

Retirement, Home Production and Labor Supply Elasticities*

Richard Rogerson

Johanna Wallenius

Princeton University

Stockholm School of Economics

August 2013

Abstract

We show that a life cycle model with home production implies a tight relationship between key preference parameters and the changes in time allocated to home production and leisure at retirement. We derive this relationship and use data from the ATUS to explore its quantitative implications. Our method implies that the intertemporal elasticity of substitution for leisure is quite large, in excess of one and possibly as high as two. JEL #'s E24, J22.

*Rogerson thanks the NSF for financial support. Wallenius thanks the Yrjo Jahnsson Foundation for financial support. We thank three anonymous referees, Robert Barro, Larry Katz, Gianluca Violante and seminar participants at several institutions for valuable comments.

1. Introduction

Preference parameters that determine labor supply elasticities are critical for many analyses, including, for example, business cycles and optimal tax policy. Two key determinants of labor supply elasticities are an individual's willingness to substitute leisure over time, and an individual's willingness to substitute between home and market produced goods. To fix ideas, consider an individual who solves a dynamic optimization problem and has a period utility function of the form:

$$u(ag_t^{1-\frac{1}{\eta}} + (1-a)h_{nt}^{1-\frac{1}{\eta}}) + \alpha l_t^{1-\frac{1}{\gamma}}$$

where g is expenditure on goods, h_n is time spent in home production, and l is leisure. In this specification these two elasticities are constant and equal to the parameters γ and η respectively.

The literatures that structurally estimate these two elasticity parameters exist independently of each other, in the sense that papers that estimate one of them do not have anything to say about the other. While there are issues in each literature, applied work often adopts values around $\gamma = .40$ and $\eta = 2.00$, on the grounds that these are supported by the available empirical evidence. In this paper we derive a simple relation that links the values of these two key structural parameters to the changes in time use and consumption expenditure at retirement. When we evaluate this restriction using data from the American Time Use Survey (ATUS) we find that it is inconsistent with these commonly adopted values of γ and η . Specifically, either γ is much larger than what is found in many studies,

or η is much smaller than what is found in studies.

Our theoretical framework considers an individual immediately before and after retirement, where retirement takes the form of moving from working full time to not working at all. Although the process of retirement varies across individuals, this is the dominant form of retirement for individuals in their 60s. In the context of such a transition, individuals need to allocate the time that is freed up upon retirement into other activities, notably leisure and home production, and adjust consumption expenditures. In this context, we show that there is a tight link between the ratio of the two elasticity parameters and the change in time allocation and consumption expenditure. While we do impose the functional form for preferences displayed above, the expression that we derive is robust to many other aspects of the economic environment. For example, we do not need to make assumptions about human capital accumulation, credit constraints when young, restrictions on choices of working hours, or how to interpret wage payments. Many standard procedures for identifying these structural parameters are known to be sensitive to each of these aspects.¹ Note that the presence of a home production decision is critical in this analysis: absent a home production margin, all of the increased time available at retirement necessarily goes to leisure, and this is independent of the individual's willingness to substitute leisure over time.

Beyond the implications that we derive for the values of two key preference parameters, another contribution of our analysis is to show how behavior at re-

¹See Imai and Keane (2004) and Wallenius (2011) regarding human capital accumulation, Domeij and Floden (2006) regarding credit constraints for younger workers, Chang and Kim (2006) and Rogerson (2011), regarding restrictions on working hours, and Ham and Reilly (2013) regarding the effect of implicit contracts.

tirement, even when it takes the form of a discrete jump from full time work to no work can be used to generate information about parameters that dictate marginal responses. Looking at the implications of structural parameters derived in one setting for behavior in different settings can be an important source of corroboration. We think this is especially true for preference parameters that shape labor supply elasticities; many of the issues noted above that have generated controversy regarding estimates of γ based on choices made by prime aged individuals (i.e., human capital accumulation, etc...) are arguably much less relevant for individuals at retirement. Moreover, whereas individuals might face restrictions on their ability to adjust their hours of market work, thereby creating a challenging context for making inference about an individual's preferences for different levels of market work, this argument is much less applicable to the allocation of time between home production and leisure. We also think that our general method can potentially be applied to other settings, for example, looking at how workers who move from employment to unemployment allocate the time previously allocated to market work.

Although our analysis focuses on the values of preference parameters that are important determinants of labor supply elasticities, it is important to note that our analysis is not about the value of aggregate labor supply elasticities *per se*. There are several reasons for this. First, aggregate elasticities reflect changes along both the intensive and extensive margins. While γ is certainly relevant for the response along the intensive margin, responses along the extensive margin are largely determined by the extent and nature of heterogeneity, and our analysis

does not provide any information about these features. Rogerson and Wallenius (2009), for example, find that the value of γ has a first order effect on the relative change along the intensive and extensive margins but not on the aggregate change. Second, even with knowledge of the preference parameters that we study, one also needs to know about many other features of the economic environment—things like the human capital accumulation technology, the presence of credit constraints and idiosyncratic shocks, fixed costs of working, labor market frictions, etc...—in order to compute the aggregate response in hours to a given change in either wages or taxes. In this regard, note that even if we knew the correct value of the aggregate labor response in some context, in order to assess welfare and design optimal tax policies it is important to understand how the various elements just mentioned combine to generate this elasticity. The objective of this paper is to learn more about the values of the preference parameters that are one key element of the overall aggregate elasticity.² Recent work on taxes and cross-country differences in market work has also focused on differences leisure and home production.³ The parameter η is very relevant for this work.

Our paper is directly related to the two literatures that provide estimates of γ and η . We summarize key papers from these literatures in Section 4. By virtue of using data on retirement to learn about preference parameters that

²Having offered these qualifications, we refer the reader to the calculations in Rogerson (2009) for a simple model that does not include an extensive margin and shows how the two parameters γ and η influence responses to permanent tax changes in the context of cross-country differences in market work and home production.

³See, for example the papers by Freeman and Schettkat (2001, 2005), Davis and Henrekson (2004), Rogerson (2007, 2008, 2009), Ngai and Pissarides (2008, 2011), Burda et al (2008), Olovsson (2009), and Ragan (2013).

shape labor supply elasticities, this paper is perhaps most related to the recent paper by Rogerson and Wallenius (2013). However, although both papers focus on properties of retirement as a source of information, the underlying sources of identification are very different. In Rogerson and Wallenius (2013), there is no home production, and inference is based on the requirement that the retirement decision is optimal, i.e., that individuals optimally choose to adjust annual hours worked from 2000 to zero despite the presence of intermediate options. In contrast, this paper does not base any inference on the optimality of the retirement decision per se, but instead focuses on how time is allocated between leisure and home production conditional on a worker transiting from full time work to no work. In a model with home production, the changing time allocation between leisure and home production provides information on preference parameters without requiring that the retirement decision is optimal.

An outline of the paper follows. In Section 2 we specify a simple model of choices before and after retirement and derive the key equation linking changes in allocations at retirement to the ratio of the two key elasticity parameters. Section 3 presents data from the ATUS and characterizes the typical changes in allocations that accompanies retirement. In Section 4 we use these estimates to explore the implications for the two elasticity parameters, and in Section 5 we consider how health and income shocks affect our conclusions. Section 6 considers an extension in which we allow for nonconvexities in how individuals value leisure in retirement and Section 7 concludes.

2. Retirement in a Life Cycle Model With Home Production

In this section we describe the life cycle model that we analyze, present first order conditions for the optimal life cycle choices of the household, and derive the key expression that we will use in our quantitative analysis.

2.1. Life Cycle Model

We consider an individual who lives for a_T periods and has preferences over sequences of consumption (c) and leisure (l) at each age a given by:

$$\sum_{a=0}^{a_T} \beta^t \left[u(c_a) + \frac{\alpha}{1 - \frac{1}{\gamma}} l_a^{1 - \frac{1}{\gamma}} \right] \quad (2.1)$$

where u has standard properties, i.e., it is strictly increasing, strictly concave and twice continuously differentiable, and $\gamma > 0$.⁴ We restrict the functional form for the utility from leisure since this is a commonly used specification and the parameter γ will be one focal point of our analysis. In the spirit of Becker (1965) we include home production and assume that the consumption that individuals care about is an aggregate of market purchased goods (g) and home production time (h_n). Following much of the literature, we assume a CES aggregator:

$$c_a = \left[\theta g_a^{1 - \frac{1}{\eta}} + (1 - \theta) h_{na}^{1 - \frac{1}{\eta}} \right]^{\frac{\eta}{\eta - 1}}. \quad (2.2)$$

⁴For simplicity here we abstract from mortality risk. As is well known, mortality risk produces an effective discount factor that differs from the true discount factor. In our empirical application we focus only on the choice between two periods, in which case we can just think of β as representing the composite effect.

where η is the elasticity of substitution between time and goods. In the spirit of Gronau (1977) we distinguish between leisure and working time, so that $l_a = 1 - (h_{ma} + h_{na})$, where we have normalized the total time endowment to unity and h_{ma} is time devoted to market work at age a .

