

NBER WORKING PAPER SERIES

HOUSEHOLD TIME USE AMONG OLDER COUPLES:
EVIDENCE AND IMPLICATIONS FOR LABOR SUPPLY PARAMETERS

Richard Rogerson
Johanna Wallenius

Working Paper 24263
<http://www.nber.org/papers/w24263>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2018

Wallenius thanks the Knut and Alice Wallenberg Foundation for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w24263.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Richard Rogerson and Johanna Wallenius. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Household Time Use Among Older Couples: Evidence and Implications for Labor Supply
Parameters

Richard Rogerson and Johanna Wallenius

NBER Working Paper No. 24263

January 2018

JEL No. E24,J22

ABSTRACT

Using the Consumption Activities Mail Survey (CAMS) module in the HRS we document how time allocations change for individuals within a household when one or more members transitions from full time work to not working. Our basic finding is that the ratio of home production to leisure time is approximately constant for both family members. We then build a model of household labor supply to understand the implications of this finding for preferences and the home production function. We conclude that this fact suggests a relatively large elasticity of substitution between the leisure of the two members. For commonly used preference specifications, this also implies a large (i.e., greater than one) intertemporal elasticity of substitution for leisure.

Richard Rogerson
Woodrow Wilson School of
Public and International Affairs
323 Bendheim Hall
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

Johanna Wallenius
Department of Economics
Stockholm School of Economics
Box 6501
113 83 Stockholm
Sweden
johanna.wallenius@hhs.se

1. Introduction

Labor supply elasticity parameters are important for a variety of positive and normative issues. Although most labor is supplied by multimember households, most estimates of labor supply elasticity parameters have historically come from settings in which household labor supply is not jointly determined.¹ In this paper we study household choices over time use in a standard life cycle setting and derive a robust relationship that links differential changes in time use by household members to some key labor supply elasticity parameters. We then estimate this relationship using household level panel data on time use from the CAMS (Consumption Activities Mail Survey) module in the HRS (Health and Retirement Study). Our estimates suggest that households have a high willingness to substitute leisure across members and over time.

The starting point for our analysis is an examination of how household time allocation changes when one or more members of the household move from full time work to no work. This transition necessarily frees up a substantial amount of time that must be allocated among other uses, notably home production or leisure. Knowing how a household chooses to allocate this newly available time is revealing about important labor supply parameters. Most time use data sets, including the ATUS, are individual based and do not include a panel component, and so cannot directly speak to this issue. The CAMS module that we use provides

¹Classic references are MaCurdy (1981), Browning et al (1985) and Altonji (1986) for studies focusing on male labor supply, and Heckman and MaCurdy (1980) for female labor supply. More recently, Attanasio et al (2008) and Attanasio et al (2015) study female labor supply taking male income as exogenous.

information on time use for both members of a household and also contains a panel component. Using this data set we study how the household and individuals within the household allocate nonmarket discretionary time between leisure and home production.

We document four key facts. First, at any point in time there is substantial heterogeneity across households in the allocation of nonmarket discretionary time. Second, these differences are persistent over time. Third, dispersion across individuals is greater than dispersion across households, suggesting some substitutability of home production among household members. And fourth, relatively little happens to this allocation at either the household or individual level when one or more individuals in the household retire. That is, although total nonmarket discretionary time increases significantly, the share of this time devoted to home production changes very little. Although the CAMS data potentially contains a significant amount of measurement error, we confirm that a key feature we document in the CAMS is also present in the ATUS.

We then develop a structural model to help us interpret this finding. The model features a two person household that makes choices about market consumption and savings, as well as the time allocation of each individual between market work, home production and leisure. The choices are linked through a single budget equation as well as a home production function in which the two time inputs are imperfectly substitutable. We derive a relationship that links relative changes in home production time and leisure of the two members to two key elasticity parameters: the elasticity of substitution between the two members time in home

production and the elasticity of substitution between the two members leisure time in household preferences. This expression is essentially the first differenced version of the optimality condition requiring the marginal rate of substitution between member's leisure to be equal to the marginal rate of transformation of their two times in producing home production. For some common preference specifications, the elasticity of substitution between the two members' leisure times will also be the intertemporal elasticity of leisure for the household. Importantly, the expression we derive is robust to many model specifications and relies purely on optimal static tradeoffs within the household.

The expression that we derive allows a simple benchmark calculation. If household time allocations of nonmarket discretionary time do not change when one or more individuals transition from full time work to not working, then the above two elasticity parameters must be equal. Assuming that time inputs into home production are substitutes rather than complements implies a lower bound of unity for the leisure elasticity parameter. Knowles (2013) suggested that the elasticity of substitution between the two members time in home production is as high as 3. When we use our expression to interpret the modest changes in time allocation that we observe after a typical transition out of full time work, we find that the leisure elasticity is about two thirds as large as the home production elasticity. We also carry out an exercise of estimating the ratio of the two elasticities directly from the micro data rather than based on a typical experience involving a transition out of full time work. While measurement error precludes any strong conclusions, these results are not inconsistent with the above findings.

Lastly, we examine time use data from the MTUS (Multinational Time Use Survey) to examine whether the salient patterns observed in the US data also appear in the data for other countries. These data do not have a panel component and so we cannot replicate our analysis using the CAMS data. But we can ask whether there are large changes in the average allocation of nonmarket discretionary time over the age range where market work decreases dramatically due to retirement. While there is some variation across countries, with some countries experiencing changes that are somewhat larger than we found for the US, the average response is quite similar to what we found in the US data.

Our paper relates to several strands of the literature. By focusing on changes in time use at retirement it relates to the analysis of Aguiar and Hurst (2005), though their focus was on distinguishing changes in consumption from changes in consumption expenditure. By considering how a decrease in market work is allocated to leisure and home production it is related to Aguiar et al (2013), though their focus was the large decrease in market work during the Great Recession. By using information about choices at retirement to infer preference parameters we follow Laitner and Silverman (2005), though their focus was on changes in consumption rather than changes in time use. Our analysis is most related to Rogerson and Wallenius (2016), who use the ATUS to study changes in time allocation at retirement to infer preference parameters. That analysis did not have panel data and considered individuals rather than households. Additionally, the expression used to infer preference parameters was based on a dynamic first order condition, whereas the current analysis only requires static first order conditions

to hold.

There is an extensive literature on various aspects of household labor supply, one that is too large to reference. By providing evidence on the substitutability of leisure between household members our paper relates to the subliteration that studies how households respond to shocks. A notable recent contribution to this literature is Blundell et al (2016), which relate to the earlier literature on the so-called “added worker effect”. See, for example, the papers by Lundberg (1985) and Cullen and Gruber (2000).

A brief outline of the paper follows. The next section describes the CAMS dataset that we will use for our analysis. Sections 3 and 4 report the key findings regarding how time use changes during a transition out of full time work, both at the individual level as well as the household level. Section 5 presents the model that we will use to interpret the salient patterns found in the data, and Section 6 reports the implications for the two key elasticity parameters. Section 7 reports the patterns we find in the MTUS and Section 8 concludes.

2. A Household Panel Data Set on Time Use

The American Time Use Survey (ATUS) is widely regarded as the highest quality data on individual time use. Nonetheless, it has two key limitations. First, because it is an individual based survey and not a household based survey, it does not provide information on the home production time of other household members. Given the potential for substitution among household members, the lack of household data is potentially critical. Second, it does not have a panel

component. In the presence of individual level heterogeneity, deriving inferences from a pure cross-section can be very difficult.

In this section we describe an alternative data set that does provide panel data on household time use—the Consumption and Activities Mail Survey (CAMS). The following two sections will document some key patterns in the changes in time use at both the individual and household level during the retirement process.

2.1. The CAMS Dataset

The Consumption and Activities Mail Survey (CAMS) is a module sent to a subset of participants from the Health and Retirement Study (HRS). The HRS is a large nationally representative panel survey of individuals in the US aged 50 and older, administered every second year, starting in 1990. The CAMS module was added in 2003 and is also administered every other year. Importantly, the HRS is a household survey, i.e., it obtains information for both spouses in the case of married individuals who are living together. This feature is also true for the CAMS module, but only starting in 2005. For this reason we restrict our analysis to the data in the CAMS modules for the years 2005, 2007, 2009, 2011 and 2013.

