side for each household. [to be added]In figure xx, we plot the lifecycle profile of this series by averaging over households and taking differences from the age group xx-xx.[12] For reference we also plot the average wage of men and women using PSID data. 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 is not as severe as that of wages. 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.

In a similar vein, the first order conditions imply that the marginal rate of transformation (MRT) between time and market goods in shopping equals the marginal rate of transformation in home production:

$$\frac{\dfrac{\partial f}{\partial (h^* + h)}}{\dfrac{\partial f}{\partial Q}} = \frac{\dfrac{\partial p}{\partial (s^* + s)} Q}{\dfrac{\partial p}{\partial Q} Q + p} \tag{4.20}$$

[12] Given the cross-sectional nature of the data, we may be combining cohort effects with pure lifecycle effects.

While this is a standard implication of optimality, it is very important for our empirical work. This is because we can estimate the MRT for shopping directly from our scanner data set. As we also observe time spent in home production and quantity purchased, we can 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). To parameterize the home production function would rely on the fact that the MRT between time and goods in home production equals the relative price of time and goods. 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.

The approach is also valid even if we observe only a subset of goods that enter nonseparably in utility. The key assumption is that for any consumption good x, we observe all of the inputs (time and market goods) that are used to produce x. It is not necessary to observe substitutes or complements in utility from consuming x. However, if both market goods y and z

are used to produce x and they are complements or substitutes in production, we need to observe both y and z.[13] Similarly with the shopping function.

As a benchmark, consider a CES home production function:

$$f(h+h^*,Q)=\left(\psi_h(h+h^*)^\rho+\psi_Q Q^\rho\right)^{\frac{1}{\rho}} \tag{4.21}$$

where the elasticity of substitution is given by $\sigma=\dfrac{1}{1-\rho}$ and . The MRT between time and goods is

$$\frac{\psi_h}{\psi_Q}\left(\frac{h+h^*}{Q}\right)^{\rho-1} \tag{4.22}$$

Taking logs on both sides of (4.20), we have

$$\ln\left((h^*+h)/Q\right)=\sigma\ln\left({\psi_h}\big/{\psi_Q}\right)-\sigma\ln\left(-\frac{\partial p}{\partial(s^*+s)}\bigg/\left(\frac{\partial p}{\partial Q}Q+p\right)\right) \tag{4.23}$$

From our price data, we can estimate and fit the right hand side for each age group. From our time use data we have measures of time spent in home production. Due to measurement error and specification error, (4.23) will not hold exactly in our data set. Averaging over a lot of households at each age group will minimize the errors in variables problem (that is, we can think of our age dummies as instruments). The averaging will also correct for idiosyncratic "productivity" shocks that are uncorrelated with age.

Estimating equation (4.23) with the unit of observation being three-year age groups (a total of 4,842 households averaged into 17 groups) yields an estimate of σ=2.14, with a standard error of 0.34.

---

[13] In particular, our measure of market goods is restricted to those reported in the Homescan database. To the extent that other goods are used as inputs, our measure of the marginal rate of transformation may be incorrect. However, as long as the omitted input impacts the marginal products of time and goods equivalently, it will not affect the ratio of marginal products and therefore our estimate of the marginal rate of transformation will be accurate.

One concern with (4.23) 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:

$$\ln\left(h^* + h\right) = \sigma \ln\left(\psi_h \big/ \psi_Q \right) - \sigma \ln\left(-\frac{\partial p}{\partial (s^* + s)} \Big/ \left(\frac{\partial p}{\partial Q} Q + p \right) \right) + \ln\left(Q\right) \qquad (4.24)$$

The estimate of σ in this case is -1.86 with a standard error of 0.66, 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 0.74 with a standard error of 0.53.

V. Model Solution and Calibration:

To simulate the model, we need to make some additional assumptions regarding functional form and parameters. The benchmark parameters are reported in table xx. We have already discussed parameter estimation and functional form for the price function and home production function. The only unidentified parameters are ψ_h and ψ_Q. As only the ratio can be estimated from the first order conditions, we set the latter equal to one and the former is calibrated so that the average time in home production matches that in the data.

For preferences, we assume adult utility is given by:

$$u(c, \ell) = \frac{c^{1-\gamma}}{1-\gamma} + \theta \ln(\ell) \qquad (5.1)$$

As a benchmark, we let $\gamma = 2$. We calibrate θ so that the average time spent shopping matches that in the time use data set. We set the labor supply ("commuting") parameters k and η so that the high wage worker devotes about one third of his/her time to market labor. We also set $\beta = \frac{1}{1+r} = 0.98$.

We assume children's marginal utilities of consumption and leisure are proportional to an adult:

$$\omega\left(c(\tau), \ell^*(\tau)+\ell(\tau)\right)=\frac{\alpha_{1,\tau} c(\tau)^{1-\gamma}}{1-\gamma}+\alpha_{2,\tau} \ln\left(\ell^*(\tau)+\ell(\tau)\right) \tag{5.2}$$

We divide children into three age groups, younger than 6, 6 through 13, and older than 13, corresponding to $\tau=1,2,3$, respectively. Somewhat arbitrarily, we assume that children in these age groups need to respectively consume 0.2, 0.3, 0.5 fraction of an adult's consumption to achieve the same utility. We set children's need for parental time at 0.3, 0.5, 0.1 as well, so that younger children value parental time more than consumption relative to their older siblings. We have experimented with alternative values. <<Discuss robustness.>>

The adult lifecycle begins at 21 and ends at 81. The number of children at each age is taken to be exogenous in the model. The number of children of age τ for a household of age t calculated using data from the PSID. Specifically, we average the number of children within each age bracket over households whose household head is age t. The number of children in the three respective age categories as a function of household age is plotted in figure xx. Due to the smallness of the sample, we do not estimate the number of children beyond age xx, which we have set to zero.

For wages, we take the lifecycle profile of wages from 21 through 6x? for men and women from the PSID adjusted for fixed effects. We average the wage of households in the PSID and feed this series through the lifecycle model. In averaging over households, we set the wage of the unemployed and those out of the labor force to zero. We assume wages are zero for households older than 6x. The relevant wage series are plotted in figure xx.

We need to account for the fact that the shopping and home production functions are estimated using a subset of household expenditures. We therefore scale the model up to total household expenditure using data from the CEX. Specifically, the observed purchases in the Homesan database represent roughly 5% of total monthly expenditures. Moreover, the NHAPS data discussed in appendix A indicates that shopping for food accounts for roughly one quarter of total shopping time. We assume other goods enter additively separably in utility, shopping and home production. This implies that the various goods interact only through the budget constraint and the time constraint. Accordingly, the budget constraint is adjusted by scaling up the expenditures from the consumer's problem by the inverse of 5% and shopping and home production are scaled up by xx. The details of this extension are discussed in Appendix xx.

We solve the model as follows. For a candidate λ we can solve the first order conditions separately for each period. We take the implied expenditure for each *t* and iterate the budget constraint forward, starting from an initial asset position of zero. This process results in a period *T+1* level of assets. This process maps a (monotonic) function from λ to final period assets. We solve for the λ that is the zero of this function.