We will not study the entire life cycle decision problem of the individual. Instead we focus entirely on how time allocations change at the time of retirement, assuming that retirement occurs. In Rogerson and Wallenius (2013) we argued that the dominant form that retirement takes in the data is an individual moving from full time work to no work, so we will focus on this type of transition and use the term “retirement” to refer to this type of transition.

Consider an individual who is observed in the data to retire at age a_R , i.e., that this individual moved from full time market work at age $a_R - 1$ to no market work at age a_R . The key equation we derive will depend only on choices made at ages $a_R - 1$ and a_R , and as a result we will not need to specify much of the detail concerning the evolution of variables over the life cycle prior to this point and the various factors that might affect this evolution. For example, we do not have to take a stand on whether human capital accumulation has been an important force in shaping hours and wages earlier in this individual’s lifecycle, on whether the individual was credit constrained when young, or whether the individual was hit by uninsurable idiosyncratic shocks, had “appropriate expectations” during his or her life about the need to accumulate savings, etc....

More specifically, we consider an individual who has reached age $a_R - 1$, has wealth level A , faces some market opportunities this period described by some

set of labor income-hours pairs denoted by Ω with generic element $\{y_i, h_i\}$ and anticipates that he/she is going to retire in the following period, so that retirement is anticipated at least one period in advance.⁵ As we will see below, the main equation that we derive will not require that the decision to retire is optimal in any particular sense. We assume that there is no uncertainty that is revealed between age $a_R - 1$ and age a_R , though we do not restrict in any way the nature of uncertainty that the individual faces beyond age a_R . We discuss how shocks might impact our results in a later section. Lastly, we assume that the individual faces an interest rate of i that applies to any movement of purchasing power between these two periods.

Our strategy is to take choices for observed hours and labor income at age $a_R - 1$ given by (\bar{h}, \bar{y}) , and the fact that retirement occurs at age a_R as given, without imposing that these choices satisfy any particular set of first order conditions. Because we do not model the choices for these variables, we are in principle throwing away some information. We do this because we feel it is more complicated to specify all of the details that influence these decisions, like, for example, the nature of the set Ω and its potential interaction with private pension programs. The downside to this strategy is that we do not separately identify γ and η . The upside is that we derive a joint restriction on the two parameters that is robust to the various details that might influence these other choices.

Given values for \bar{h} , \bar{y} and assuming that retirement is anticipated to take place

⁵We parameterize the hours and income choices in this abstract fashion to emphasize that our derivation allows for various rigidities in the opportunity set for workers. In fact, this feature plays no role in the equation that we derive.

next period, we focus on the optimal choices for hours of home production and consumption expenditure during the last period of work and the first period of retirement. To simplify notation, we will use h_w to represent home production time in the last period of work and h_r to represent home production in the first period of retirement, and similarly for c and g . It follows that the optimal solutions for these values will be solutions to:

$$\max u(c_w) + \frac{\alpha}{1 - \frac{1}{\gamma}}(1 - \bar{h} - h_w)^{1 - \frac{1}{\gamma}} + \beta[u(c_r) + \frac{\alpha}{1 - \frac{1}{\gamma}}(1 - h_r)^{1 - \frac{1}{\gamma}}] + \beta^2 \dots$$

$$s.t. g_w + \frac{g_r}{1 + i} + \dots = \bar{y} + \frac{V(\bar{h}, \bar{y})}{1 + i} + A + \dots$$

$$c_j = [\theta g_j^{1 - \frac{1}{\eta}} + (1 - \theta)h_j^{1 - \frac{1}{\eta}}]^{\frac{\eta}{\eta - 1}}, j = w, r$$

$$\bar{h} + h_r \leq 1$$

The function V captures any transfer payments that the individual might receive at age a_R and allows them to depend explicitly on the choices of hours and income at age a_{R-1} , as well as implicitly on the previous history of the individual, since the function V is defined conditional on the situation of the individual at age $a_R - 1$. In particular, this function would capture both Social Security and private pensions.

Letting μ denote the Lagrange multiplier on the budget equation, and assuming an interior solution for home production time in both periods, we obtain the following four first order conditions:

$$g_w : \theta u'(c_w) c_w^{\frac{1}{\eta}} g_w^{-\frac{1}{\eta}} = \mu \quad (2.3)$$

$$g_r : \beta \theta u'(c_r) c_r^{\frac{1}{\eta}} g_r^{-\frac{1}{\eta}} = \frac{\mu}{1+i} \quad (2.4)$$

$$h_w : (1-\theta) u'(c_w) c_w^{\frac{1}{\eta}} h_w^{-\frac{1}{\eta}} = \alpha (1 - \bar{h} - h_w)^{-\frac{1}{\gamma}} \quad (2.5)$$

$$h_r : (1-\theta) u'(c_r) c_r^{\frac{1}{\eta}} h_r^{-\frac{1}{\eta}} = \alpha (1 - h_r)^{-\frac{1}{\gamma}} \quad (2.6)$$

Divide (2.3) by (2.4) to get:

$$\left[\frac{g_w}{g_r} \right]^{\frac{1}{\eta}} = \frac{u'(c_w)}{u'(c_r)} \left[\frac{c_w}{c_r} \right]^{\frac{1}{\eta}} \beta (1+i) \quad (2.7)$$

Divide (2.5) by (2.6) to get:

$$\left[\frac{1-h_r}{1-\bar{h}-h_w} \right]^{\frac{1}{\gamma}} \left[\frac{h_w}{h_r} \right]^{\frac{1}{\eta}} = \frac{u'(c_w)}{u'(c_r)} \left[\frac{c_w}{c_r} \right]^{\frac{1}{\eta}} \quad (2.8)$$

Dividing (2.7) by (2.8), taking logs and rearranging gives:

$$\frac{\gamma}{\eta} = \frac{\log(1-h_r) - \log(1-\bar{h}-h_w) + \log(\beta(1+i))}{\log(g_w/g_r) - \log(h_w/h_r)} \quad (2.9)$$

For our benchmark results we focus on the case in which $\beta(1+i)$ is assumed to equal unity, giving:

$$\frac{\gamma}{\eta} = \frac{\log(1 - h_r) - \log(1 - \bar{h} - h_w)}{\log(g_w/g_r) - \log(h_w/h_r)} \quad (2.10)$$

Equation (2.10) is the relationship that will be the focus of our analysis in the next section. Given values for all of the variables from the right hand side, this expression pins down the relative value of the two key preference parameters, γ and η .

In the next section we explore the quantitative implications of this relationship. But before doing so it is of interest to discuss the intuition associated with this expression. Equation (2.10) tells us what combinations of values are required for γ and η in order to rationalize observed values for the right hand side variables. Since the set of such values involve a constant ratio, it follows that a higher value of the intertemporal substitution of leisure necessarily implies a higher value of the substitution between time and goods. Why is this?

Intuitively, taking \bar{h} and h_w as given, when an individual retires they have more time available that must be allocated between increased leisure and increased home production. The greater is η , the easier it is for the individual to use this time in home production to substitute for market goods. Intuitively, therefore, a higher value of η will create greater incentives for the individual to allocate the additional time to home production relative to leisure. If we want to hit the same targets for time allocation, we would need to adjust γ in such a way that undoes the incentive for additional home production. It turns out that this is accomplished by increasing the value of γ . To see why a higher value of γ is associated with less incentives for home production, note that the higher the value of γ , the less

the individual desires a smooth profile for leisure. Put somewhat differently, the marginal utility from additional leisure declines less rapidly as γ is increased. It follows that a higher value of γ increases the incentive to take additional leisure, thereby decreasing the incentive for additional home production.