The CAMS module provides information on both time use and consumption spending. We focus solely on the time use component. In contrast to the ATUS and other time use surveys that rely on a diary method and have individuals detail all of their activities over a single day, the CAMS asks people to recall how much time was allocated to a set of activities over the previous week. For a subset of activities which are thought to be performed on a more occasional basis the

survey asks about time allocated to them in the last month. For a more extensive discussion of the CAMS data as well as a comparison with the ATUS along several dimensions we refer the reader to Hurd and Rohwedder (2007).

Our analysis will focus on a simple statistic: the ratio of home production time to the sum of home production plus leisure time. That is, we examine how discretionary time not allocated to market work is allocated between home production and leisure. In standard models that abstract from home production all of this time is viewed as leisure time, and explicitly modelling home production amounts to adding this dimension to the time allocation problem. In what follows we use the letter z to refer to the ratio of nonmarket discretionary time devoted to home production, i.e.,

$$z_{jt} = \frac{h_{jt}}{l_{jt} + h_{jt}}$$

where h_{jt} and l_{jt} are home production time and leisure time respectively, for individual j at date t .

While this ratio is perhaps a natural statistic of interest in models with home production, our choice is also motivated by a feature of the CAMS data set. In particular, whereas a time diary survey allocates each interval over the course of a day to some activity and so will necessarily have total time allocated equal to the total time available, this is not necessarily true in the CAMS, which is based on recall. The extent to which total time is underreported varies across households at a point in time and for a given household across time. If we were to focus on levels of leisure and home production time, cross-sectional and time

variation would be driven both by the extent of differences in underreporting and true differences in time use. For now, we note that the ratio observed in the data would be an unbiased estimate of the true ratio if the extent of underreporting is proportional for the two categories. Importantly, this is true even if the extent of underreporting varies across individuals. When we turn to formal estimation later in the paper we will show that we can accommodate differential underreporting across categories.

Although the CAMS module is a household survey, our first set of results focuses on the time allocations of individuals, and so does not make use of the household feature of the data. This provides an opportunity to assess whether the key patterns we document also appear in the ATUS, which does not have the measurement issue noted above.

Our definition of leisure time includes the following categories in the CAMS data: watching TV, reading papers & magazines, reading books, listening to music, walking, sports/exercise, visiting friends/neighbors/relatives, communicating by phone/letter/email with friends/neighbors/relatives, playing cards or games, attending concerts/movies/museums, singing or playing instrument, doing arts and crafts, eating out. We have also studied a specification in which we classify sleep in excess of 49 hours per week as leisure, but as the findings were effectively unchanged we do not report any results for this case. Our definition of home production time includes the following categories: house cleaning, laundry, yard work & gardening, shopping and errands, meal prep and clean up, caring for pets, managing household finances, home repairs, and vehicle maintenance. In terms of

cleaning the data we drop observations if homework is greater than 100 hours per week, if market work is greater than 100 hours per week, or if any of homework, market work or leisure is missing.² We apply respondent weights for the individual level analysis and household weights for the household level analysis, and in all cases use the weights from the initial year when looking at individuals or households over time. Note that because the CAMS is a subsample of the HRS and is conducted in between the main HRS surveys, the weights for the 2005 CAMS are the 2004 weights from the HRS. After cleaning the data and matching individuals across consecutive waves we have 4668, 4581, 4266, and 5043 observations for the 2005-2007, 2007-2009, 2009-2011 and 2011-2013 pairs of waves.

We will be particularly interested in how time allocations change when one or more members of a household moves from full time work to not working, since this is a case where we know there are large changes in the total time being allocated between leisure and home production, thereby increasing the signal to noise ratio. For ease of exposition we will refer to such transitions as “retirement”, though of course they need not represent retirement since they could simply reflect temporary fluctuations in work due to a variety of factors. Importantly, from the perspective of the model that we will use to interpret the data, this is not an issue. What matters most for our purposes is to isolate a situation in which we think there is a high signal to noise ratio in terms of changes in home production time.

Nonetheless, we will consider two different criteria for “retirement”. The first criterion will identify someone as working full time if they report working at

²If a particular subcategory of home production or leisure is missing, we simply replace it by a zero. But if all subcategories are missing, or leisure is zero, we delete the individual.

least 35 hours per week in the initial wave, and identify them as being retired in the subsequent wave if they report working no more than 5 hours per week. The second criterion will examine market work over three consecutive waves and require that the individual work at least 35 hours per week in the initial wave, and then no more than 5 hours per week in *each* of the next two waves. When applying this criterion we lose any retirements occurring between the last two waves since we cannot check whether the decrease in market work is persistent. Using the first criterion, the number of retiring individuals in each of the pairs of surveys is 204, 188, 158, and 215. Using the second criterion the corresponding numbers of retiring individuals are 136, 126 and 106 for the 2005-2007-2009, 2007-2009-2011 and 2009-2011-2013 periods, so that about two-thirds of the retirements under criterion 1 also satisfy the more stringent condition for criterion 2.

3. Patterns for Individuals

In this section we document some key patterns for changes in the allocation of nonmarket discretionary time at the individual level. We first report patterns found in the entire sample, and then consider the sample consisting of individuals who transition to retirement.

3.1. Patterns in the Overall Sample

We begin by examining what happens to the variable z across consecutive surveys for individuals that we can match across consecutive pairs of surveys. Results are in Table 1, presented separately for males and females.

Table 1

Value of z for Matched Individuals by Gender

	Males					Females				
	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$
2005-07	.25	.24	.15	.14	.46	.30	.30	.14	.15	.48
2007-09	.25	.25	.14	.14	.53	.30	.30	.15	.15	.52
2009-11	.24	.24	.14	.14	.55	.30	.29	.15	.15	.54
2011-13	.25	.24	.14	.14	.51	.30	.29	.14	.15	.49

Several patterns emerge. Because they are so similar for males and females, here we focus on the results for males. First, the average value of z is remarkably stable over time, both for a fixed group of males from one survey to the next (i.e., going from the first to the second column) and samples (i.e., moving down the rows of either the first or second column). Second, there is substantial dispersion of this ratio in the population, with a coefficient of variation equal to roughly .60 for males. Keeping in mind that the data on time use is essentially for one week and that time amounts based on recall are expected to be noisy, one might suspect that a large part of the dispersion simply reflects a combination of measurement error and sampling variation. However, the fourth column of the table shows that the correlation of these ratios at the individual level two years apart is strongly positive, suggesting that a substantial amount of the dispersion reflects true dispersion, and is persistent. The one difference between males and females in Table 1 is that the value of z is higher for females by about .05, though the standard deviations are effectively identical.

As a way to gauge the amount of true dispersion consider the following simple but suggestive calculation. Assume that the true value of z for individual i , denoted by z_i is constant over time and that the distribution across individuals is distributed normally with mean μ_z and variance σ_z^2 . Assume also that measurement error is iid over time and across individuals, and is normally distributed with mean 0 and variance σ_ε^2 , so that the observed data on z_{it} , denoted by z_{it}^d are given by³:

$$z_{it}^d = z_i + \varepsilon_{it}$$

Using data on the mean and standard deviation of z_{it}^d and the correlation between z_{it}^d and z_{it+1}^d allows us to estimate all three parameters. Given that the correlation between z_{it}^d and z_{it+1}^d is around 0.50, the assumed structure implies that σ_z^2 and σ_ε^2 are approximately equal, in which case the true standard deviation of z in the population is around .10. This still implies a very substantial degree of dispersion in the value of z across individuals, with a 90 – 10 ratio in excess of 2.

To provide some further information about the structure of the variation over time, we construct a balanced panel consisting of all individuals who can be matched across all five surveys. We then examine the cross-correlations of z across the five surveys, separately for males and females. Consistent with their being an important temporary component (which includes measurement error), we find that the correlation of consecutive first differences is strongly negative, in

³The term ε_{it} may also include temporary variation in z in addition to measurement error. In this case we are estimating the variance in the permanent component.

the range of $-.40$ to $-.45$. But for first differences that are not consecutive, the correlation is very close to zero.

3.2. Patterns in the Retiree Subsample

We now examine the behavior of z for retiring individuals. Here we present the results based on the first criterion, which only requires a transition from more than 35 hours per week to no more than 5 hours per week across consecutive surveys. The summary statistics for this group are presented in Table 2, once again presented separately for males and females.