VI. Results

Expenditure

The simulated lifecycle results for our benchmark model are plotted in figures 10 through 14. Figure 10 plot expenditure over the lifecycle as well as expenditure per adult equivalent. The latter series is expenditure divided by the number of adults plus the number of children times their consumption weights ($\sum n(\tau)\alpha_1(\tau)$). We see that household consumption peads in middle age, with the young spending roughly xx% less than 45 year olds and retirees spending yy% less than 45 year olds. The shape and magnitude of the hump in expenditure tracks the data fairly well (c.f. figure 1).

On a per adult basis, expenditure remains humped-shape.and tracks the shape found in the data. The reason expenditure is higher in middle age even after controlling for family size stems from the fact that the middle aged face a higher opportunity cost of time. This implies that prices are highest during middle age, as discussed below. Secondly, in producing consumption goods through home production, middle aged households substitute away from time and increase the inputs of market goods. Thus, any given level of consumption will be produced with more expenditures during middle age.

Figure 11 plots prices paid over the lifecycle in our benchmark model. Prices peak around 40? at a level roughly x% higher than the young and y% higher than retirees. Again this corresponds closely to the data depicted in figure 5.

Time spent on shopping, home production and childcare is depicted in figure 12. The model produces the double peaks in shopping time we documented in the data, with the first peak at xx and the second peak in retirement. The magnitudes of the peaks are in line with the data as well. The first peak in shopping id driven by shopping needs, while the second peak is due to the low opportunity cost of time in retirement. This corresponds to the fact that prices are high in middle age despite the increased frequency of shopping, and prices are at their lowest in retirement and early in adulthood.

Time spent in home production mirrors time spent shopping, as both track changing family size and the cost of time. We also depict the ratio of time to goods in home production and see the clear substitution away from time and toward goods during middle age.

Labor supply in the model is hump shaped as in the data, peaking around xx when wages are highest. Asymmetry – human capital, physical changes with aging.

Finally, we plot the consumption of an adult in our model. The fact that the total cost of consumption (including time) is highest during middle age implies the agent consumes least at this time. This is a direct consequence of the nature of our utility function and the relative cost of time over the lifecycle. There is little direct evidence on consumption with which to compare this implication of our model. The fact that we have matched the inverse "U" in expenditure implies that our model is already consistent with observed expenditure. To assess the plausibility of this implication requires a direct measure of consumption. In an earlier paper (Aguiar and Hurst 2004), we constructed one such measure of consumption distinct from expenditure using food diaries. In that paper, we documented that utility from food consumption increases slightly in early adulthood, is relatively flat until retirement, and then increases during retirement. The model fits this pattern well later in life, but overstates the consumption of the young. This may call for an alternative utility function, such as habit formation, to alter the response of consumption to lifecycle movements in the price of consumption.

VII. Conclusion

References

Appendix A – Time Use

Appendix B – Extension of Model to Additional Goods

Figure 1: Lifecycle Profile of Total Food Expenditure and Total Food Expenditure Adjusted for Family Composition (Relative to 24-26 Year Olds)