Before proceeding to the quantitative analysis, we think it is important to note some advantages associated with our method for imposing discipline on preference parameters relative to more standard methods used in the literature. Following MaCurdy (1981), many researchers have studied changes in hours of work and wages for continuously employed individuals to estimate the value of γ . The subsequent literature has shown that these estimates can be very sensitive to assumptions about human capital accumulation (Imai and Keane (2004), Wallenius (2011)), credit constraints facing younger workers (Domeij and Floden (2006)), and constraints on individual choices (Chang and Kim (2006), Rogerson (2011)). In contrast, our method is robust to allowing for all of these features.

3. Changes in Time Use and Consumption Expenditure at Retirement

Ideally we would estimate equation (2.10) using individual level panel data, but such data is unfortunately not available. We will instead adopt the strategy of evaluating the right hand side of equation (2.10) using what we view as “average” values for individuals that move from full time work to retirement. In this section we describe how we determine empirically reasonable values for the changes in time use and consumption expenditure at retirement. Since the existing litera-

ture provides estimates of the change in hours of market work and the change in consumption expenditure at retirement, our primary focus in this section will be to estimate the change in time devoted to home production at retirement.

3.1. Data

Because time spent in home production is a key data input into our exercise, our main data source will be the ATUS, since it is the highest quality data set on time use. We use all waves of the ATUS from the years 2003-2011, and pool all of the observations for both males and females between the ages of 60 and 70.⁶ We aggregate the time use data into four main categories: market work (MW), home production (HP), leisure (L), and personal care (PC). Market work includes hours spent doing the specific tasks required of one's main or other job, regardless of location. Work-related activities include activities that are not obviously work but are done as part of one's job, such as having a business lunch or playing golf with clients.⁷ Commuting time to and from work is included. Home production is defined to include time spent on housework, food and drink prep, presentation and cleanup, interior maintenance, repair and decoration, exterior repair, maintenance and decoration, lawn, garden and houseplants, animals and pets, vehicles, appliances, tools and toys, household management and shopping, which in turn

⁶Although in what follows we report statistics based on the sample for both males and females, we note that our key results also hold if we restrict attention to the sample of males only.

⁷Time spent on job search is also included in this category. Given that our framework does not attribute any role to job search one might want to exclude this category. It turns out that this category is sufficiently small, on average accounting for xx hours per week, that it does not influence any of our findings.

includes time spent purchasing consumer goods, groceries, professional and personal care services, financial services and banking, medical services, household services, home and vehicle maintenance (not done self) and government services, plus the time spent commuting to make these purchases. We define leisure to include time spent socializing and communicating, attending and hosting social events, relaxing and leisure, arts and entertainment, sports and exercise, as well as eating and drinking. Personal care consists of sleeping, grooming, health related self-care and personal activities. These four categories collectively account for over 90% of the total time allocation, and this fraction is very stable across the age range. The remaining categories include time spent caring for household and non-household members, time spent in educational activities, time spent on organizational, civic and religious activities and time spent communicating (i.e., phone, mail, email). These categories are not so easily allocated among the four main categories that we focus on, but for our purposes this is not an issue since the overall contribution of each group is both small and stable across age. For all of our tabulations and regressions we use the sample weights as provided by the ATUS.⁸

Table 1 shows the allocation of weekly hours by age for each of the four main categories mentioned above plus the residual. The first feature to note is that market work drops by almost 80 percent between the ages of 60 and 70. Of

⁸All data was extracted using the ATUS-X extractor. (Katharine G. Abraham, Sarah M. Flood, Matthew Sobek, and Betsy Thorn. 2011. American Time Use Survey Data Extract System: Version 2.4 [Machine-readable database]. Maryland Population Research Center, University of Maryland, College Park, Maryland, and Minnesota Population Center, University of Minnesota, Minneapolis, Minnesota. <http://www.atusdata.org>)

this drop, roughly half occurs between the ages of 60 and 63, which includes the early retirement age, with the other half coming between the ages of 63 and 66, a range that includes the full retirement age for these individuals and the age at which individuals become eligible for Medicare. Concurrently, we see that leisure increases quite markedly, while time devoted to home production also increases, but by somewhat less. There is a modest increase in time devoted to personal care, and as previously noted, the residual category is both small and stable, especially so for the 60-66 age range.

The model that we developed in the previous section emphasized transitions from working to not working. Data in the ATUS allows us to confirm that this is indeed the dominant dynamic behind the aggregate change in hours by age. Specifically, respondents in the ATUS report whether they currently work full time (at least 35 hours per week), part time (less than 35 hours per week) or are not working. Table 2 shows how the proportions of full time workers and non-workers changes with age.

The table shows that as we move from age 60 to age 70, the fraction of individuals working full time decreases by 40 percentage points, while the fraction not working increases by this same exact amount. From this we conclude that the movement from full time work to not working is indeed the dominant dynamic. As a check on these patterns found in the ATUS, we also report similar statistics based on data from the CPS Annual March Survey, which allows us to compute a measure of hours worked last year.⁹ We pool data from the surveys from 2004-

⁹We compute this as the product of weeks worked last year with usual hours per week. We apply sample weights provided by the CPS. The source for the CPS data was Miriam King,

2012, so that the data refers to 2003-2011, the same period covered by the ATUS. We label all of those workers who have annual hours worked of at least 1750 as full-time, and all of those with positive annual hours worked but less than 1750 as part-time. The patterns are remarkably similar, though not surprisingly, the CPS shows a slightly smaller fraction of individuals who report no work and a higher fraction who report working part-time, since it refers to an entire year whereas the ATUS is at a point in time. But notably, in both data sets virtually all of the aggregate change is accounted for by the movement of individuals from full-time work to no work.

3.2. Empirical Evidence on Life Cycle Allocations

We begin with evidence on the drop in consumption expenditure at retirement, since we will simply take this value from the literature. Aguiar and Hurst (2005) emphasize that a drop in consumption expenditure is not the same as a drop in true consumption, a distinction that is captured by our specification, since the drop in consumption expenditure corresponds to g_R/g_{R-1} , whereas the change in the flow of consumption corresponds to the ratio c_R/c_{R-1} . Our reading of the literature suggests that the drop in consumption expenditure at retirement is in the neighborhood of 15% (see, for example the estimates in Laitner and Silverman (2005) and the references contained therein). To the extent that part of the decrease in consumption expenditure represents fixed consumption costs associ-

Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick. Integrated Public Use Microdata Series, Current Population Survey: Version 3.0. [Machine-readable database]. Minneapolis: University of Minnesota, 2010.

ated with work that do not generate utility, the target in our model should be adjusted appropriately.¹⁰ Also, allowing for nonseparability between consumption and leisure, as in Laitner and Silverman (2005), would induce an additional channel through which consumption of market goods would drop at retirement. From this we conclude that .85 should be considered a lower bound for g_R/g_{R-1} . For our benchmark specification we assume $g_R/g_{R-1} = .90$, but we also consider a value of .85. As shown below, lower values of this ratio reflect conservative choices given our findings.

Next we consider time allocations. In the model we normalized the time endowment to equal one, which we will interpret as the amount of discretionary time that an individual has. We will in turn view discretionary time as the total time allocated to market work, home production and leisure. For individuals aged 60-66 in our pooled ATUS data, the average of this value is 92.92 hours per week. To calibrate \bar{h} and h_w we compute average values of market work and home production for individuals aged 60-64 who report that they work full time. These values are 44.27 and 16.37 hours per week, respectively, implying that $\bar{h} = .4764$ and $h_w = .1762$.

The next value that we need to assign is the value of h_r . This value will turn out to be quite important in the calculations that follow, and since it is the most challenging to measure given the limitations of the ATUS, we will devote a considerable amount of space to it. Given values for \bar{h} and h_w , we will find it

¹⁰Specifically, if there are fixed costs \bar{g} associated with working that do not generate utility, then our expressions would all involve $g_{R-1} - \bar{g}$ instead of g_{R-1} and we would be targeting $g_R/(g_{R-1} - \bar{g})$ rather than g_R/g_{R-1} .

useful to parameterize h_r by thinking about the quantity $f_{HP} = (h_r - h_w)/\bar{h}$, i.e., the fraction of the increased time available at retirement that is allocated to home production. In what follows we will aim to estimate f_{HP} and then compute h_r as $h_w + f_{HP} \cdot \bar{h}$.