Table 2
Value of z for Individuals Retiring Across Surveys

	Males					Females				
	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$
2005-07	.26	.27	.14	.15	.37	.29	.29	.14	.13	.23
2007-09	.24	.27	.13	.14	.47	.30	.31	.14	.15	.33
2009-11	.25	.28	.14	.15	.60	.33	.33	.15	.14	.44
2011-13	.23	.25	.14	.13	.30	.30	.31	.14	.16	.34

Remarkably, the same basic patterns found in Table 1 for the overall sample also appear when restricting attention to retiring individuals. In particular, both the mean and standard deviation of z change very little as the sample moves from working full time to retirement; there is a modest increase in the mean of z for males, and an even smaller increase for females. The standard deviation seems effectively unchanged by retirement. Notice also that the statistics for individuals

who are about to retire are almost identical to the population averages in the sample. Lastly, it remains the case that z at the individual level is positively correlated between the two periods, though the correlation is somewhat lower than for the overall population.

The picture that emerges is that very little seems to happen to the value of the z when an individual moves from working full time to retirement. It is important to emphasize that this transition necessarily involves a sharp decrease in the amount of time devoted to market work and so also involves a substantial increase in the amount of discretionary time that individuals allocate between leisure and home production. The fact that z is roughly constant does not imply that there is no change in overall time allocation; rather, it simply implies that time spent in leisure and home production increase proportionately.

To document these patterns more formally we pool the data from all of the surveys and run a panel regression of the following form:

$$z_{it} = \bar{z}_i + \beta I_{Rt} + \varepsilon_{it} \tag{3.1}$$

where \bar{z}_i is an individual fixed effect, and I_{Rt} is an indicator function which takes on the value of 1 if individual i satisfies our criterion for being retired in period t . We run this specification for samples constructed using each of our two criteria of retirement. Specifically, in the first sample, we consider all of the individuals in the sample used to generate Table 2, i.e., all consecutive pairs of observations for an individual that moves from at least 35 hours per week in the initial period to five or fewer hours in the second period. For males the estimated value of β

is .022 with a standard error of .009, and for females the estimated value of β is .004 with a standard error of .011. Consistent with our summary of results in Table 2, these estimates suggest a modest increase in z for retiring men and no statistically significant change for retiring women.

The second sample instead focuses on individuals who meet our second criterion for retirement. That is, we consider those cases in which an individual has three consecutive observations, works at least 35 in the first survey, and five or fewer in each of the next two surveys. For this criterion, the results are basically the same; the estimated value of β for males is .029 with a standard error of .010, while for females it is .003 with a standard error of .013.

3.3. Patterns for Individuals in the ATUS

As noted previously, time use measures derived from surveys that rely on time diaries may reasonably be viewed to be more reliable than those that rely on recall. This argument would suggest that patterns found using the ATUS are more reliable than patterns found using the CAMS. Because the ATUS does not contain a panel component, we cannot replicate the above analysis; specifically, we cannot assess the extent to which z changes when an individual moves from full time work to retirement using the ATUS. However, in this subsection we argue that patterns found using the ATUS strongly support the key patterns we have highlighted in the CAMS data. We view this as evidence in favor of taking the patterns for the behavior of z at retirement in the CAMS data at face value.

As noted, the ATUS does not have a panel component and so we cannot follow

what happens to z at the individual level when an individual retires from full time work. However, we can use the ATUS to create a synthetic panel. With this in mind, Table 3 shows the behavior of mean market work (μ_m), mean z (μ_z) and the standard deviation of z (σ_z) by age using pooled data from the ATUS samples for the years 2003-2015, both in aggregate as well as separately for men and women.

Table 3

Market Work (m) and z in the ATUS

	Total			Male			Female		
Age	μ_m	μ_z	σ_z	μ_m	μ_z	σ_z	μ_m	μ_z	σ_z
55	32.1	.34	.25	36.8	.28	.25	27.8	.39	.24
56	29.4	.32	.25	33.5	.27	.24	25.5	.38	.24
57	29.2	.33	.25	32.5	.28	.25	25.9	.39	.24
58	26.9	.33	.25	33.0	.27	.24	21.3	.39	.24
59	25.6	.34	.25	29.1	.27	.24	22.3	.39	.25
60	25.7	.33	.25	31.0	.27	.23	20.8	.40	.25
61	22.4	.32	.24	26.2	.27	.24	18.8	.38	.23
62	20.6	.33	.24	24.4	.28	.24	17.3	.38	.24
63	16.1	.33	.24	17.8	.28	.24	14.5	.38	.23
64	15.4	.33	.24	18.2	.28	.24	12.9	.37	.24
65	11.7	.33	.24	13.7	.28	.23	9.9	.37	.23
66	10.8	.33	.23	12.1	.28	.24	9.6	.38	.22
67	10.1	.31	.23	12.7	.27	.22	7.9	.36	.22
68	9.8	.32	.24	11.3	.27	.24	8.5	.35	.24
69	7.0	.31	.23	9.3	.26	.22	5.1	.36	.23
70	7.0	.30	.24	9.2	.23	.22	4.9	.36	.23

Notably, Table 3 shows that mean market work decreases dramatically with age, especially between the ages of 60 and 67. Rogerson and Wallenius (2016) show that the dominant source of this decrease in market hours is the movement

of individuals from full time work to retirement. It follows that examining the changes in z with age are effectively providing information on the changes in z associated with retirement. Interestingly, we see the same two features in the synthetic cohort constructed from the ATUS that we saw in the panel component of the CAMS: both the mean and standard deviation of z are virtually constant in the face of the dramatic decrease in time devoted to market work as individuals leave full time work. This is true both in the aggregate as well as for each gender considered separately.

Although the ATUS data possesses the same qualitative properties found using the CAMS data, we note two quantitative differences. First, mean z is higher in the ATUS than in the CAMS. As noted previously, the two surveys use very different methods, and these statistics suggest that there is some systematic differences in levels of home production and leisure in the two surveys. Second, the standard deviation of z is higher in the ATUS. This difference is to be expected, at least at a qualitative level. The reason is that the unit of observation in the ATUS is one person for a particular day of the week. It follows that at least part of the standard deviation reflects variation across days of the week. In contrast, the unit of observation in the CAMS module is one person for a week (and to some degree the month). It follows that dispersion due to variation across days of the week is implicitly removed in the CAMS, leading one to expect a smaller standard deviation. Of course, to the extent that measurement error is larger in the CAMS there is also a factor leading to the opposite pattern, though as noted previously, one could have substantial measurement error in each of leisure and

home production time without necessarily having a lot of measurement error in z .

The ATUS cannot speak to all of the patterns that we found using the CAMS. In particular, the ATUS cannot tell us if the near constancy of the mean of z reflects persistence for a given individual over time as opposed to simply a constant distribution over time with individuals moving within the distribution. Our analysis using the CAMS data found evidence for the former. But the key message that we take away from our analysis of the ATUS is that despite some concerns with data quality in the CAMS, the patterns we have documented and will make use of going forward appear to be robust.

4. Patterns for Households

In this section we proceed to use both the panel and household features of the CAMS to examine what happens to household time allocation when one or more members of the household retire. For this analysis we use the same criterion as before applied to the household unit. That is, we only include data for two member households and we require that both individuals satisfy our criterion in both periods. Some individuals are removed from the sample because they are not part of a two member household, and others are removed because their partner has missing observations. The resulting sample of matched two member households contains 1395, 1356, 1218, and 1409 observations for the 2005-2007, 2007-2009, 2009-2011 and 2011-2013 pairs of waves. As before, we will also focus on households in which at least one member experiences a move from full time

work to retirement.

We begin by documenting some properties of the household’s aggregate time allocation, i.e., the ratio of total household home production time to the sum of total household home production time plus total household leisure time. Results are presented in Table 4.

Table 4
 z for Matched Households

	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$
2005 – 07	.29	.29	.11	.11	.52
2007 – 09	.29	.28	.11	.11	.56
2009 – 11	.28	.28	.11	.11	.55
2011 – 13	.28	.27	.10	.11	.47

Perhaps not surprisingly, the key finding in this table is that the same patterns found in the individual level data are also present at the household level, i.e., both the mean and standard deviation are unchanged across surveys and the level of z at the household level is highly positively correlated across surveys. Note that the standard deviation at the household level is about one third smaller than at the individual level, suggesting that a significant part of the variation found in the individual data is across individuals within households.⁴ To the extent that the time of different members are substitutes in household production and that there is some comparative advantage for market versus home work across household members, this is to be expected. Nonetheless, the data still indicate a very

⁴It is also the case that if measurement error is iid across household members, the variance of household level measurement error will be smaller in a two member household.

significant degree of dispersion in the level of the z across households. Repeating the same simple calculation as earlier to estimate the part of the dispersion that is not due to measurement error implies a standard deviation of around .08.