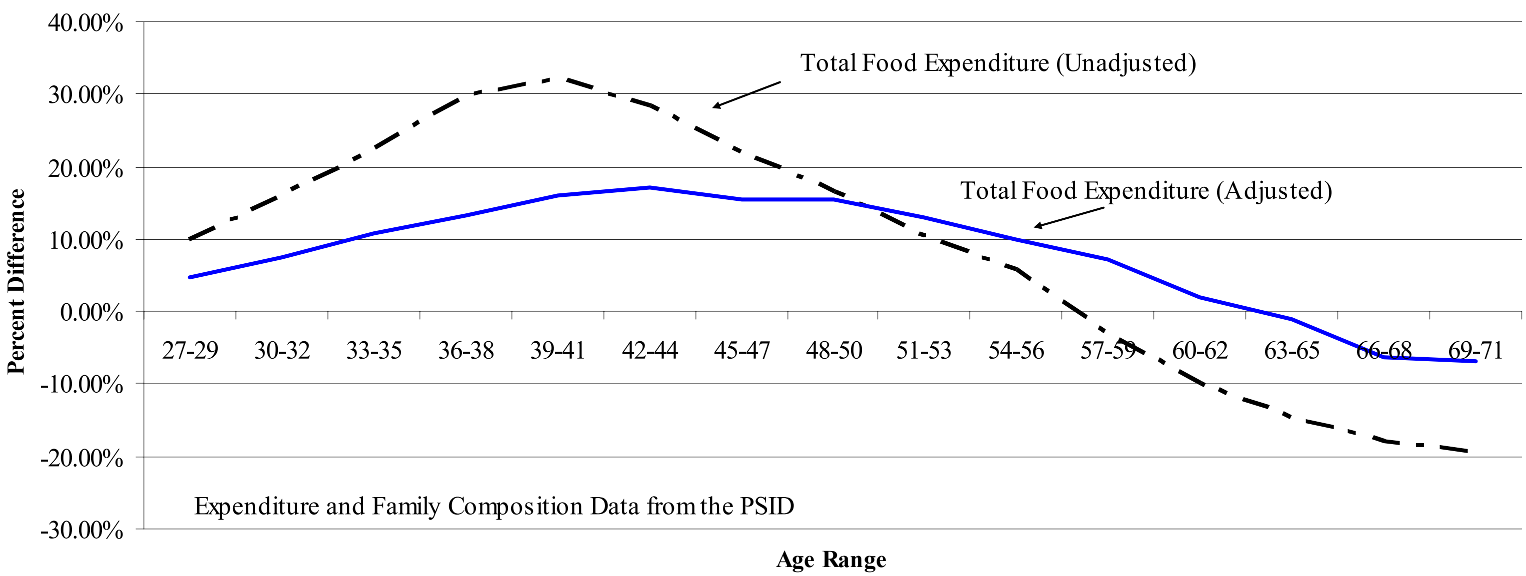

Figure 2: Price Paid by Income Categories

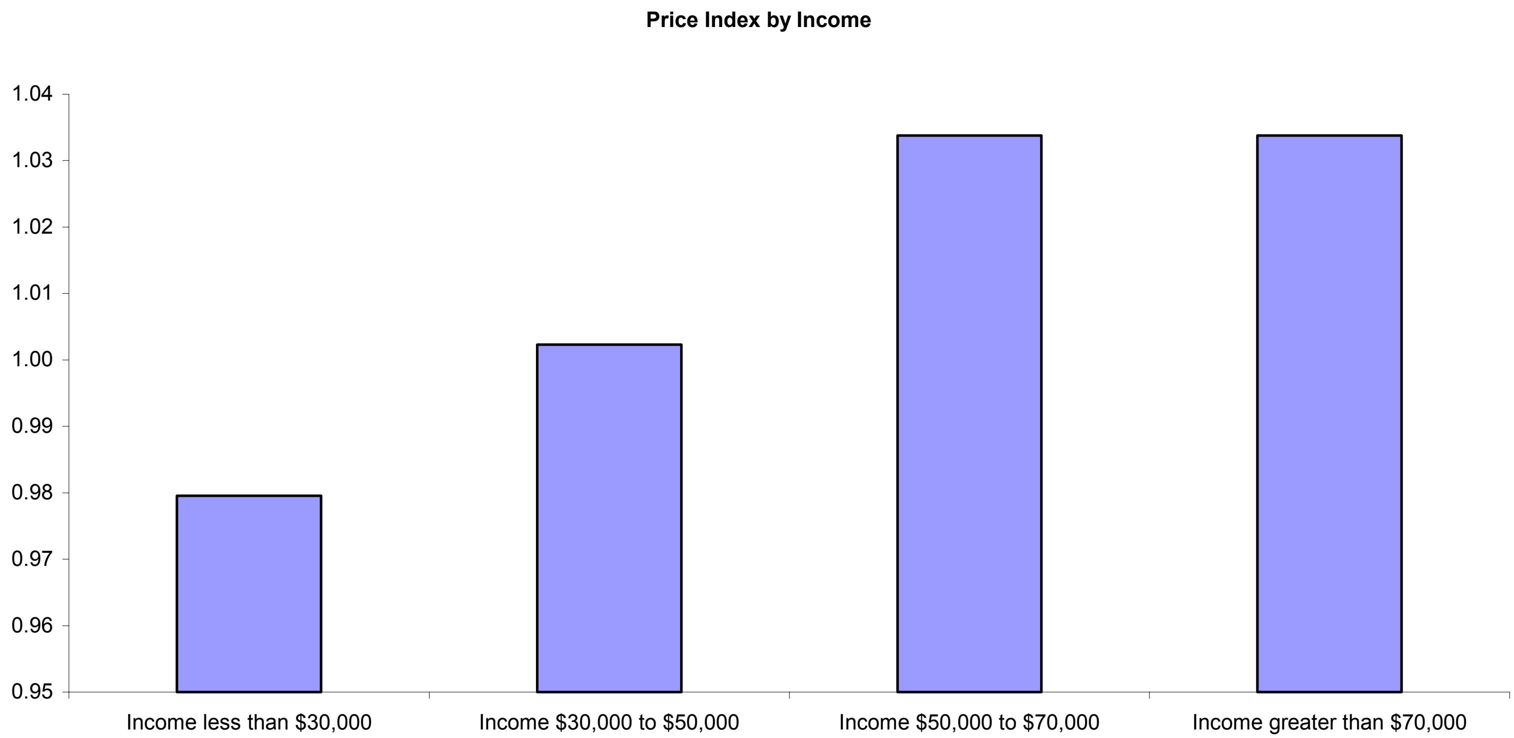

Figure 3: Price Paid by Household Size

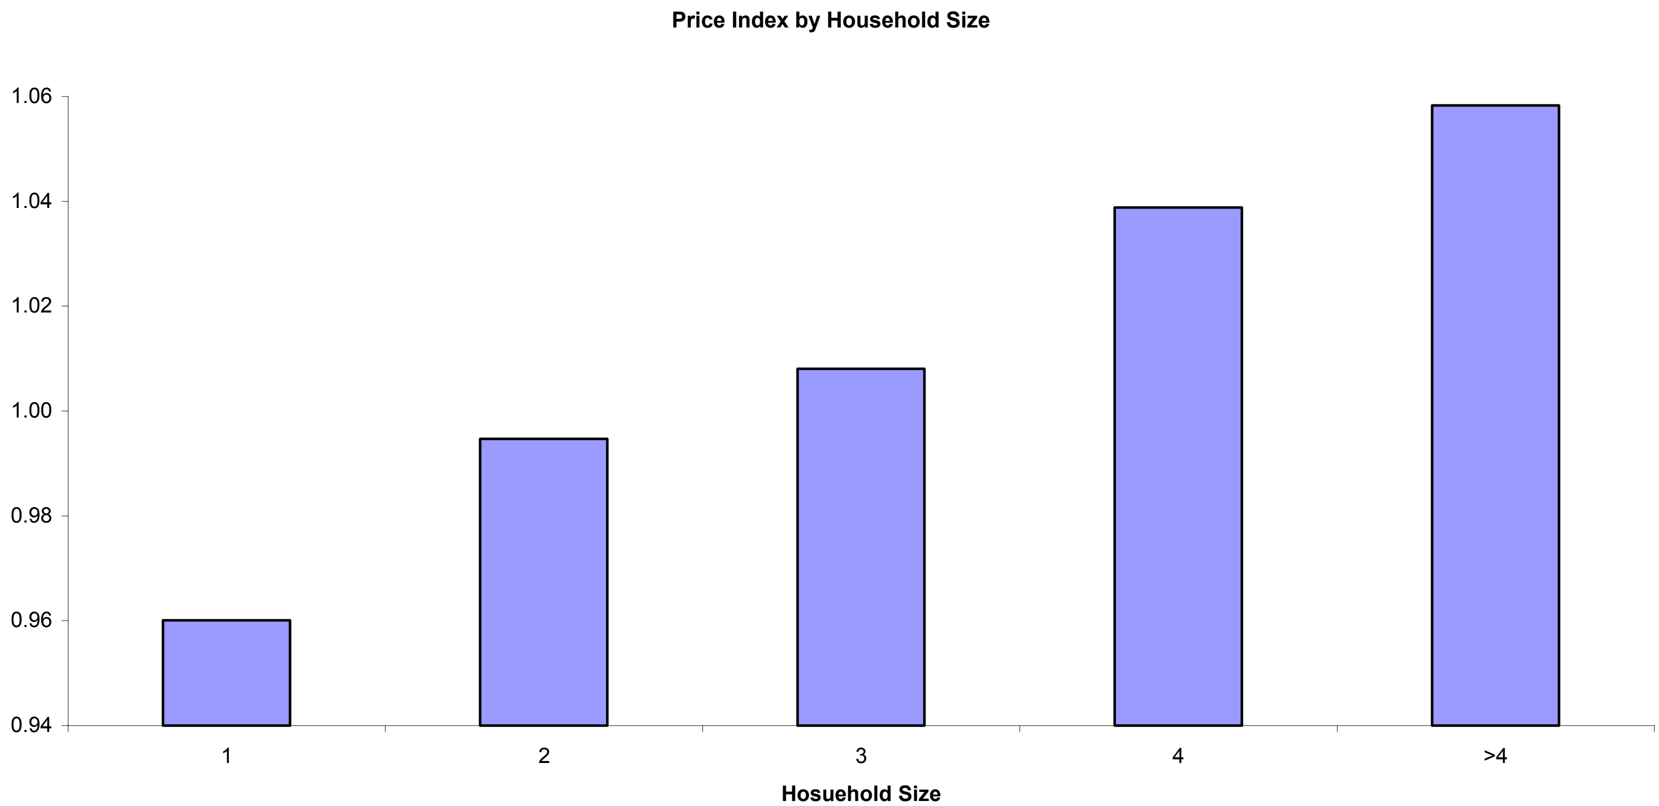

Figure 4: Price Paid by Household Composition

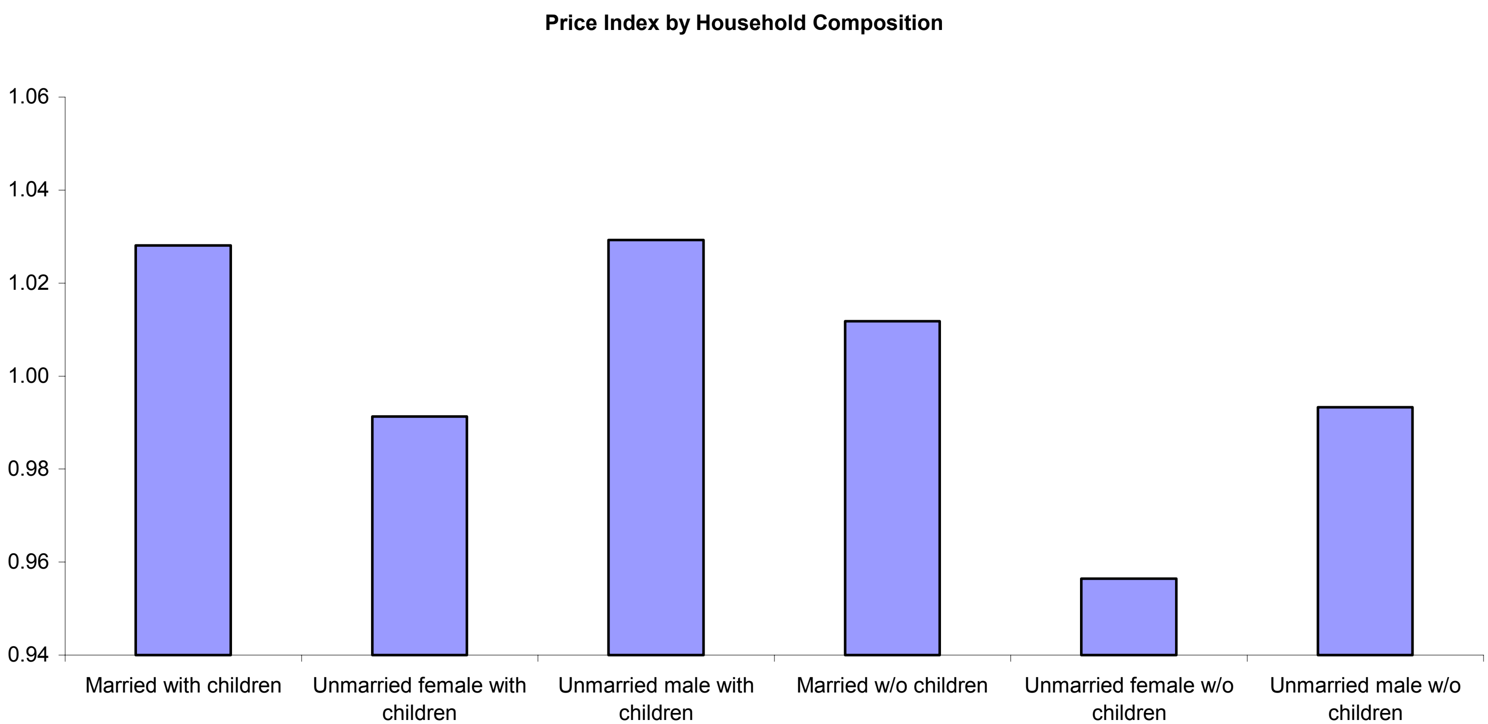