One would like to have micro panel data on time use in order to estimate f_{HP} , but the ATUS does not have a panel component.¹¹ Instead, as is often done when faced with this situation, we will form a synthetic cohort from the pooled cross-section data and focus on the changes in market work and home production across age groups. As we noted previously, the dominant source of changes in market work as we move from age 60 to older ages is accounted for by a net movement of individuals from full-time work to no work. By looking at how time allocations vary with age we can trace out how the time made available by the decrease in market work is allocated to the other time use categories, and use this average variation to estimate f_{HP} . Specifically, using the data in Table 1 we will run the following regression:

$$HP_a = \text{constant} - f_{HP}MW_a + \varepsilon_a \tag{3.1}$$

for different age ranges, where a denotes age. We include the negative sign in the regression since decreases in market work are associated with increases in time devoted to home production, and as noted above, we want to interpret f_{HP} as a

¹¹In the next subsection we report an estimate using limited panel data on time use that is part of the Health and Retirement Survey. We do not use this as our benchmark estimate since the quality of the time use data seems substantially worse than in the ATUS. Nonetheless, we will see that it produces an estimate that is very similar to our benchmark.

fraction.

While this is our preferred method for estimating f_{HP} , we also report results based on cross-sectional variation in market work at each age. Ex ante there is cause for concern that this method could suffer from large selection effects—the home production time of those individuals who work few market hours at a particular age may be a poor indicator for the home production and leisure times that individuals with high values for market work at that same age will engage in when they reduce their market hours at a later age. The selection effects could easily be imagined to go in either direction, depending upon whether people who drop out of market work at early ages are more inclined toward higher home production or higher leisure. We carry out this regression for each age, allowing us to gauge whether the results change significantly as the pool of non-workers increases, perhaps providing some evidence about the extent of selection effects. We carry out these regressions despite concerns about possible selection effects simply to check whether this method would provide any evidence for higher values of f_{HP} than those obtained using the synthetic cohort approach. As already noted, the sharpness of our key result will depend on the upper bound for f_{HP} .

We begin with results using the synthetic cohort approach. Our choice of an age range balances three considerations. First, we want to choose an age range that captures a large share of the movement into retirement. Second, a larger age range increases the amount of data. But third, as the age range increases, especially into older ages, the more one is concerned that there may be drift in preference parameters associated with aging. Balancing these concerns, our

preferred estimate is for the age range 60 – 66, since this is the smallest age range that contains the two key eligibility ages for Social Security and Medicare, thereby capturing the heaviest concentration of the movement into retirement.¹² Moreover, because of these eligibility thresholds, we believe it is reasonable to think that many of the movements from working to not working are anticipated at least one year ahead, as assumed in our derivation.

Running the regression in equation (3.1) using the ATUS data averaged by age in Table 1 for the age range 60-66 delivers a point estimate for f_{HP} of .267 with a standard error of .037. If we instead used the age ranges of 60-64 or 60-65 we get similar point estimates (.276 and .274 respectively), but with slightly larger standard errors (.077 and .048, respectively). On the other hand, if we extend the age range to 60-67 or 60-68 the point estimates drop (to .2227 and .2014 respectively) with standard errors of .050 and .045. In view of these results we will use the estimates for the age range 60-66 as our preferred estimate.¹³ Allowing for a two-standard error deviation, the largest value for h_r would correspond to $f_{HP} = .341$.

In the model of the previous section we assumed that home production and leisure are the only categories that the time freed up by retirement can be allocated

¹²As we extend the upper age limit beyond 66 we obtain decreasing absolute values of f_{HP} , which we interpret to possibly reflect the effect of drift in parameter values at older ages. We discuss this further later on in this section.

¹³Hamermesh and Donald (2007) carry out a somewhat similar exercise using the ATUS data from 2003 and 2004, contrasting the change in average market work and average home production for individuals aged 55 – 59 and 65 – 69. They find that 40 minutes of the roughly 170 minute decrease in market work is allocated to home production, implying an estimate of $\beta_{HP} = .23$, which is quite similar to our estimates when we extend the age range to either 67 or 68.

to, but this need not be the case in the data, since there are both changes in time devoted to personal care and there is also a residual category. If we repeat the above analysis but instead use the sum of leisure and home production on the left hand side, we find that the coefficient on market work for the age range 60-66 is .86, implying that 14 percent of the time freed up by the reduction in market work is being allocated to something other than leisure or home production. The vast majority of this (12 percentage points) is accounted for by an increase in time devoted to personal care, and all of this increase in personal care is in turn accounted for by an increase in sleeping time. Moving from age 60 to 66, this corresponds to a little less than two additional hours per week. To the extent that a little extra sleep at the margin might be viewed as discretionary and hence best included as a form of leisure, this observation might not call for any adjustment. In fact, holding age constant, we also find that a one hour decrease in market work is associated with an increase in sleep between .12 and .15 hours, thereby supporting this interpretation. Nonetheless, we will also consider the polar extreme case in which we interpret the additional time devoted to sleep as a decrease in the amount of discretionary time. This is easily incorporated into equation (2.10), by replacing the term $\log(1 - h_r)$ with $\log(1 - .14\bar{h} - h_r)$ to reflect that fact that the amount of discretionary time is reduced by 14% of the hours of work prior to retirement.

Next we present the cross-sectional estimates. We run regressions separately for each age, in order to assess whether there is any systematic change with age. When running this cross-section regression of home production time on (the negative of) market work we also include controls for gender and four educational

groups (less than high school, high school, some college, and college and above).¹⁴ Table 3 presents the results.

The point estimates are all in the range of .25 to .30, just as in the first exercise, though the standard errors are now much smaller. There is little evidence of a substantial systematic change with age, despite the fact that the older cross-sections will have a much larger fraction of recent retirees. For our purposes, the key result from Table 3 is that the cross-section regressions do not provide evidence for larger absolute values of f_{HP} . In particular, given that we will be considering values of f_{HP} as large as .341, the cross-section results do not call for consideration of larger values.

3.3. Evidence From Other Sources

The key advantage of the ATUS is that it is widely regarded as the highest quality data on time use. But for our purposes, a key limitation of the ATUS is that it does not contain a panel component. In this section we report results presented in Hurd and Rohwedder (2005) based on the Consumption and Activities Mail Survey (CAMS) which is a supplemental survey sent to a random subsample of roughly 5000 individuals within the Health and Retirement Survey (HRS). This survey is biannual and contains a panel component. The results in their paper are based on the 2001 and 2003 waves. While this survey does contain a panel component and provide data on time use, the quality of the time use data is of much lower quality than in the ATUS. For example, time use information in

¹⁴We note that including these controls has very little effect on either the point estimates or standard errors.

CAMS is based on recall as opposed to a time diary, and there are many missing values in the survey.¹⁵ Nonetheless, Hurd and Rohwedder (2007) argue that the CAMS time use data still match the salient features present in the ATUS.

Hurd and Rohwedder (2005) consider those individuals who report not being retired in the first wave of CAMS and retired in the second wave of CAMS and were between the ages of 52 and 71 in 2002, providing a sample of 201 individuals. They consider seven categories of time use that they think represent activities which are substitutes for spending on market goods and use this as their definition of home production. These categories are house cleaning, washing/ironing, yardwork/gardening, shopping, meal preparation, money management and home improvements. For this sample they report the change in work for pay and home production time between the two surveys, separately for males and females. The reduction in work for pay among the males was 21.25 hours, and the increase in time devoted to home production was 6.02 hours.¹⁶ For females the two values were 17.42 and 5.15. These produce estimates for f_{HP} of .28 and .30. These values are virtually identical to the those that we reported above using the two different procedures on the ATUS.

In a different but related exercise, Aguiar et al (2013) use the ATUS to ask how the decrease in market work during the most recent recession has been allocated to other time use categories. They conclude that roughly 30 percent of foregone market work during the Great Recession was allocated to home production, and

¹⁵We refer the reader to Hurd and Rohwedder (2005, 2007) for more detail on this survey.