Next we consider in more detail what happens inside the household when one or both members retires. Five different cases are possible. One case is when both members move from full time work to retired. The other cases involve one member retiring but conditioned on whether the other member is retired or working full time.⁵ As we cut the sample of retirees into finer categories the sample sizes tend to become somewhat small, so in what follows we pool the observations across the four pairs of consecutive waves.

Table 5
 z for Households With a Retiring Member

	Male		Female		Gap		#obs
	Pre	Post	Pre	Post	Pre	Post	
$M \rightarrow R, F \rightarrow R$.233	.257	.325	.346	.092	.088	37
$M \rightarrow R, FW$.256	.275	.319	.310	.063	.035	78
$M \rightarrow R, FR$.216	.229	.331	.329	.115	.101	115
$F \rightarrow R, MW$.221	.191	.304	.303	.082	.113	47
$F \rightarrow R, MR$.288	.263	.310	.328	.022	.065	62

Several patterns are present. When the male retires there is a modest increase in his z . Similarly, when the female retires, there is a modest increase in her z in

⁵There are also cases in which one or both members are working an intermediate number of hours in the market, i.e., between 5 and 35 hours per week. These observations are excluded from the table.

two of the three cases, with effectively no change in the third. When the female retires and the status of the male is unchanged, there is a modest decrease in the value of z for the male. A similar pattern is found for the female z when the male retires and the status of the female is unchanged, though the decrease is even more modest. The mean value of z is greater for females than males in all cases. This gap decreases when the male member retires and increases when the female member retires. The gap is greatest when the male member is working and the female member is not working.

To document these effects more formally, we next present results from a panel regression analysis similar to what we did in the case of individuals, though here we focus on the largest group in the above table. Specifically, we pool all of the household level data across surveys and focus on households in which the male moves from full time work to retired while the female is retired in both periods. For these households we then run a fixed effects regression for the z of the male and female members. That is, we run a regression of the form:

$$z_{it} = \bar{z}_i + \beta I_{Rt}^m + \varepsilon_{it} \tag{4.1}$$

where \bar{z}_i is an individual fixed effect and I_{Rt}^m is a dummy variable equal to one if the male member of the household meets our criterion for retirement in period t . Once again we consider samples based on both of our retirement criterion. Based on criterion 1, the estimated value of β is .013 for males and $-.002$ for females, with standard errors of .017 and .019 respectively. Based on the second criterion, the estimated values of β are .026 for males and $-.013$ for females, with standard

errors of .019 in both cases. Importantly, the standard errors are quite small.

In summary, when the male member of a household moves from working full time to retired in a household in which the female household member is not working, the point estimates suggests a very modest increase in the male z and a very modest decrease in the female z , though in neither case are the results statistically significant at the 5% level. In the next section we develop a model to help us infer the implications of this finding.

5. A Model of Household Time Allocation

In this section we present a model of household time allocation for a multi-member household and derive its implications for the optimal profile of home production and leisure across household members and how it changes over time. Our analysis focuses entirely on first order conditions that characterize static choices within a given period for a subset of variables. For this reason our analysis can accommodate a great deal of generality along several dimensions. For ease of exposition we first develop the key relationship of interest in the context of a fairly standard deterministic formulation of the household life cycle optimization problem, and later discuss robustness to allowing for many alternative features.

5.1. Model

We consider a household that consists of two members, that we refer to as the male and female members. The period utility function for household i is written as:

$$u^i(c_{it}, \frac{\alpha_{im}}{1 - (1/\gamma)} l_{imt}^{1-\frac{1}{\gamma}} + \frac{\alpha_{if}}{1 - (1/\gamma)} l_{ift}^{1-\frac{1}{\gamma}})$$

where c_{it} is the flow of consumption services for the household in period t , and l_{imt} and l_{ift} are male and female leisure in period t respectively. The function u^i is allowed to vary across households and is assumed to be C^2 , increasing in each argument, weakly concave jointly in both arguments and strictly concave in each argument individually. The parameters α_{im} and α_{if} are household specific positive constants. While this functional form imposes some structure on how leisure enters into the utility function, it is very flexible in terms of how the leisure aggregate and consumption interact. In particular, we do not impose separability between household consumption and household leisure. One special case of interest commonly found in the macro literature would be:

$$\log(c_{it}) + \frac{\alpha_{im}}{1 - (1/\gamma)} l_{imt}^{1-\frac{1}{\gamma}} + \frac{\alpha_{if}}{1 - (1/\gamma)} l_{ift}^{1-\frac{1}{\gamma}}$$

in which case the parameter γ governs both the elasticity of substitution between leisure of the two household members as well as the intertemporal elasticity of substitution of leisure.⁶ A slightly more general specification would be:

$$\log(c_{it}) + v(\frac{\alpha_{im}}{1 - (1/\gamma)} l_{imt}^{1-\frac{1}{\gamma}} + \frac{\alpha_{if}}{1 - (1/\gamma)} l_{ift}^{1-\frac{1}{\gamma}})$$

⁶This special case raises the possibility that one might want to consider gender specific values of γ . We carry out an exercise later in the paper that allows for gender specific values of γ and find no evidence to support this.

where v is some increasing and concave function, in which case γ governs the elasticity of substitution between leisure of different members but not necessarily the intertemporal elasticity of substitution.

The flow of household consumption is a CES aggregate of household expenditure (g_{it}) and household efficiency units of home production time (h_{it}):

$$c_{it} = [a_i g_{it}^{1-\frac{1}{\eta}} + (1 - a_i) h_{it}^{1-\frac{1}{\eta}}]^{\frac{\eta}{\eta-1}}$$

Efficiency units of home production time at the household level are in turn a CES aggregate of male and female home production time, denoted by h_{mt} and h_{ft} respectively:

$$h_{it} = [A_{im} h_{imt}^{1-\frac{1}{\rho}} + A_{if} h_{ift}^{1-\frac{1}{\rho}}]^{\frac{\rho}{\rho-1}}$$

where $\rho \geq 0$ is the elasticity of substitution between the time of the two members in household production. Although this specification nests the special case of perfect substitutes, i.e., ρ tending to infinity and $A_{im} = A_{if}$, it allows for much more generality. The special case of perfect substitutes is empirically problematic because it creates a tendency for corner solutions in home production time, a property that is not found in the data.

As emphasized with our notation, we allow the α , a , and A parameters to all be household specific. The α 's can reflect true differences in preferences for leisure across individuals within the household, or could reflect the differential weights that the household places on the utility of its different members. Similarly, dif-

ferences in a across households could reflect differences in their ability to combine goods and home production time or differences in preferences.

We normalize the total amount of discretionary time to equal unity for each member of the household, so that leisure is equal to one minus the sum of time spent in market work (m) plus home production (h):

$$l_{ijt} = 1 - m_{ijt} - h_{ijt}, j = m, f$$

We assume that the household maximizes utility over a T period horizon, using a discount factor β . The household faces a sequence of budget constraints given by:

$$g_{it} + a_{it} = w_{imt}m_{imt} + w_{ift}m_{ift} + (1 + r_t)a_{it-1}$$

where w_{ijt} is the wage for member j in household i in period t , and m_{ijt} is hours of market work for member j in household i in period t .

5.2. Optimal Home Production Decisions

In what follows we focus entirely on the implications of the optimal choice of home production time in a given period taking as given the choices for market work and spending on goods. As we discuss in greater detail below, while this approach does not utilize all of the structure of the household problem, its advantage is that it is robust to a wide variety of specifications regarding some aspects of the household problem, including some that may be controversial and/or complicated.

In each period the household chooses how much time each member should allocate to home production given all of the other variables, yielding two first order conditions, one for h_{imt} and one for h_{ift} . Assuming interior solutions for each of these choices and abstracting from the household index i for notational convenience, these two first order conditions are:

$$\begin{aligned} h_{mt} &: u_1(c_t, l_t)(1-a)c_t^{\frac{1}{\eta-1}}h_t^{\frac{1}{\rho}-\frac{1}{\eta}}A_m h_{mt}^{-\frac{1}{\rho}} = u_2(c_t, l_t)\alpha_m(1-m_{mt}-h_{mt})^{-\frac{1}{\gamma}} \\ h_{ft} &: u_1(c_t, l_t)(1-a)c_t^{\frac{1}{\eta-1}}h_t^{\frac{1}{\rho}-\frac{1}{\eta}}A_f h_{ft}^{-\frac{1}{\rho}} = u_2(c_t, l_t)\alpha_f(1-m_{ft}-h_{ft})^{-\frac{1}{\gamma}} \end{aligned}$$

where $l_t = \frac{\alpha_{im}}{1-(1/\gamma)}l_{imt}^{1-\frac{1}{\gamma}} + \frac{\alpha_{if}}{1-(1/\gamma)}l_{ift}^{1-\frac{1}{\gamma}}$.