Figure 5: Price Paid by Age

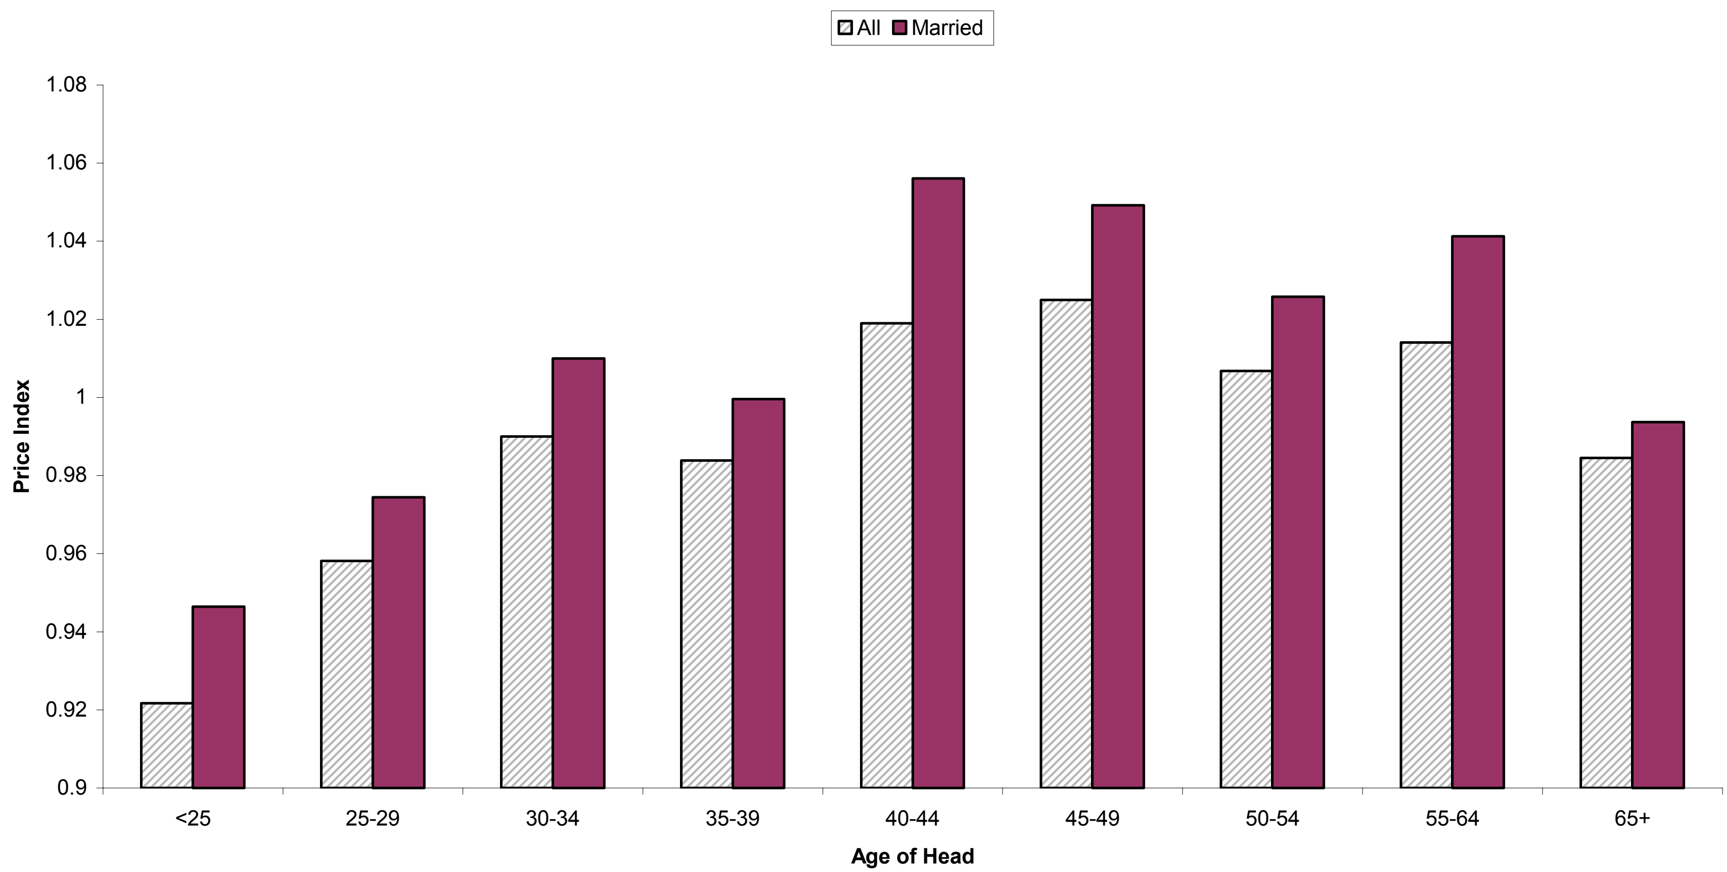

Figure 6: Within Age Price Dispersion

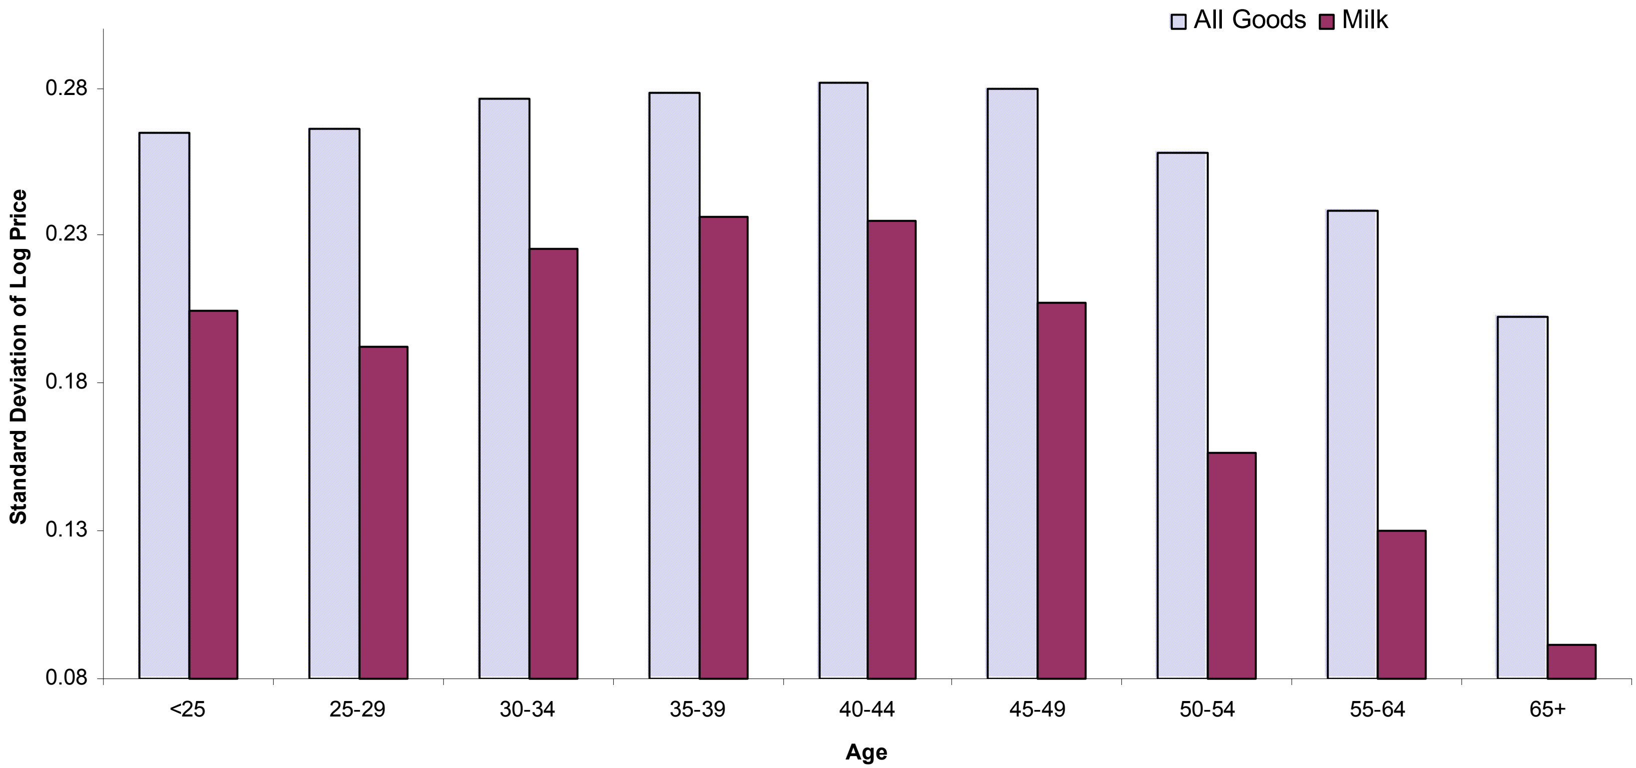

Note: For each good and month, we calculate the standard deviation of price paid within each age group (where an observation is a shopping trip). We