¹⁶Interestingly, Hurd and Rohwedder (2007) find that the cross-sectional difference in time devoted to home production between retired and non-retired people is very similar to the difference found in the panel data when an individual moves from working to retired.

another 50 percent was allocated to leisure. There is no reason that these values should be the same as ours, given that the context and the samples are different. For example, one would expect that an increase in time devoted to job search might be more relevant in their context, though it turns out that this magnitude is not that large. Nonetheless, it is interesting to note that they obtain responses that are quite similar to ours. We will discuss some implications of this later in the paper.

4. Results

Given values for \bar{h} , g_r/g_w , h_w , and h_r , equation (2.10) tells us the value of γ/η that is required to support those values and therefore implicitly defines a one dimensional curve in (γ, η) space. Given that there are distinct literatures that provide estimates for each of these parameters individually, this calculation can shed light on the extent to which the existing ranges of estimates are consistent, and whether this joint restriction imposes additional empirical discipline on the two elasticity parameters.

The literature on estimating γ is large and has been the subject of some controversy. The early literature, including MaCurdy (1981), Browning et al (1985) and Altonji (1986), typically found relatively small estimates, less than .40 and sometimes close to zero.¹⁷ Chetty (2012) offers a synthesis of this literature and

¹⁷Note that in our formulation, γ measures the intertemporal elasticity of substitution for leisure, whereas many studies estimate the intertemporal elasticity of substitution for labor supply. If labor supply and leisure are approximately equal, as is the case for many groups, then the two are basically the same. Given we are discussing some broad ranges of estimates as opposed to specific values, we will not focus on this distinction.

suggests that estimates on the order of .4 to .5 are probably most reasonable. We note that using a very different source for identification, Pistaferri (2003) obtained estimates as high as .75. A recent literature has also investigated the extent to which these estimates of γ are biased due to neglected model features, such as incomplete markets for borrowing and lending, and the accumulation of human capital. Accounting for these features has produced estimates of γ that are as large as 1.4.¹⁸ We also note that these estimates are all obtained in specifications that abstract from home production. Rupert et al (2000) have argued that abstracting from home production in life cycle models can lead to a downward bias in estimates of γ .

Though the literature on estimating η is much smaller, there are several estimates of η in the literature. The most appropriate estimate would seem to come from Aguiar and Hurst (2007), who explicitly considered older individuals, though they focused specifically on how shopping time could substitute for expenditure. For their benchmark specification they report estimates in the range of 2.0 to 2.5.¹⁹ Using aggregate data, McGrattan, Rogerson and Wright (1997) find a value of η in the range of 1.67 to 1.8, while Chang and Schorfheide (2003) find a value in the range of 1.8 to 2.5. These estimates all suggest relatively high values for η . While there is no logical reason that η cannot be very small, we think the essence of the

¹⁸In fact, Imai and Keane (2004) obtained estimates larger than 3 when incorporating human capital accumulation and using data for individuals aged 20-36. But more recently, Wallenius (2011) shows that the estimates are significantly reduced when a longer life cycle profile is used in the estimation.

¹⁹In an earlier paper and also using micro data, Rupert, Rogerson and Wright (1995) obtain several estimates, some of which had large standard errors. Their preferred regression produced an estimate of 1.8.

idea of home production is that it is a substitute for market goods, suggesting that one might reasonably consider unity as a reasonable lower bound for the value of η . We will see below that even this weak restriction is somewhat informative.

Table 4 presents implied values for γ for specific values of η given our benchmark values for the values on the right-hand side of equation (2.10). A striking message emerges from Table 4: viewed from the perspective of our simple choice problem, changes in allocations at retirement are not jointly consistent with what many would view as “consensus” estimates for the two preference parameters γ and η . That is, values of γ around .4 and η around 2 are not mutually consistent in this context. Moreover, the extent of this tension is dramatic: even if one contemplates values of η as low as 1.00, we see that the benchmark specification would still require a value of γ that exceeds unity. Only if one takes values of γ from the very upper end of estimates (say 1.4), and estimates of η from the very bottom end of the range of estimates (say 1.5), can one come close to reconciling the data in the context of our model and be jointly consistent with at least some estimates for the two preference parameters.

Next we consider how sensitive these results are to alternative values for some of the objects on the right hand side of equation (2.10). If we consider smaller values for g_r/g_w the main message remains relatively unaltered: if we use $g_r/g_w = .85$ instead of .90 as in the benchmark, the ratio γ/η is reduced from 1.07 to .99, so that the values of γ in Table 4 would be reduced by roughly 8 percent. This does not have a first order impact on the message.

In contrast, changes in f_{HP} can matter quite a bit. In Table 5 we report

the locus of (γ, η) pairs when we use $f_{HP} = .3414, .40,$ and $.45$. The first value represents a two standard error deviation from the value used in Table 4, while the other two values are used to illustrate the importance of this parameter. As the table indicates, higher absolute values of f_{HP} ease the tension between the available estimates for the two elasticity parameters. Nonetheless, assuming that $f_{HP} = .3414$ still generates very little overlap between existing ranges for estimates, and even with $f_{HP} = .45$, one would have to consider values of η much below one (in fact, $.64$) in order to rationalize a value for γ of around $.4$. Conversely, even when $f_{HP} = .45$, one would need a value for γ that is near unity in order to rationalize $\eta = 1.50$. In short, even assuming values of f_{HP} that are substantially larger than those implied by our analysis of the data does not change the main message.

It is also instructive to ask what value for f_{HP} would rationalize $\gamma = .50$ and $\eta = 1.5$, which we have previously suggested are commonly adopted values. The answer turns out to be $f_{HP} = .67$. To rationalize a value of γ that is one half as large as η would require $f_{HP} = .53$.

Lastly, we consider the effect of incorporating a change in the amount of discretionary time when the individual moves from working to not working, as discussed in the previous section. Specifically, we assume that the amount of discretionary time is reduced by the amount of $.14 \cdot \bar{h}$ when moving from working to not working. Results are reported in Table 6. While the effect of this change is quantitatively significant, it does not alter the main message. In this regard, note that considering the same reduction in discretionary time with $f_{HP} = .3414$ is somewhat

extreme, since we are implicitly assuming that more time is allocated to home production without changing the implications for discretionary time.

In summary, the main message that we take away from the above calculations is that viewed through the lens of our simple decision problem, the changes in allocations at retirement are inconsistent with the joint prior that γ is around .40 and η is around 2.0, which are values that one might extract as “consensus” values from readings of the two respective literatures on estimating these elasticities. Moreover, even if one believes existing estimates of η are biased upward by a substantial amount, so that the true value is much closer to 1.0 than it is to 2.0, one would still need a value of γ above .70 and possibly above unity.

5. The Effect of Shocks

The framework that we developed in Section 2 considered how an individual’s choices for consumption expenditure and home production time change at retirement taking as given that the individual knew that he or she was going to retire at least one period in advance and that there were no changes in preferences or other shocks over these two periods. In this section we discuss two plausibly important reasons for concern with this assumption and how they might influence our conclusions. The first of these is health shocks that influence the disutility associated with working, both in the market and at home. The second concerns the possibility of income shocks that are correlated with retirement.

5.1. Health Shocks

If an individual retires from market work due to a health shock, it is possible that this shock would also cause the individual to do less home production in retirement. If this is the case, our estimate of f_{HP} may significantly underestimate the extent to which an individual substitutes from market work into home production in the absence of health shocks. Note that in terms of the theoretical analysis in Section 2, the case of an individual who experiences a one-time permanent health shock at some earlier age, but who then continues to work for some time before retiring does not constitute an issue for the theoretical derivation, even if the health shock does lead them to retire sooner than they would have in the absence of the health shock. The reason for this is that our derivation did not place any restrictions on the shocks that might have occurred prior to the ultimate period of work. If the health shock is associated with an ongoing deterioration of health status it would be an issue for our framework.

We present several pieces of evidence to help us address the potential significance of health shocks. First, we note that in their analysis using the Health and Retirement Survey, Blau and Shvydko (2011) find that less than 14% of the transitions from employment to retirement among individuals aged 52-71 are associated with a change in self-reported health status from “good” to “bad”. While it is the case that employed individuals who experience such a change are almost twice as likely to retire as those who remain in good health, these cases account for only a small fraction of retirements.