Dividing the two first order conditions by each other and recalling that $1 - m_{jt} - h_{jt} = l_{jt}$ gives:

$$\frac{\alpha_m}{\alpha_f} \left[\frac{l_{mt}}{l_{ft}} \right]^{-\frac{1}{\gamma}} = \frac{A_m}{A_f} \left[\frac{h_{mt}}{h_{ft}} \right]^{-\frac{1}{\rho}} \quad (5.1)$$

This equation reflects a purely static condition for household optimization: the marginal rate of substitution between leisure of the two members must equal the marginal rate of transformation of the two members' time spent in home production. As discussed in the introduction, our goal is to use data on time allocations to learn about key elasticity parameters. While the above expression imposes some structure, this structure effectively provides no information if the only data we have is from a single cross-section. To see this, note that given any values for γ and ρ , we can always rationalize any pattern of time allocation

within the household by appealing to an appropriate profile of preference and productivity parameters, i.e., the α_j 's and the A_j 's. Put somewhat differently, cross-section data alone provides effectively no information about the preference parameters γ and ρ .

However, the situation is very different if we have access to panel data on time allocations. To see why, begin by rewriting equation (5.1) as:

$$\left[\frac{l_{mt}}{l_{ft}}\right]^{\frac{1}{\gamma}} \left[\frac{h_{mt}}{h_{ft}}\right]^{-\frac{1}{\rho}} = \frac{\alpha_m A_f}{\alpha_f A_m} \quad (5.2)$$

Assuming that the values on the right-hand side of this equation do not change over time, we can eliminate them by first differencing. More specifically, taking logs and first differencing yields:

$$\Delta \log l_m - \Delta \log l_f = \frac{\gamma}{\rho} [\Delta \log h_m - \Delta \log h_f] \quad (5.3)$$

In words, the relative change in leisure across household members should be proportional to the relative change in home production time. The key point is that given access to panel data on time allocations, the household specific values can be removed by first differencing, and the theory imposes quite a bit of structure on the changes in household allocation over time and the two elasticity parameters γ and ρ . It is notable that this expression includes a curvature parameter from both preferences and technology. As Gronau (1997) noted in his survey paper, there is a fundamental identification problem in the home production literature that has often been avoided by abstracting from curvature in the production function.

Although our specification of the home production function is constant returns to scale, it does feature curvature with regard to each of the individual time inputs.

Although one might think that a multi-member household makes analysis more complicated, it is interesting to note that the assumption of a multi-member household is key to the ability to derive a condition that involves only changes in time allocation and the two preference parameters γ and ρ based purely on static first order conditions. In Rogerson and Wallenius (2016) we performed a similar analysis in the context of a single individual household. But in that case our final expression involved both time allocations and consumption expenditure (i.e., g_t), and required that the household's choices satisfied the consumption Euler equation.

As one special case, note that if we assumed perfect substitutes in home production, i.e., the limiting case as ρ tends to infinity, then equation (5.3) implies:

$$\Delta \log l_m - \Delta \log l_f = 0 \tag{5.4}$$

and the parameter γ disappears from the expression. This expression is inconsistent with the main finding from the previous section.⁷

More generally, the above expression implies a value for γ/ρ given changes in time allocations. Given a value for ρ we can then recover the implied value for γ . For given changes in time allocations, the implied value of γ is increasing in the value of ρ , so that the further we move away from perfect substitutes the

⁷It is important to recall our previous comment about the tendency for perfect substitutes to lead to corner solutions, given that this expression assumes interior solutions. Recall, however, that we take market hours as given in this derivation.

smaller is the implied value for γ . It seems reasonable to assume that the time inputs in home production are substitutes rather than complements, so that the lower bound for γ would correspond to the case in which the home production aggregator is Cobb-Douglas, i.e., $\rho = 1$

The result that a higher value of ρ implies a higher value of γ holding the changes in time use fixed is intuitive and straightforward. Taking the change in time use as given, suppose we have a pair of values for γ and ρ such that equation (5.3) holds, and assume the nature of the change across time is that both home production and leisure time increase for the male member of the household. At the given values of γ and ρ the male choices are such that the increase in household utility from marginally higher male leisure is exactly equal to the increase in household utility from a marginal increase in male home production time. If we consider a higher value of ρ then the marginal utility from increasing male home production increases as the extent of decreasing returns is lessened. To maintain equality we must have that the marginal utility of leisure must also increase, which means less curvature in leisure.

5.3. Interpreting Transitions Out of Full Time Work

As we have noted previously, we will be particularly interested in using equation (5.3) to study changes in time allocations in the context of transitions out of full time work. In order to justify this it is important to have some discussion of what drives these transitions. One possibility is an adverse shocks to market opportunities. Note that because our derivation placed no restrictions on how market

opportunities change over time, this creates no problems. A second possibility is that transitions out of full time work are not in response to a contemporaneous shock, but are instead the result of choosing an optimal path of labor supply over the life cycle in an environment with some sort of frictions or nonconvexities. Hurd (1996) argued that restrictions on the ability to choose hours are a key driving force behind retirement, and Blau and Shvydko (2011) and Ameriks et al (2017) are recent works supporting this view. Alternatively, the models of French (2005) and Rogerson and Wallenius (2009) generate retirement in response to nonconvexities in the compensation structure. Once again, our derivation placed no restrictions on these features.

Our strategy would be invalid if the transition out of full time work was driven by changes in either the A'_i s or the α'_i s. If these transitions were driven by health shocks then it would be natural that they were systematically related to changes in these parameters.⁸ Note that health shocks are only a problem if they are *contemporaneous* with the transition out of full time work. In particular, consider an individual of age 55 who “plans” to retire at age 65. Suppose this individual experiences an adverse health shock at age 55 and as a result ends up retiring at age 62 instead of 65. Although the health shock in this case strongly influences the timing of retirement, this case is not problematic for our strategy if health is stable between 60 and 62. That is, the presence of permanent health shocks per se is not a challenge to our strategy.

Blau and Shvydko (2011) present evidence that very few retirements are the

⁸Note that *iid* shocks to these parameters is not an issue as this would simply show up as an error term in equation (5.3).

direct result of health shocks. Nonetheless, in order to address this issue we use the self-reported health status question in the HRS to create a subsample in which all individuals are in good health or better.⁹ At the individual level, this healthy subsample is about two-thirds of the overall sample. In the Appendix we report the equivalent of Tables 1, 2 and 4 for this subsample, and show that the key patterns remain unchanged. We also repeat the fixed effect regressions for this subsample, and again find that it does not affect our key finding, though the point estimates are a bit larger in absolute value.

5.4. Extensions

In this section we note a variety of extensions to which our key estimating equation is robust. We emphasize that our analysis is robust to a large number of extensions which are now well known to have first order effects on estimates of γ in other contexts.¹⁰ While some of the robustness in the current framework mirrors the discussion in Rogerson and Wallenius (2016), the fact that our current analysis does not rely on any dynamic choices renders it robust to even more factors. First, our equation depends in no way on the set of choices for market hours that household members face (e.g., indivisible labor), whether the observed choices for market hours are optimal (i.e., whether individuals are on their labor supply

⁹The health status question in the HRS ask individuals to rate their current health as excellent, very good, good, fair or poor.

¹⁰See for example, Imai and Keane (2004) and Wallenius (2011) for analyses with human capital accumulation, Domeij and Floden (2006) for an analysis with credit constraints, Low (2005) for an analysis with incomplete insurance markets, Chang and Kim (2006) and Rogerson (2011) for analyses with restrictions on working hours, and Ham and Reilly (2013) for an analysis with implicit contracts. See also the survey in Keane and Rogerson (2015).

curve for market work), whether market work is associated with human capital accumulation, and whether there are non-linearities in the compensation structure.

Because our analysis does not rely on dynamic first order conditions, it is also invariant to the presence of credit constraints. Although we formulated the household problem without any sources of randomness, our approach is robust to allowing for stochastic market opportunities and whether there are incomplete markets to insure against randomness in market prices. In particular, our key equation is robust to embedding our analysis in the standard Aiyagari style model.

Our model description did not include any tax and transfer programs, but our key equation is invariant to any form of tax and transfer policies that are functions of market work and market income. In particular, given that we will be focusing on older individuals and the transitions that they make when retiring, it is important to know that the presence of a realistic social security system has no impact on our key equation.