then average across goods weighted by number of shopping trips. “Milk” is the within-age group price dispersion for that single good.

Figure 7a: Time Use over the Lifecycle

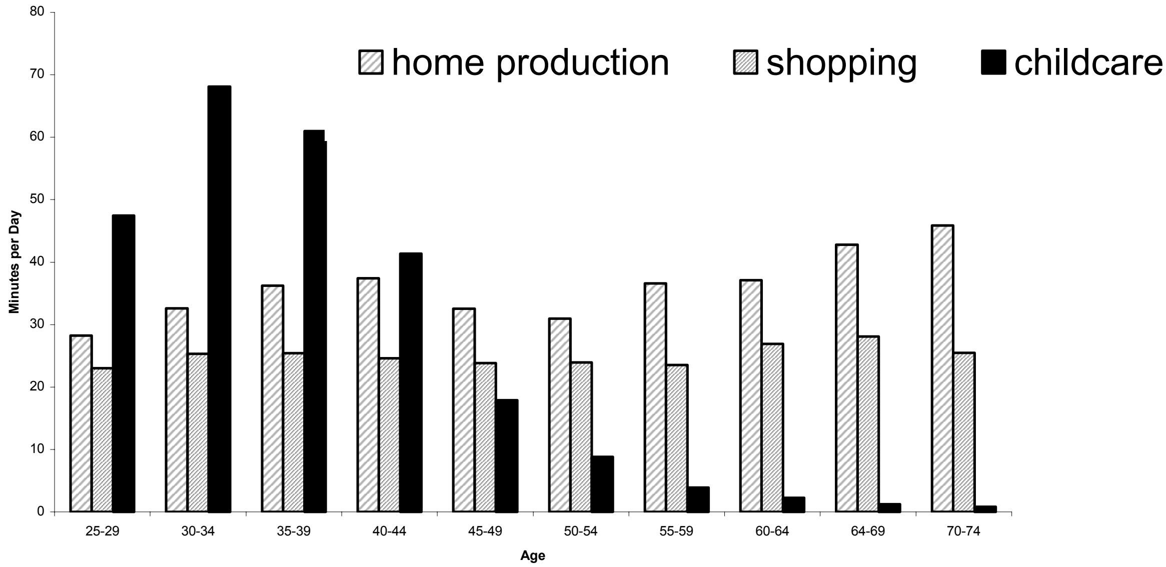

Figure 7b: Time Use – Married Households
(Sum of Married Men and Married Women)