Next we address this issue using the ATUS. For the years 2006-2008 the ATUS

includes questions on health status. In particular, in each of these survey years individuals were asked to self-report their health status on a scale of 1 to 5, with 1 denoting excellent health and 5 denoting poor health. The first thing we examine is the extent to which (self-reported) health status deteriorates between the ages of 60 and 66, the interval of ages that we used to estimate f_{HP} . It turns out that the distribution of individuals across health states is quite stable over this age range. For example, defining “healthy” to correspond to a health status of 1, 2 or 3, Table 7 shows the fraction of the sample that is healthy at each age.

While there is a decrease going from 60 to 61, there is virtually no change as we go from age 61 to 66.²⁰ Note that the distribution of health status could remain stable at the same time that individuals experience deteriorating health status if there is higher mortality from the worst health states. However, this would not be a problem with our synthetic cohort estimates as long as the distribution of health status is not changing. While not definitive, this is at least suggestive that deteriorating health is not a key driving force over this age range.

To pursue this further, we have repeated the earlier analysis used to estimate f_{HP} using only those individuals who report a health status of 3 or better. Results are presented in the Appendix; here we simply summarize them. To begin, we note that Table A1 in the appendix shows that the aggregate changes in time use are very similar for the healthy subsample and for the entire sample used earlier. One issue with restricting attention to the healthy subsample is that the sample size shrinks dramatically, by more than a factor of 4. This reflects the fact

²⁰This is also true if we consider a higher threshold for what constitutes healthy.

that the sample shrinks by roughly a factor of three when we restrict attention to the years 2006-2008, and that we additionally lose about 30 percent of these observations when we restrict attention to the healthy subsample. Perhaps not surprisingly, with dramatically smaller sample sizes, the averages by age seem to exhibit substantially more noise, which will lead to much higher standard errors on the estimates.

Table A2 in the Appendix provides synthetic cohort estimates of f_{HP} using the healthy subsample. The point estimates are very similar to those obtained using the whole sample, though the standard errors are much larger. Table A3 in the Appendix provides cross-sectional estimates of f_{HP} using the healthy subsample. While some of the point estimates are slightly higher than for the whole sample, the average values are quite similar. In the previous subsection we reported calculations assuming that $f_{HP} = .40$, and none of the point estimates obtained for the healthy subsample would support using a higher value.

The data in Hurd and Rohwedder (2007), again based on the CAMS, shed some light on why health status does not seem to play such a large role in our estimates. They compare work for pay and home production time for two age groups, 50-64 and 65-74, for three different health statuses: excellent or very good, good, and fair or poor. They find that for the age group 50-64, those in fair or poor health are already doing very little work: less than 40 hours per month on average, as compared to more than 100 hours per month on average for the other two groups. Moreover, they do about the same amount of home production time. For the 65-74 group, differences in work for pay shrink considerably as those

in better health move into retirement, but again, differences in home production time are minimal.

Based on the health status information in the ATUS, we believe that our earlier conclusions appear to be robust to considering changes in health. This is subject to the caveat that the sample sizes are small and the health status information is somewhat crude.

5.2. Income Shocks

A second issue that we consider is the possibility that an individual might retire earlier than was anticipated the previous year, due perhaps to the fact that they experienced a job loss and retired earlier than originally planned. The specific situation that we have in mind is someone who planned to work for one or more additional years beyond the current year, but unexpectedly loses their (good) job in a plant closing and then retires rather than take a much lower paying job that is their next best option in the labor market. Note once again that our analytic result is valid if, for example, an individual suffers job loss at age 60 and then works at some lesser job until he or she reaches the early retirement age, at which time he or she retires, since we did not make any assumptions about what shocks might have occurred prior to what we called age $a_R - 1$. This is true even if the job loss led them to retire earlier than they would have had they not suffered the job loss.

The ATUS does not allow us to estimate the prevalence of this type of event. Nonetheless, we argue here that in this situation our calculation will actually

underestimate the appropriate value of the right hand side of equation (2.10). The reason for this is that if someone loses their job and retires earlier than previously planned, this individual will presumably retire with a lower level of wealth than previously planned. Given the lower level of wealth the individual will adapt by choosing a lower level of consumption expenditure spending more time in home production than previously planned, each of these are ways to economize on expenditure. In this case, the data on changes in time spent in home production would show a larger increase than if the individual had retired as planned, and a larger drop in consumption of market goods. The appropriate values to plug into the right hand side of equation (2.10) are the values that the individual planned in light of not incorporating the reality of the job loss and the decision to retire early. As shown in the sensitivity analysis in the previous section, both of these adjustments would actually lead to even higher values of the right hand side of (2.10), and since our main message is only reinforced in this case, we conclude that our message is robust to these types of shocks playing a large role.

In our earlier discussion of health shocks we focused on the possibility that health shocks lead to changes in preference parameters. But health shocks that lead to premature retirement are likely to also be associated with lower levels of wealth, and thus will also generate these types of effects that bias findings in the other direction.

6. Nonconvexities in the Utility from Leisure

In the previous analysis we have assumed that retirement does not affect the manner in which individuals experience utility from leisure. While this seems to be a natural benchmark, some authors in the literature (see, for example, Hamermesh and Donald (2007)) have suggested that the state of retirement may influence the utility flow associated with leisure. Market work typically involves restrictions on the timing of market work, and as a result restricts the timing of leisure. Some leisure activities, such as reading the morning newspaper during a leisurely breakfast, are associated with specific times, and may not be feasible when an individual is working full time. Similarly, retirement allows individuals much more freedom to travel, also providing them access to additional leisure activities and effectively increasing the marginal benefit of leisure. Alternatively, perhaps the stress associated with working makes it difficult for individuals to enjoy their leisure time.

The possibility that the marginal utility of leisure increases when individuals move from full time work into retirement is of particular interest for the calculations that we have carried out, since any factor that leads to greater marginal utility of leisure upon retirement will create a force for an increase in leisure upon retirement, and thereby lessen the tension between relatively low labor supply elasticities and the increase in leisure upon retirement. In this section we attempt to quantify the potential impact of this channel.²¹

²¹Some readers may also consider plausible stories for why the marginal utility from leisure is less during retirement. We do not focus on this possibility for the simple reason that this possibility would only strengthen our previous conclusions.

A simple way to capture these types of possibilities in our framework is to assume that the marginal utility of leisure is higher during retirement, i.e., when market work is zero as opposed to “full time”. If we assume that $\gamma > 1$, then a simple way to capture this in our discrete choice model of market work is to modify the period utility function from our earlier analysis to:

$$\alpha \log(c) + (1 + I_R D)(1 - \alpha) \frac{1}{1 - \frac{1}{\gamma}} (1 - h)^{1 - \frac{1}{\gamma}}$$

where $D \geq 0$ and I_R is an indicator function that takes the value 1 if the individual chooses zero hours of market work.

Our previous analysis represents the special case in which $D = 0$. Here we consider the case in which $D > 0$. It remains true that the optimal life-cycle profile is described by the same five values as previously: h_w , h_r , e , g_w , and g_r . However, the lifetime utility associated with these choices is now given by:

$$e[u(c_w) + \frac{\alpha}{1 - \frac{1}{\gamma}} (1 - \bar{h} - h_w)^{1 - \frac{1}{\gamma}}] + (1 - e)[u(c_r) + \frac{\alpha(1 + D)}{1 - \frac{1}{\gamma}} (1 - h_r)^{1 - \frac{1}{\gamma}}]$$

where c_w and c_r are as defined earlier.

One way to interpret the magnitude of D is to evaluate by what fraction you could decrease the amount of leisure that the individual enjoys as they move into retirement to generate the same utility flow that they receive from leisure when working. Specifically, we solve for the value Δ that solves:

$$\frac{\alpha}{1 - \frac{1}{\gamma}} ((1 - \bar{h} - h_w))^{1 - \frac{1}{\gamma}} = \frac{\alpha(1 + D)}{1 - \frac{1}{\gamma}} (\Delta(1 - \bar{h} - h_w))^{1 - \frac{1}{\gamma}}$$

This yields:

$$\Delta = (1 + D)^{-1/(1-\frac{1}{\gamma})}$$

For example, if weekly hours of leisure were equal to 40 when working and D and γ are such that $\Delta = .5$, then it would only require 20 hours of leisure when retired to produce the same utility flow as the 40 hours of leisure when working. We will make use of this mapping in what follows. Note that the value of Δ that one obtains from this expression is independent of the level of leisure when working.