6. Implications for Parameter Values

In this section we use the CAMS data and our model of household decision-making to generate information on the joint values of γ and ρ . We use two different methods to produce estimates using equation (5.3). The first method will use the key property documented earlier: the relative constancy of the z values for both household members when one of the members transitions from working to retired. We will show that this “average response” has a sharp prediction for the value of γ/ρ . The second method generates estimates of γ/ρ by directly estimating

equation (5.3) using the full set of household observations. A key issue for this second method is the concern that the CAMS data features significant amounts of measurement error.

6.1. Estimates Based on a Typical Transition

We begin by asking what an “average” transition implies for the values of ρ and γ . Specifically, we use the fixed effect panel regression results for households from the previous section to impute values for the left and right hand side variables in equation (5.3) and infer a value for γ/ρ .

We start with a benchmark calculation that delivers a very sharp result. In particular, one of the patterns documented in the previous section was that the value for z is very close to constant for both household members, even when one of them transitions from full time work to retirement. Because $z = h/(h + l)$, it follows that a constant value of z implies that the percentage change in h is the same as the percentage change in l , i.e., that for each member of the household $\Delta \log l = \Delta \log h$. It thus follows that $\Delta \log l_m - \Delta \log l_f = \Delta \log h_m - \Delta \log h_f$. Viewed through the lens of equation (5.3) the implication is that $\gamma/\rho = 1$. Note that this conclusion holds independently of what the initial value of z 's were for the two household members, what their hours of market work were prior to retirement, and whether there was a change in total discretionary time.

Given that $\gamma/\rho = 1$, it is trivial to trace out the value of γ for different values of ρ . Previously we argued that $\rho = 1$ is a reasonable lower bound for ρ , which would then also imply that unity is a lower bound for γ . In a model with a similar

production technology, Knowles (2013) estimates a value of $\rho = 3$ based on time series changes in the male-female wage differential and the male-female differential in time spent in home production. Because his estimate attributes all of the fall in the home production time gap across genders to changes in relative wages it is reasonable to view this as an upper bound on ρ . But strikingly, this value of ρ would of course imply a value of three for γ .

The above calculation assumed that z was constant for both individuals. While the evidence in the previous section suggests that the data closely conforms to this pattern, we did provide some evidence of small changes in z that were marginally statistically significant. Here we examine the extent to which allowing for changes of the magnitude estimated in the previous section affect the implications for the value of γ/ρ implied by equation (5.3). For concreteness we focus on the case of a household in which the female is not working in both surveys and the male goes from full time work to no work across the surveys. We impute the following values for the variables in equation (5.3) using the information at the end of the previous section. Because there is no statistically significant change in z for the female member we set $\Delta \log h_{ft}$ and $\Delta \log l_{ft}$ equal to zero. For the male household member, we assume that market work when working full time is 40 hours, total discretionary time is 100 hours, the value of z before retirement is .21, and that the increase in z following retirement is equal to .020.

The implied value for γ/ρ is 0.80, so that allowing for a modest increase in z for the male tends to decrease the implied value of γ for any given value of ρ . This estimate is only modestly affected by changes in the other values assumed in this

calculation. Assuming that total discretionary time is 90 hours per week instead of 100 hours per week increases the estimate of γ/ρ to 0.82. And increasing the working time before retirement to 45 hours per week produces a change of the same magnitude. Assuming that z increases by .04 instead of .02 implies a value for γ/ρ of .69. Finally, we consider modest increases in the male value of z in combination with a modest decrease in the female value of z . Specifically, we consider an increase of .03 for the male and a decrease of .03 for the female. The implied value of γ/ρ is .58. It remains true that even very modest values of ρ would suggest values of γ in excess of unity. If we consider an increase of .05 for the male and a decrease of .03 for the female the implied value for γ/ρ is .54.

We conclude that the key pattern that we document in the CAMS—that the value of z is nearly constant for each gender even when a household member moves into retirement—suggests a reasonably high value for the labor supply elasticity parameter γ . In particular, the implied value of γ is likely at least as high as unity and potentially significantly higher, depending upon the value of ρ .

6.2. Estimates Based on Panel Regression

We now turn to providing estimates of γ/ρ from panel regression estimates of equation (5.3). Recall that our derivation implied that this condition should hold in the face of any changes in the economic environment that generate changes in some component of time allocations holding parameters fixed. In this sense we can run this regression for the entire sample of matched households.

A key issue that we have discussed earlier is that both right and left hand side

variables in equation (5.3) are likely to be measured with considerable error. Measurement error in the left hand side variables will of course not bias the estimates, but measurement error in the right hand side variable will bias the estimated coefficient toward zero. As is well known, one can run both the specification in (5.3) as well as the reverse specification with left and right hand sides switched in order to generate an interval of estimates, but this of course does not eliminate the effect of measurement error. Attempts to deal with measurement error will be a major focus of the exercises in this subsection.

Before proceeding we discuss the type of measurement error that our specification can accommodate. One source of measurement error in the CAMS is that total time use need not add up to total time available. We previously suggested that if the extent of this problem were the same across leisure and homework that our variable z would be unaffected on average. Here we want to note that equation (5.3) is robust to a much more general specification of measurement error. In particular, we can assume that each gender and each category have its own proportional error in addition to an iid term that reflects classical measurement error. That is, we can assume for example, that reported time spent in home production by a member of gender g in period t in the survey, denoted by \hat{h}_{gt} , is related to true time spent in home production, denoted h_{gt} by:

$$h_{gt} = B_{gh} \hat{h}_{gt} \varepsilon_{hgt}$$

where B_g reflects the fact that only a fraction of total time is reported in this category and $\log(\varepsilon_{hgt})$ is classical measurement error. Because our estimating

equation takes log differences by gender over time, the B_{gh} terms will all cancel, leaving only classical measurement error. More generally, we could even allow for a deterministic trend in the B_{gh} terms to capture some systematic component of measurement error associated with aging by including a constant term in our estimating equation.

As a starting point, we run specification (5.3) using our sample of matched households across consecutive surveys. Each matched pair leads to one observation. A given household that appears in all of the surveys could contribute 4 observations to our sample, subject to there being no missing values that exclude them. The resulting sample size is 4922.

Consistent with the discussion above, we run both the specification in this equation as well as its mirror image with the right and left hand variables reversed. When estimating equation (5.3) we get a point estimate of .130 for γ/ρ with a standard error of .013. When we run the regression with the left and right hand side variables reversed we obtain a point estimate of .259 for ρ/γ with a standard error of .026. Both point estimates are statistically significant and of the appropriate sign, thus supporting the basic economic mechanism in our model. Absent measurement error and assuming that the model were correct, the two point estimates should be the inverse of each other, so that one of the estimates would be smaller than one and the other would be larger than one. The presence of measurement error biases both of them toward zero, and thus can explain why both are smaller than one.

The implied range of values for γ is very large, and the lower bound for γ is

increasing in the value of ρ . Taking $\rho = 1$ as a reasonable lower bound for ρ , the two implied values for γ would be .130 and 3.86. Note that for a given value of ρ , measurement error biases the implied value of γ toward 0 when running the regression as in equation (5.3), but biases the implied value of γ upward when running the reverse regression. For this reason we regard these two values as natural bounds. For $\rho = 3$, the implied values for γ become .390 and 11.58. To the extent that the latter value is biased upward and we view 3.86 as the lowest upper bound we do not view the 11.58 value as particularly relevant. But the value of the lower number is significant to the extent that it is biased downward and so represents a lower bound.

The sample used to run the above regressions included all observations in which household members could be matched over time. It is perhaps to be expected that a lot of the variation in time allocations in this sample might reflect measurement error. One way to dampen the potential effect of measurement error is to select a subsample where the relative importance of measurement error might be lower. To do this we focus on households in which one of the members transitions from full time work to not working. For such an individual we expect there to be large changes in both leisure and home production, thus hopefully increasing the signal to noise ratio. For this exercise we construct the sample in the following manner. The data for a household from the surveys in t and $t + 2$ will be included if the household is in the survey at each of t , $t + 2$, and $t + 4$, and at least one member works at least 35 hours in the survey at t and no more than 5 hours in the surveys at both $t + 2$ and $t + 4$. That is, our sample consists of households

that experience at least one member moving in a persistent way from full time work to retirement. Note that we do not place any restrictions on the choice of the other member. The resulting sample size is 190 observations.