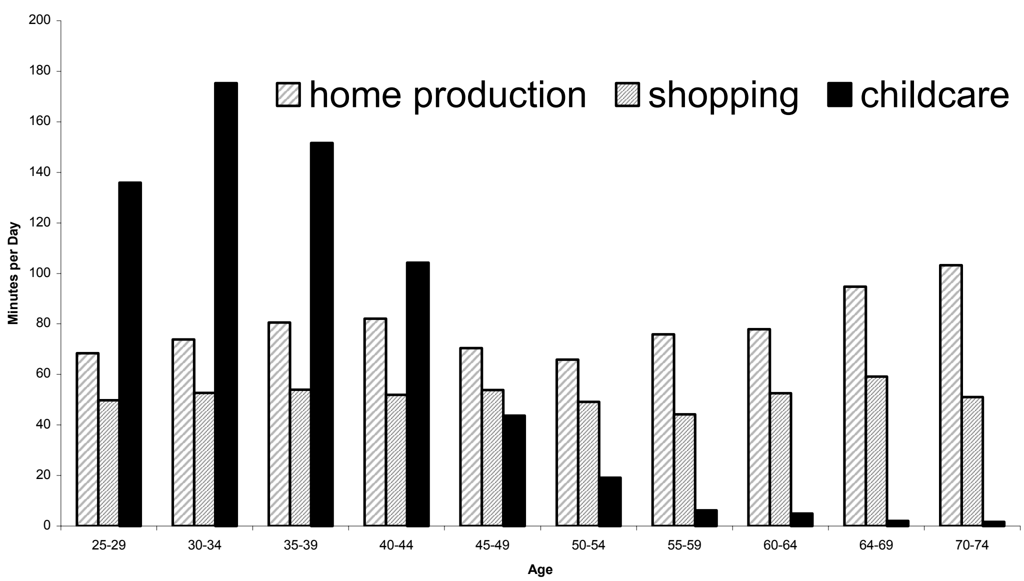

Figure 7c: Hours Worked

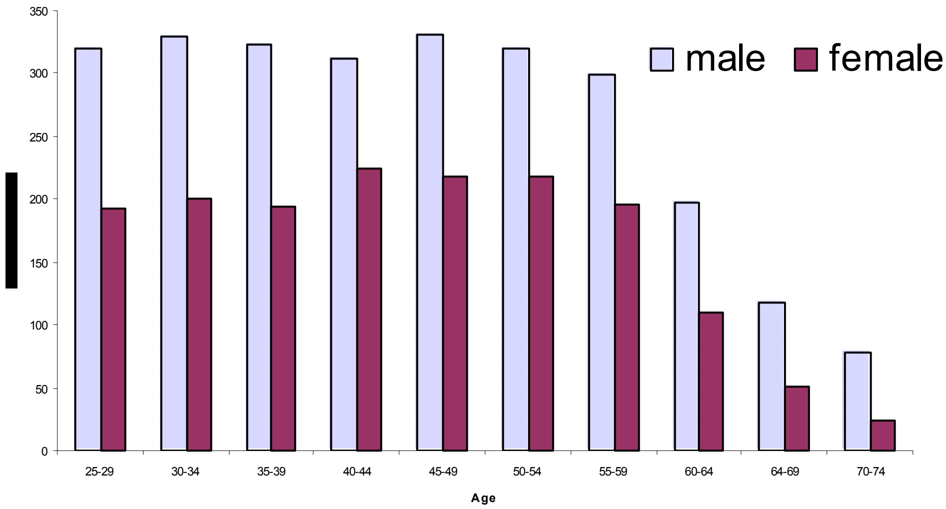

Figure 7d: Hours Worked – Married Households

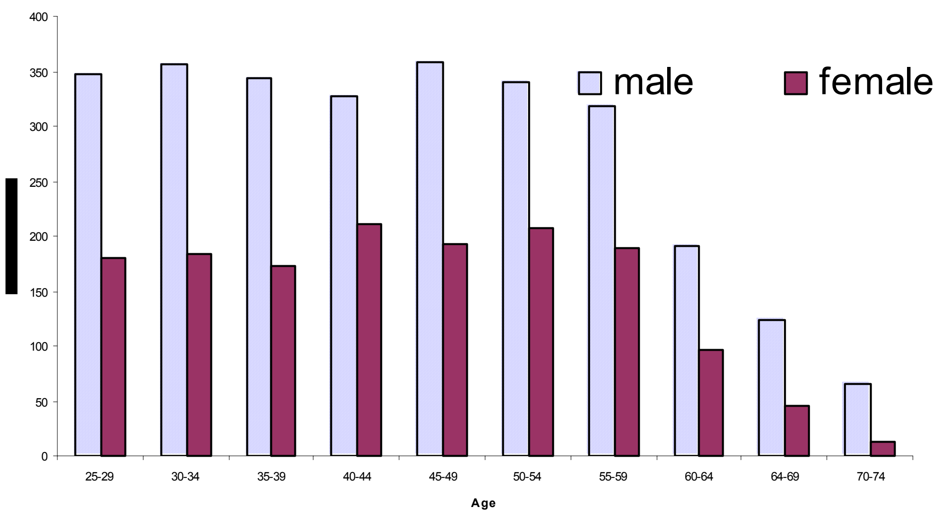