Proceeding exactly as in section 2, one obtains the following analogue of equation (2.10) :

$$\frac{\gamma}{\eta} = \frac{\log(1 - h_r) - \log(1 - \bar{h} - h_w)}{\log(g_w/g_r) - \log(h_w/h_r) + \eta \log(1 + D)} \quad (6.1)$$

Given a value for D , we can again map out a locus of (γ, η) pairs that are consistent with observed values for g_w/g_r , h_w , h_r , and \bar{h} .

Implementing this procedure obviously requires a value for D . We are aware of no empirical evidence on this issue, and so cannot offer any definitive assessment of how this possibility influences our overall conclusions. However, we think some illustrative calculations tied to values of parameter Δ are informative. Specifically, in what follows we assume the same targets as were used in Table 4, and assume that D is such that $\Delta = 2/3$, i.e., that it takes 50% more leisure time when working to generate the same utility flow as a given amount of leisure when retired. While we just noted that we have no empirical evidence on the reasonableness of this value, we feel this is a rather dramatic magnitude. Table 8 shows the implied

values of γ for various values of η , for $\Delta = 2/3$ in addition to the benchmark values from Table 4, which corresponds to $\Delta = 1$.

The basic pattern is that allowing for this type of nonconvexity in the enjoyment of leisure pushes the implied values of γ toward one, with the percentage effect being relatively larger for higher values of η . Whereas the magnitude of the reduction is quite substantial for higher values of η , the main message that we take away from this table is that even with what seems to be relatively large nonconvexities in the utility from leisure, the implied values of γ are still in excess of 1. We conclude that while nonconvexities of the sort considered here may be quantitatively significant, for the magnitudes we have considered they do not overturn the main implications of our previous analysis.

7. Conclusion

In a life cycle model that features home production and retirement, we show that the changes in home production time, leisure and consumption expenditure imposes a tight relationship between two key preference parameters that help to determine labor supply elasticities: the willingness of an individual to substitute leisure over time and the elasticity of time and goods in generating current utility. This relationship is robust to allowing for many features, such as human capital accumulation, borrowing constraints for younger workers, non-linear taxation and the presence of private pensions and social security. We estimate how allocations change at retirement using data from the recently available ATUS, and use these estimates to explore the quantitative implications of the restriction implied by our

model. We conclude that there is a sharp tension between what many researchers view as empirically reasonable values for the two elasticity parameters. Specifically, an intertemporal elasticity of around .40 and an elasticity of substitution between time and goods of around 2.0 are not at all consistent with the measures that we estimate in the data. One possible resolution is that the intertemporal elasticity of substitution for leisure is much greater, perhaps even larger than one, as suggested by some recent studies. Alternatively, the elasticity of substitution between time and goods might be much smaller than existing estimates, though rationalizing values of γ as low as .40 would require values of η substantially below unity, implying that time and goods are actually complements in production rather than substitutes. Of course, a third possibility is that the structure that we have imposed to derive our key restriction is missing some important feature. In this regard we do stress that we have imposed relatively little structure beyond functional forms.

Our general method can also potentially be used in other settings to derive similar restrictions. One promising possibility is the case of workers who move from employment to unemployment. Previous work has argued that a typical consumption drop for someone who moves from employed to unemployed is around 15%, and the recent work of Aguiar et al (2013) argues that roughly 30% of working time was allocated to home production in the recent recession. While the case of a worker who moves into unemployment will involve some additional factors relative to the case of retirement, like the need to engage in job search and the effect on future earnings, these figures suggest that this context may produce

similar findings to what we found in the context of retirement.

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] Becker, G. "A Theory of the Allocation of Time." *The Economic Journal* 75 (1965): 493-517.
- [6] Blau, D, and Shvydko, T. 2011. "Labor Market Rigidities and Employment Behavior of Older Workers." *Industrial and Labor Relations Review* 64, Article 3.
- [7] 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.

- [8] Burda, M., Hamermesh D., and Weil. P. 2008. "The Distribution of Total Work in the EU and US." in *Are Europeans Lazy? or Americans Crazy?*, edited by Tito Boeri, Oxford University Press, Oxford.
- [9] 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.
- [10] Chang, Y., and Schorfheide, F. 2003. "Labor Supply Shifts and Economic Fluctuations." *Journal of Monetary Economics* 50, 1751-1768.
- [11] Chetty, R. 2012. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." *Econometrica* 80(3): 969-1018.
- [12] 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.
- [13] Domeij, D., and Floden, M. 2006. "The Labor Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased." *Review of Economic Dynamics* 9, 242-262.
- [14] Freeman, R., and Schettkat, R. 2001. "Marketization of Production and the US-Europe Employment Gap." *Oxford Bulletin of Economics and Statistics* 63, 647-670.

- [15] _____ . 2005. “Marketization of Household Production and the EU-US Gap in Work.” *Economic Policy* 20, 6-50.
- [16] Gronau, R. 1977. “Leisure, Home Production and Work—The Theory of the Allocation of Time Revisited.” *Journal of Political Economy* 75, 1099-1124.
- [17] 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*.
- [18] Hamermesh, D., and Donald, S. 2007. “The Time and Timing Costs of Market Work.” NBER Working Paper #13127.
- [19] 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.
- [20] Hurd, M., and Rohwedder, S. 2003. “The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement.” NBER Working Paper #9586.
- [21] _____ . 2005. “Changes in Consumption and Activities in Retirement.” Michigan Retirement Research Center Working Paper 2005-096.
- [22] _____ . 2007. “Time-Use in the Older Population. Variation by Socio-economic Status and Health.” RAND Working Paper WR-463.

- [23] Imai, S., and Keane, M. 2004. "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review* 45, 601-641.
- [24] Laitner, J., and Silverman, D. 2005. "Estimating Life-Cycle Parameters from Consumption Behavior at Retirement." NBER Working Paper #11163.
- [25] MaCurdy, T. 1981. "An Empirical Model of Labor Supply in a Life Cycle Setting." *Journal of Political Economy* 89, 1059-1085.
- [26] McGrattan, E., Rogerson, R., and Wright, R. 1997. "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy." *International Economic Review* 38, 267-290.
- [27] Ngai, R. and Pissarides, C. 2008. "Trends in Hours and Economic Growth." *Review of Economic Dynamics* 11, 239-256.
- [28] _____ . 2011 "Taxes, Social Subsidies and the Allocation of Work Time." *American Economic Journal: Macroeconomics* 3, 1-26.
- [29] Olovsson, Connie. 2009. "Why Do Europeans Work So Little?" *International Economic Review* 50, 39-61.
- [30] Pistaferri, L. 2003. "Anticipated and Unanticipated Wage Changes, Wage Risk and Intertemporal Labor Supply." *Journal of Labor Economics* 21, 729-754.
- [31] Ragan, K. 2013. "Taxes and Time Use: Fiscal Policy in a Home Production Model." *American Economic Journal: Macroeconomics* 5, 168-192.

- [32] Rogerson, R. 2007. "Taxes and Market Work: Is Scandinavia an Outlier?" *Economic Theory* 32, 59-85.
- [33] _____. 2008. "Structural Transformation and the Deterioration of European Labor Market Outcomes." *Journal of Political Economy* 116, 235-259.
- [34] _____. 2009. "Market Work, Home Work and Taxes: A Cross Country Analysis." *Review of International Economics* 17,588-601.
- [35] _____. 2011. "Individual and Aggregate Labor Supply in a Model With Coordinated Working Times." *Journal of Money, Credit and Banking* 43, 7-37
- [36] Rogerson, R., and Wallenius, J. 2013. "Nonconvexities, Retirement and the Elasticity of Labor Supply." *American Economic Review* 103, 1445-1462
- [37] Rupert, P., Rogerson R., and Wright R. 1995. "Estimating Substitution Elasticities in Household Production Models." *Economic Theory* 6, 179-193.
- [38] _____. 2000. "Homework in Labor Economics: Home Production and Intertemporal Substitution." *Journal of Monetary Economics* 46, 557-579.
- [39] Rust, J., and Phelan, C. 1997. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets." *Econometrica* 65, 781-831.

- [40] 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.