We run the same two regressions for this sample as we did for the original sample. When we run the regression as in equation (5.3) we get a point estimate for γ/ρ of .234 with a standard error of .065. When we run the reverse regression we obtain a point estimate for ρ/γ of .426 with a standard error of .097. For $\rho = 1$ the implied values for γ are .234 and 2.35, while for $\rho = 3$ the values are 0.702 and 7.04. As expected, these point estimates are somewhat larger than for the larger sample, though the effect is much smaller for the first estimate.

In order to further dampen the effect of measurement error we carry out one additional exercise. Specifically, using the same sample as above, we now construct an instrument for the right and left hand side variables. Our instrument is motivated by the fact that in the absence of any large “shock” time allocations should stay relatively constant. In the context of our sample, for a household who contributes data to our sample between the surveys at t and $t + 2$, we would expect that there is relatively little change in the true time allocations between the two surveys at $t + 2$ and $t + 4$ because the large “shock” for this household is the retirement of one member between t and $t + 2$. That is, changes between $t + 2$ and $t + 4$ might represent mostly measurement error. This motivates us to use the change in time allocations between t and $t + 4$ as an instrument for the changes of the right-hand side variable between t and $t + 2$. Once again, we can run this regression in both directions. When we do this we obtain an estimate for

γ/ρ of .286, with a standard error of .132, and an estimate of .684 for ρ/γ with a standard error of .177. In both cases the first stage is significant. For equation (5.3) the estimate in the first stage is .467 with a standard error of .111, while in the reverse regression the estimate in the first stage is .646 with a standard error of .080.

The estimates that result from these exercises change in the expected way, in that when we make an effort to reduce the effect of measurement error the estimates seem to move away from zero in absolute value. However, in all cases it remains true that both estimates are smaller than one, which is consistent with the notion that considerable measurement error remains. While it is perhaps disappointing that we do not obtain sharper results from the panel estimation, it is important to note that there is no inconsistency between the estimates from the two methods. That is, the results of the panel regression estimates are consistent with the value for γ/ρ that we inferred from simply evaluating equation (5.3) under the assumption that z is constant for both household members during a retirement transition in that they lie within the interval generated by the two different estimates.

6.3. Heterogeneous γ by Gender

Earlier in the paper we noted that if one interprets γ as evidence about the intertemporal elasticity of substitution it might be of interest to allow for this value to differ by gender. In this subsection we report the results of one exercise that speaks to this possibility. In the interests of space we do not go through

the derivation here, but it is straightforward to show that if we had allowed for heterogeneous values of γ , then going through the same derivation as before we would have ended up with the expression:

$$\Delta \log h_m - \Delta \log h_f = \frac{\rho}{\gamma_m} \Delta \log l_m - \frac{\rho}{\gamma_f} \Delta \log l_f$$

It follows that when estimating the inverse of equation (5.3) on panel data it is straightforward to allow γ to vary by gender. When we estimate this expression allowing for gender specific coefficients we obtain point estimates of .254 and .265 for ρ/γ_m and ρ/γ_f respectively, both with standard errors of .034. We conclude that imposing $\gamma_m = \gamma_f$ is consistent with our data.¹¹

7. Evidence from Other Countries

To this point we have focused on time use patterns in the US. Given that time use data is available for many countries it is of interest to ask whether the patterns that we find for the US are also present in other countries. In this section we examine this using time use data from the MTUS covering France, Germany, Italy, the Netherlands, Norway, Spain and the UK . For each of these countries we can present summary statistics similar to what we did for the ATUS. That is, we can analyze how the ratio of discretionary non-market time at the individual level is divided between home production and leisure as we vary age.

For this analysis we concentrate on what happens during the age range from

¹¹This result also holds when we consider the other estimation results considered above, though the standard errors become larger when we allow the γ 's to vary by gender.

50 to 65 since this age range covers the most substantial decreases in market work associated with the process of retirement. Our main goal is to assess whether there is any tendency for our ratio z to change with age over this age range. To present the evidence we run a simple regression of age specific mean values of z against a constant and age, separately for each gender:

$$z_{acg} = a_{cg} + b_{cg}a + \varepsilon_{acg}$$

where a is age, c is country, and g is gender.

Table 6 presents the results of this regression exercise. The column headed by #obs reflects the average number of observations for each age.

Table 6

Change in z by Age: MTUS

	Male				Female			
	a	b	#obs	σ_z	a	b	#obs	σ_z
France	.2454 (.0094)	.0051 (.0011)	98.5	.227	.5321 (.0105)	−.0049 (.0012)	103.3	.197
Germany	.2818 (.0100)	.0033 (.0011)	241.9	.221	.4549 (.0086)	−.0009 (.0010)	246.0	.202
Italy	.2115 (.0068)	.0057 (.0008)	326.1	.221	.6068 (.0064)	−.0022 (.0007)	342.3	.187
Netherlands	.2365 (.0113)	.0038 (.0013)	286.8	.214	.4068 (.0082)	−.0020 (.0009)	355.8	.199
Norway	.2856 (.0117)	−.0004 (.0013)	90.3	.202	.3895 (.0115)	−.0015 (.0013)	85.5	.171
Spain	.2141 (.0042)	−.0006 (.0005)	510.7	.215	.5342 (.0058)	−.0028 (.0007)	571.2	.209
UK	.2591 (.0124)	.0016 (.0016)	175.4	.213	.4282 (.0091)	−.0028 (.0010)	189.7	.202

Several aspects of the results are worth noting. First, there are large differences in the constant term a_{cg} across countries for a given gender, which basically translates into differences in mean z . For males the constant term varies from a low of .2115 in Italy to a high of .3895 in Norway. For females the constant term varies from a low of .3895 in Norway to a high of .6068 in Italy. Accordingly, the

gender gap in a_{cg} varies quite substantially across countries, from a low of .1039 in Norway to a high of .3935 in Italy. Second, the dispersion of z even among individuals of a given gender and age in a given country is substantial. The column labeled σ_z presents the average value of the age specific standard deviation of z . Similar to what we found in the ATUS, this value tends to be around .20, with relatively little variation across gender or country. Third, although the modal tendency is for male z to increase with age and for female z to decrease with age, the magnitude of these effects is for the most part quite modest. For example, a point estimate for b of .003 implies an increase of .045 for z over a period of fifteen years. To the extent that our age range captures most of the retirement in the data this magnitude is comparable to our point estimate for the change in z for a male moving from full time work to retirement. The point estimates for France and Italy are both a bit larger, while those for Norway and Spain are smaller. The point estimates for females are even smaller in absolute value.

Accepting that these estimates do not come from panel data and are only at the individual level, a simple calculation suggests the same message as the CAMS data. In particular, consider a household in which both individuals have discretionary time of 100 hours, the male moves from full time work (40 hours) to retirement, the female moves from part time work to retirement (20 hours), and that for the male z increases from .250 to .295 while for the female z falls from .450 to .420. Using equation (5.3), the implied value for γ/ρ is .60. If we instead assumed that these changes happened with no change for market work for the female the implied value would increase to .75. We conclude that the available

data from the MTUS suggests estimates of γ that are similar to those implied by the US data.

8. Conclusion

We study what happens to household time allocations when one or more members retires. Unlike the vast majority of studies of time use, we examine these changes using panel data that contains information about both household members in two person households. The striking pattern that we find is that very little happens to the way that individuals allocate their nonmarket discretionary time between leisure and home production in response to retirement. Additionally, we find that there is considerable heterogeneity across households in the way that this time is allocated, and that this heterogeneity is very persistent.

We then develop a multi-member household model of time use and show how the key pattern found in the data can be used to infer information about two key elasticity parameters: the elasticity of substitution between the time of household members in home production and the elasticity of substitution between the leisure time of household members. In some commonly studied settings, this latter elasticity will also be equal to the household's intertemporal elasticity of substitution for leisure. Our theory places a joint restriction on these two elasticity parameters and changes in household time use. This restriction is robust to a variety of model features, and for what we view as very conservative values for the production elasticity of substitution, we still find that the preference elasticity is quite large, most likely greater than unity. We also present data from several

other countries to suggest that the key pattern we document in US data appears to hold more generally.

Appendix

In this appendix we present results when we restrict the analysis to individuals who have good health or better for each of the relevant surveys. The health data is from the HRS. For the wave pair 2005-07, we impose the health restriction in 2004, 2006 and 2008. We will refer to this as the healthy subsample. The first two tables reproduce results for Tables 1 and 2 in the paper. Specifically, Table A1 considers individuals that are matched across consecutive surveys, requiring that individuals report health status of good or better in three consecutive surveys of the HRS.