Figure 8: Number of Children over the Lifecycle

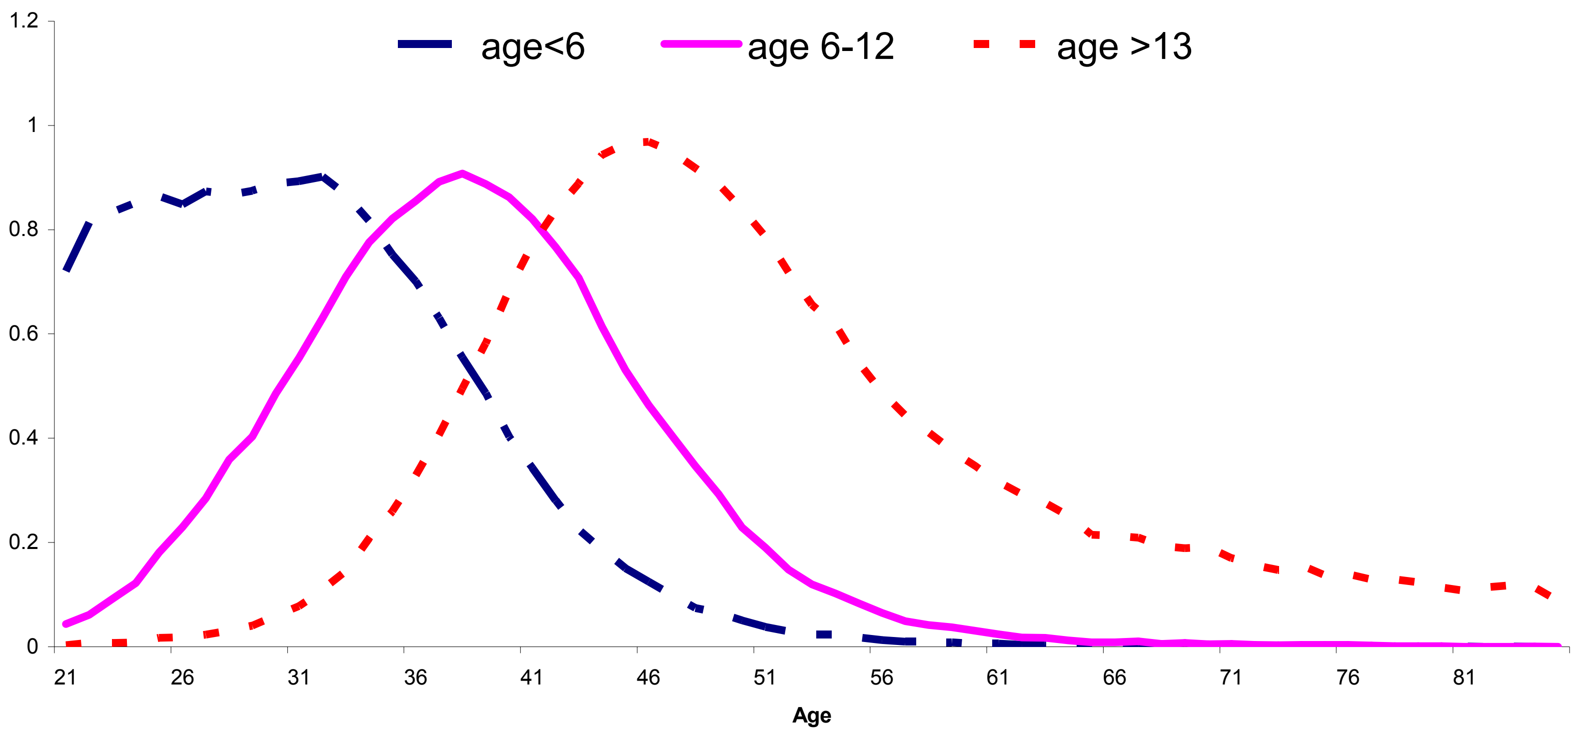

Figure 9: Wages over the Lifecycle for Married Households

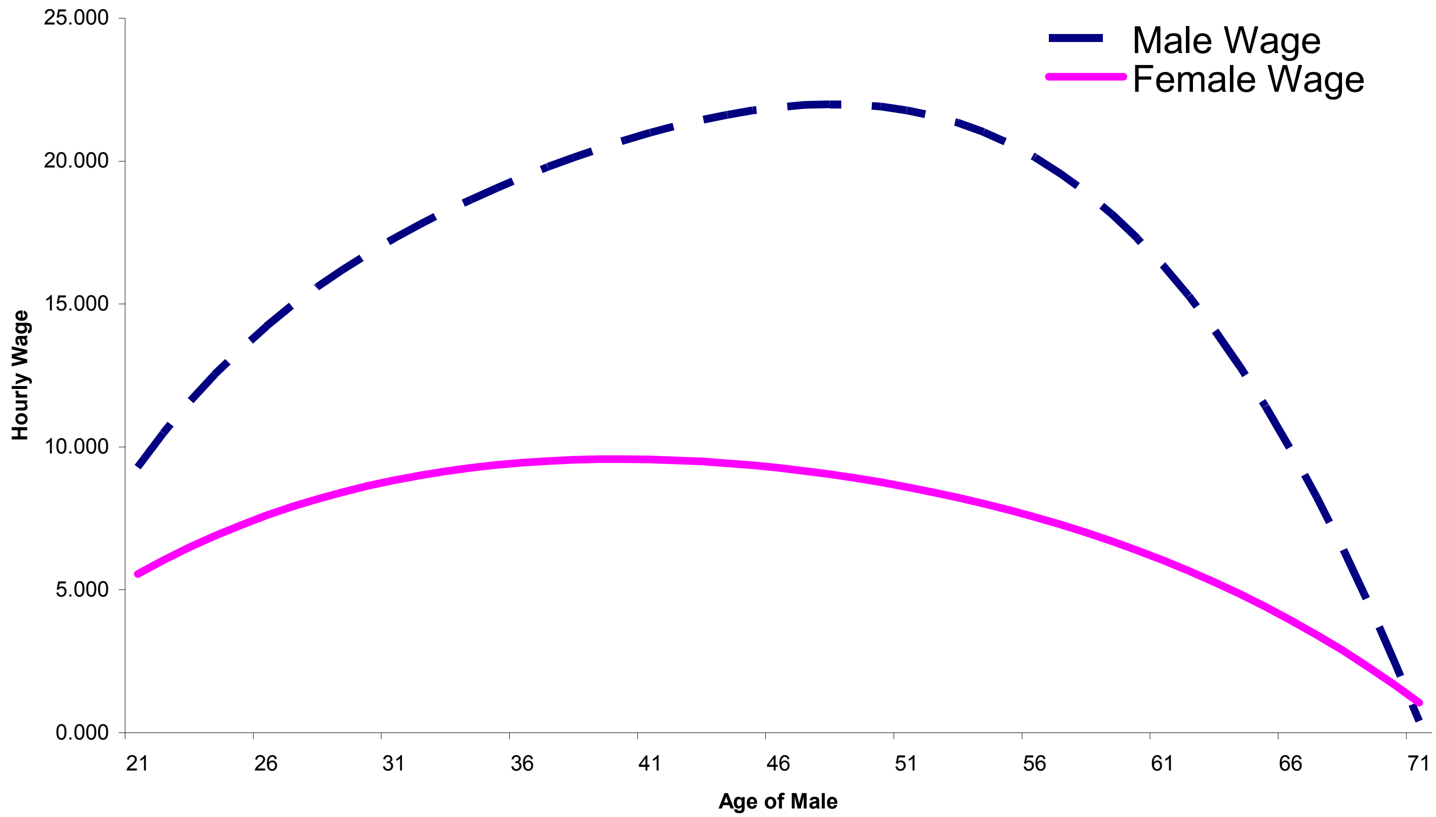

Note: PSID. Smoothed with a 5th order polynomial.