Tables

Table 1

Time Use by Age in the ATUS 2003-2011 (hours/week)

Age	MW	HP	L	PC	Residual	#obs
60	26.17	20.18	47.65	64.27	9.73	1748
61	22.72	21.24	49.68	64.82	9.54	1694
62	22.08	22.02	49.56	64.18	10.16	1616
63	16.67	22.17	53.24	65.63	10.29	1609
64	15.96	23.79	53.88	64.65	9.72	1483
65	11.20	24.46	55.65	66.09	10.60	1520
66	11.28	24.27	56.93	65.97	9.55	1384
67	10.91	22.71	57.84	66.33	10.21	1292
68	9.18	23.46	58.31	66.46	10.59	1271
69	7.58	23.27	60.03	67.05	10.07	1176
70	6.73	23.96	58.84	67.61	10.86	1106

Notes: Values are averages for pooled ATUS surveys from 2003-2011, computed using sample weights. MW is market work, HP is home production, L is leisure, PC is personal care. Last column shows sample size for each age.

Table 2
Work Status by Age, CPS and ATUS 2003-2011

Age	ATUS			CPS		
	Not Working	Part-time	Full-time	Not Working	Part-time	Full-time
60	.38	.12	.50	.32	.18	.50
61	.44	.13	.43	.36	.17	.47
62	.50	.15	.35	.40	.19	.41
63	.54	.14	.32	.47	.19	.34
64	.58	.13	.29	.51	.19	.30
65	.65	.14	.21	.57	.18	.25
66	.67	.16	.17	.62	.17	.20
67	.72	.13	.15	.65	.18	.17
68	.72	.14	.14	.68	.17	.15
69	.75	.14	.11	.72	.15	.13
70	.78	.12	.10	.74	.15	.11

Notes: Data for ATUS is based on pooled surveys from 2003-2011, using sample weights. Data for CPS is based on pooled March CPS for 2004-2012 using data on weeks worked last year and usual hours per week, and using sample weights. In the ATUS we define full time to be at least 35 hours per week. In the CPS we define full time to be at least 1750 hours last year.

Table 3

Cross-Sectional Estimates of f_{HP} by Age (standard errors in parentheses)

60	61	62	63	64	65	66
.27(.01)	.26(.02)	.27(.02)	.28(.02)	.30(.02)	.26(.02)	.29(.02)

Notes: The table reports the negative of the coefficient on hours of market work per week with dependent variable being hours of home production per week. Controls for gender and four education levels (less than high school, high school, some college, at least college) are also included. Data is the pooled ATUS surveys from 2003-2011. Sample weights are used.

Table 4

Implied Values of γ For Benchmark Specification

	$\eta = 1.00$	$\eta = 1.25$	$\eta = 1.50$	$\eta = 2.00$	$\eta = 2.50$
$f_{HP} = .2674$	1.07	1.34	1.61	2.14	2.68

Notes: Table shows implied values of γ for several values of η based on equation (2.10) from the text for benchmark values of right hand side variables: $\bar{h} = .4764$, $h_w = .1762$, $g_r/g_w = .90$, $h_r = h_w + .2674 \cdot \bar{h}$.

Table 5

Effect of f_{HP} on Implied Values of γ

	$\eta = 1.00$	$\eta = 1.25$	$\eta = 1.50$	$\eta = 2.00$	$\eta = 2.50$
$f_{HP} = .3414$.85	1.06	1.27	1.70	2.12
$f_{HP} = .40$.72	.89	1.07	1.43	1.79
$f_{HP} = .45$.62	.78	.94	1.25	1.56

Notes: Table shows implied values of γ for several values of η based on equation (2.10) from the text for the following values of right hand side variables: $\bar{h} = .4764$, $h_w = .17362$, $g_r/g_w = .90$, $h_r = h_w + f_{HP} \cdot \bar{h}$.

Table 6

Implied Values of γ With Reduction in Discretionary Time					
	$\eta = 1.00$	$\eta = 1.25$	$\eta = 1.50$	$\eta = 2.00$	$\eta = 2.50$
$f_{HP} = .2674$.92	1.14	1.37	1.83	2.29
$f_{HP} = .3414$.71	.88	1.06	1.42	1.77

Notes: Table shows implied values of γ for several values of η based on equation (2.10) from the text for the following values of right hand side variables: $\bar{h} = .4764$, $h_w = .1762$, $g_r/g_w = .90$, $h_r = h_w + f_{HP} \cdot \bar{h}$, but assuming that total discretionary time in retirement is reduced from 1.00 to $1.00 - .14 \cdot \bar{h}$.

Table 7

Fraction Healthy by Age in the ATUS									
60	61	62	63	64	65	66	67	68	69
.76	.71	.70	.72	.71	.70	.71	.72	.73	.69

Notes: Table shows percent of sample with self reported health of 1, 2 or 3 in the pooled ATUS samples for 2006-2008, using sample weights.

Table 8

Implied Values of γ With Non-Convexities in Leisure					
	$\eta = 1.00$	$\eta = 1.25$	$\eta = 1.50$	$\eta = 2.00$	$\eta = 2.50$
$\Delta = 1$	1.07	1.34	1.61	2.14	2.68
$\Delta = 2/3$	1.04	1.19	1.31	1.51	1.65

Notes: Table shows implied values of γ for several values of η based on equation (6.1) from the text for the following values of right hand side variables: $\bar{h} = .4764$, $h_w = .1762$, $g_r/g_w = .90$, $h_r = h_w + .2674 \cdot \bar{h}$, for different values of Δ .

Appendix

Table A1

Time Use by Age in the ATUS (hours/week)

Age	Entire Sample					Healthy Subsample				
	MW	HP	L	PC	#obs	MW	HP	L	PC	#obs
60	26.17	20.18	47.65	64.27	1748	26.89	20.29	47.31	63.36	395
61	22.72	21.24	49.68	64.82	1694	24.96	22.42	47.36	63.33	347
62	22.08	22.02	49.56	64.18	1616	27.36	23.49	45.65	60.81	332
63	16.67	22.17	53.24	65.63	1609	17.22	22.54	51.75	64.02	326
64	15.96	23.79	53.88	64.65	1483	19.99	24.46	50.91	62.67	324
65	11.20	24.46	55.65	66.09	1520	11.43	27.76	54.32	64.40	314
66	11.28	24.27	56.93	65.97	1384	13.86	24.58	56.02	64.11	289
67	10.91	22.71	57.84	66.33	1292	14.19	21.54	56.08	64.66	280
68	9.18	23.46	58.31	66.46	1271	10.39	25.60	55.65	65.33	275
69	7.58	23.27	60.03	67.05	1176	9.04	23.14	61.48	63.97	250
70	6.73	23.96	58.84	67.61	1106	7.96	25.59	58.26	64.87	227

Notes: Values are sample averages computed using sample weights. Entire sample refers to the pooled surveys for 2003-2011. Healthy subsample refers to the pooled surveys from 2006-2008 for all individuals who reports a health status of 1, 2 or 3. MW is market work, HP is home production, L is leisure, PC is personal care. Last column shows sample size for each age.

Table A2

Synthetic Cohort Estimates of f_{HP} (standard errors in parentheses)

Age Range	Entire Sample	Healthy Subsample
60 – 65	–.274 (.066)	–.311 (.174)
60 – 66	–.264 (.047)	–.279 (.135)
61 – 65	–.255 (.109)	–.250 (.220)
61 – 66	–.246 (.069)	–.219 (.159)

Notes: Table reports the coefficient on market work when running the regression in equation (3.1) in the text for different age ranges. Entire sample refers to the pooled surveys for 2003-2011. Healthy subsample refers to the pooled surveys from 2006-2008 for all individuals who reports a health status of 1, 2 or 3. Sample weights are used.

Table A3

Cross-Sectional Estimates of f_{HP} by Age (standard errors in parentheses)

	Entire Sample	Healthy Subsample
60	–.270 (.014)	–.298 (.030)
61	–.258 (.016)	–.328 (.038)
62	–.268 (.015)	–.301 (.036)
63	–.277 (.015)	–.277 (.035)
64	–.296 (.017)	–.288 (.033)
65	–.256 (.018)	–.319 (.036)
66	–.291 (.019)	–.276 (.033)

Notes: The table reports the negative of the coefficient on hours of market work per week with dependent variable being hours of home production per week. Controls for gender and four education levels (less than high school, high school, some college, at least college) are also included. Entire sample refers to the pooled surveys for 2003-2011. Healthy subsample refers to the pooled surveys from 2006-2008 for all individuals who reports a health status of 1, 2 or 3. Sample weights are used.