Table A1

Value of z for Matched Individuals by Gender, Healthy Subsample

	Males					Females				
	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$
2005-07	.25	.25	.14	.14	.46	.31	.31	.13	.14	.49
2007-09	.26	.25	.14	.13	.54	.31	.30	.14	.14	.51
2009-11	.25	.25	.13	.14	.58	.30	.30	.13	.13	.52
2011-13	.25	.24	.13	.13	.53	.30	.29	.13	.14	.49

The patterns are effectively unchanged from Table 1 in the text. Table A2 restricts attention to those individuals in the previous sample who transition from full time work to retirement across the two consecutive surveys.

Table A2

Value of z for Individuals Retiring Across Surveys, Health Subsample

	Males					Females				
	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$
2005-07	.26	.29	.14	.15	.38	.30	.30	.14	.13	.18
2007-09	.27	.26	.13	.14	.61	.32	.34	.15	.14	.20
2009-11	.28	.31	.15	.16	.71	.33	.34	.13	.14	.24
2011-13	.23	.25	.12	.13	.25	.29	.31	.13	.15	.33

Once again, the patterns in this table are virtually identical to those in Table 2. The one small difference is that for this sample there is evidence for a modest increase in z both for males and females. We have also run the fixed effects regression specification for this sample. For the sample in Table A2 the estimates of β are .015 for males and .003 for females, with standard errors of .011 and .013 respectively. As in the text, we have also created a sample based on matching individuals across three surveys in order to have a stricter definition of retirement, in which the individual needs to work no more than 5 hours in each of the second and third surveys. In this case we also require that the individual self report being in at least good health in four consecutive surveys of the HRS. For this sample the estimates of β are .019 for males and $-.011$ for females, with standard errors of .013 and .015 respectively. These results mirror those found for the entire sample.

Next we consider data at the household level. As in the text, we consider all observations in which we can match both members of a household across consecutive surveys, but with the added restriction that both household members exhibit

health status of good or better in the relevant surveys.

Table A3

z for Matched Households, Healthy Subsample

	μ_{z_t}	$\mu_{z_{t+1}}$	σ_{z_t}	$\sigma_{z_{t+1}}$	$\rho_{z_t, z_{t+1}}$
2005 – 07	.29	.28	.11	.10	.54
2007 – 09	.28	.28	.10	.11	.58
2009 – 11	.28	.27	.10	.10	.58
2011 – 13	.27	.27	.10	.10	.49

Once again, the results strongly mirror those in Table 4 in the text. When we condition on various types of transitions within the household, some of the sample sizes now become quite small. As before, the largest subsample is those households in which the male goes from full time work to retirement while the female member works no more than five hours in both surveys. This subsample has 43 observations. When we run the fixed effects regression for this subsample we obtain estimates of β equal to .032 for males and $-.034$ for females, with standard errors of .021 and .033 respectively.

References

- [1] Aguiar, M., and Hurst, E. 2005. "Consumption and Expenditure." *Journal of Political Economy* 113, 919-948.
- [2] _____ 2007 "Life Cycle Prices and Consumption," *American Economic Review* 97, 1533-1559.
- [3] Aguiar, M., Hurst, E., and Karabarbounis, L. 2013. "Time Use During the Great Recession." *American Economic Review* 103, 1664-1696.
- [4] Altonji, J., 1986. "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *Journal of Political Economy* 94, S176-S215.
- [5] Ameriks, J., Briggs, J., Caplin, A., Shapiro, M., and Tonetti, C. 2017. "Older Americans Would Work Longer if Jobs Were Flexible." Working Paper, Stanford University.
- [6] Attanasio, O., Low, H., and Sanchez-Marcos, V., 2008. "Explaining Changes in Female Labor Supply in a Life Cycle Model." *American Economic Review* 98, 1517-1552.
- [7] Attanasio, O., Levell, P., Low, H. and Sanchez-Marcos, V. 2015. "Aggregating Elasticities: Intensive and Extensive Margins of Female Labor Supply." NBER Working Paper #21315.
- [8] Becker, G. "A Theory of the Allocation of Time." *The Economic Journal* 75 (1965): 493-517.

- [9] Blau, D, and Shvydko, T. 2011. "Labor Market Rigidities and Employment Behavior of Older Workers." *Industrial and Labor Relations Review* 64, Article 3.
- [10] Blundell, R., Pistaferri, L., and Saporta-Eksten, I. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review* 106, 387-435
- [11] Browning, M., Deaton, A. and Irish, M. (1985). "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle." *Econometrica* 53, 503-44.
- [12] Chang, Y., and Kim, S. 2006. "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy." *International Economic Review* 47, 1-27.
- [13] Cullen. J., Gruber, J. 2000. "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" *Journal of Labor Economics* 18, 546-572.
- [14] Davis, S., and Henrekson, M. 2004. "Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich Country Comparisons." NBER Working Paper #1059.
- [15] Domeij, D., and Floden, M. 2006. "The Labor Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased." *Review of Economic Dynamics* 9, 242-262.
- [16] French, E., 2005. "The Effect of Health, Wealth and Wages on Labour Supply and Retirement Behavior." *Review of Economic Studies* 72 (2005), 395-427.

- [17] Gronau, R. 1977. "Leisure, Home Production and Work—The Theory of the Allocation of Time Revisited." *Journal of Political Economy* 75, 1099-1124.
- [18] _____. 1997 "The Theory of Home Production: The Last Ten Years." *Journal of Labor Economics* 15, 197-205.
- [19] Ham, J., and Reilly, K. 2013. "Using Micro Data to Estimate the Intertemporal Substitution Elasticity for Labor Supply in an Implicit Contract Model." forthcoming in *International Economic Review*.
- [20] Hamermesh, D., and Donald, S. 2007. "The Time and Timing Costs of Market Work." NBER Working Paper #13127.
- [21] Heckman, J. and MaCurdy, T. 1980. "A Life Cycle Model of Female Labor Supply." *Review of Economic Studies* 47, 47-74.
- [22] Hurd, M. 1996. "The Effect of Labor Market Rigidities on the Labor Force Behavior of Older Workers." in *Advances in the Economics of Aging*, edited by D. Wise, University of Chicago Press, Chicago.
- [23] Hurd, M., and Rohwedder, S. 2003. "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement." NBER Working Paper #9586.
- [24] _____. 2005. "Changes in Consumption and Activities in Retirement." Michigan Retirement Research Center Working Paper 2005-096.

- [25] _____ . 2007. “Time-Use in the Older Population. Variation by Socio-economic Status and Health.” RAND Working Paper WR-463.
- [26] Imai, S., and Keane, M. 2004. “Intertemporal Labor Supply and Human Capital Accumulation.” *International Economic Review* 45, 601-641.
- [27] Keane, M., and Rogerson, R. 2015. “Reconciling Micro and Macro Labor Elasticities: A Structural Perspective.” *Annual Review of Economics* 7, 89-117.
- [28] Knowles, J. 2013. “Why Are Married Men Working So Hard? An Aggregate Analysis of Intra-Household Bargaining and Labor Supply.” *Review of Economic Studies* 80, 1055-1085.
- [29] Laitner, J., and Silverman, D. 2005. “Estimating Life-Cycle Parameters from Consumption Behavior at Retirement.” NBER Working Paper #11163.
- [30] Low, H. 2005. “Self-insurance in a Life Cycle Model of Labour Supply and Savings.” *Review of Economic Dynamics* 8, 945-975.
- [31] Lundberg, S. 1985. “The Added Worker Effect.” *Journal of Labor Economics* 3, 11-37.
- [32] MaCurdy, T. 1981. “An Empirical Model of Labor Supply in a Life Cycle Setting.” *Journal of Political Economy* 89, 1059-1085..

- [33] Rogerson, R. 2011. "Individual and Aggregate Labor Supply in a Model With Coordinated Working Times." *Journal of Money, Credit and Banking* 43, 7-37
- [34] Rogerson, R., and Wallenius, J., 2009. "Micro and Macro Elasticities in a Life Cycle Model with Taxes," *Journal of Economic Theory* 144, 2277-2292.
- [35] _____., 2016. "Retirement, Home Production and Labor Supply Elasticities." *Journal of Monetary Economics* 78, 23-34.
- [36] Wallenius, J. 2011. "Human Capital Accumulation and the Intertemporal Elasticity of Substitution of Labor: How Large is the Bias?" *Review of Economic Dynamics* 14, 577-591.