Figure 10: Predicted Expenditure over Lifecycle

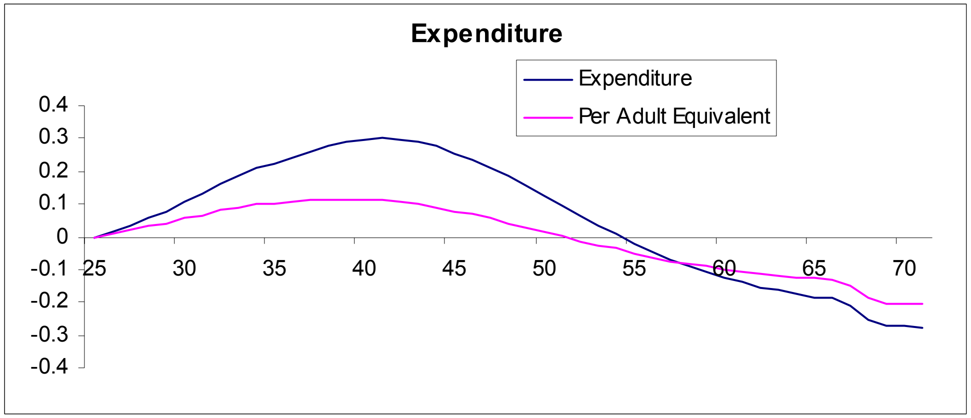

Figure 11: Predicted Price Paid over Lifecycle

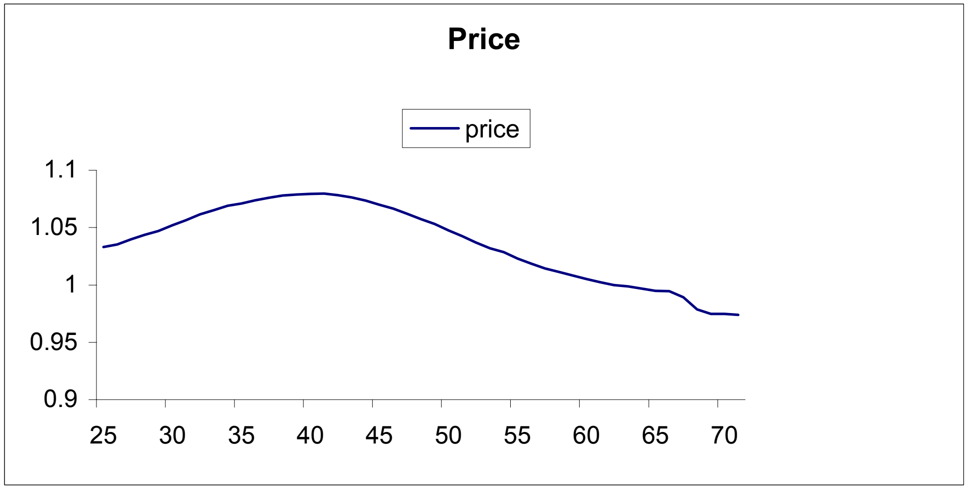

Figure 12a: Predicted Shopping and Home Production

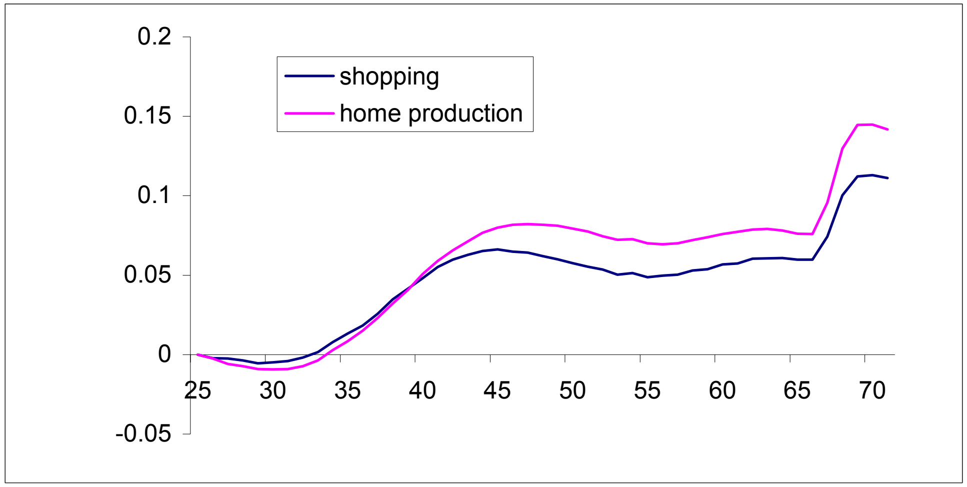

Figure 12b: Predicted Labor Supply

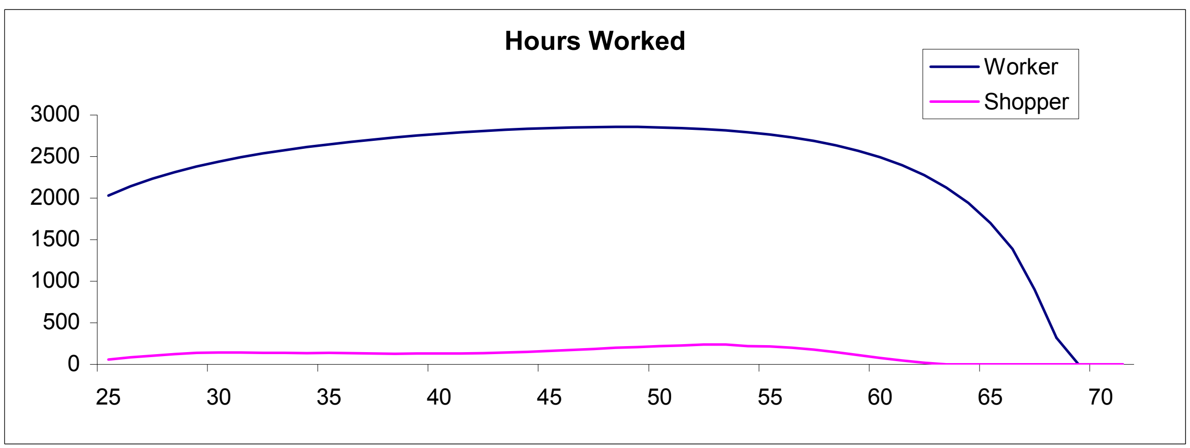

Figure 13: Opportunity Cost of Time

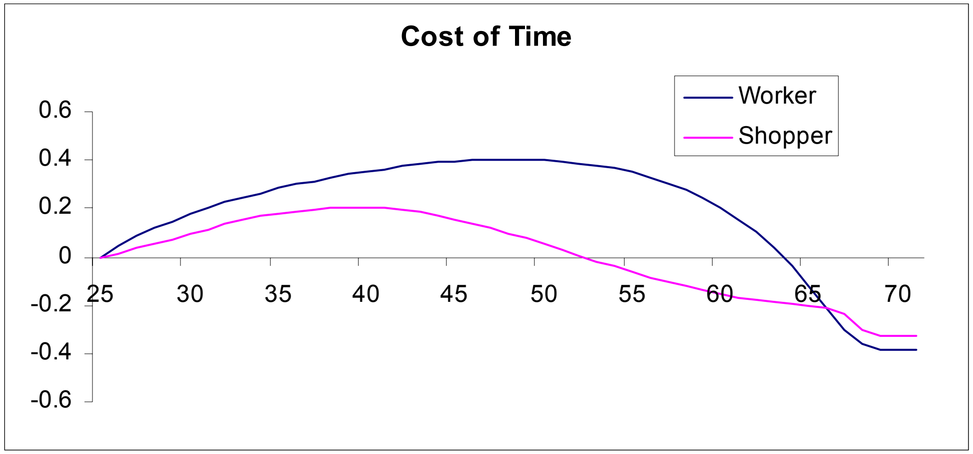

Figure 15: Implied Consumption

Table 1: Average Price Paid as a Function of Shopping Frequency and Total Quantity

	(1)	(2)	(3)	(4)
Dependent Variable:	p	p	ln(p)	ln(p)
ln(shopping frequency)	-0.10 (0.01)		-0.07 (0.01)	-0.18 (0.02)
Shopping frequency		-0.02 (0.005)		
(Shopping frequency)^2		$3x10^{-4}$ ($3x10^{-4}$)		
Elasticity with respect to shopping frequency (at average p ,frequency):	-0.10	-0.11	-0.07	-0.18
Additional Terms	$Q,\ldots,Q^5$	$Q,\ldots,Q^5$	ln(Q)	$Q,\ldots,Q^5$ (IV)
Elasticity with respect to Q (evaluated at average p, Q):	0.10	0.10	0.06	0.22
N	41,872	41,872	41,872	41,872
R-squared	0.07	0.07	0.04	NA

Notes – An observation is a household in a particular month. There were 2,065 total households. Robust standard errors clustered on household in parantheses. Last column instruments for $Q,\ldots,Q^5$ with age and household composition dummies.