Retirement, Home Production and Labor Supply Elasticities*

Richard Rogerson

Johanna Wallenius

Princeton University

Stockholm School of Economics

December 2011

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.

*Rogerson thanks the NSF for financial support. Wallenius thanks the Yrjo Jahnesson Foundation for financial support.

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. Following MaCurdy (1981), a large literature has estimated the willingness of individuals to substitute leisure over time by examining changes in hours worked and wages for continuously employed prime aged individuals, typically males. Most of these analyses conclude that this willingness is very low. Recent work shows that these estimates are quite sensitive to a variety of plausible extensions, including for example, human capital accumulation, credit constraints faced by younger workers, restrictions on hours worked and optimization frictions.¹ In view of this situation it is of interest to explore additional settings that can yield information about the key preference parameters that determine labor supply elasticities. In this paper we argue that studying how time allocations change at the time of retirement is an important alternative source of information about these parameters. Importantly, we show that by focusing on what happens at the time of retirement, the features mentioned above that are known to greatly affect estimates based on looking at continuously employed prime age individuals, are much less relevant.

¹See, for example Imai and Keane (2004) and Wallenius (2011) for analyses that include human capital accumulation, Domeij and Floden (2006) for an analysis that includes credit constraints, Chang and Kim (2006) and Rogerson (2011) for analyses that include hours restrictions, and Chetty (2011) for a discussion of optimization frictions. See also Keane and Rogerson (2011) for a general discussion of these issues.

Although the process of retirement varies across individuals, a typical pattern involves individuals moving from full time market work of roughly 2000 hours per year to no market work. We build a simple model of life cycle labor supply that includes home production and captures this feature of retirement. Specifically, we assume that individuals face constraints on working hours: either they work full time or not at all.² Our main result does not rely on whether the retirement decision is optimal, i.e., it is consistent with either an individual who optimally chooses the timing of retirement subject to the above restriction on hours, or an individuals who is exogenously forced to retire.

Independently of whether the retirement decision is optimal, we study the optimal response of time allocation to home production and leisure following the transition from full time work to no work. The relative response of these two uses of time is very dependent on the two key preference parameters that we noted above. Specifically, the increase in leisure time at retirement is increasing in the individual's willingness to substitute leisure over time, and decreasing in the individual's willingness to substitute between home and market produced goods. Conversely, the increase in time devoted to home production at retirement is decreasing in the individual's willingness to substitute leisure over time and increasing in the elasticity of substitution between home and market goods. We derive a simple relationship that links the relative value of these two elasticity parameters and the change in time allocations at retirement. It is important to

²By linking retirement to a restriction on choices of working hours we are implicitly following Rust and Phelan (1997) and Laitner and Silverman (2005). See also the discussion in Hurd (1996) and Blau and Shvydko (2007).

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. A large literature has documented the empirical importance of the home production margin. And as we will see later, the fact that time devoted to home production does increase at retirement suggests that this margin should be explicitly included.

Having derived an expression that links preference parameters to changes in time allocations at retirement, we examine data from the recently available American Time Use Survey (ATUS). Based on this data we find that only about 15–20% of the additional time that becomes available at retirement is devoted to home production. Additionally, the literature on estimating the elasticity of substitution between time and goods suggests elasticities in the vicinity of 2. (See, for example, Aguiar and Hurst (2007).) Given this estimate, our model implies that the intertemporal elasticity of substitution for leisure is also around 2. Even if we take a value of one as a lower bound on the elasticity of substitution between time and goods, our model implies a value for the intertemporal elasticity of substitution of leisure that is also around one. Our estimates are consistent with the studies that conclude that earlier estimates of this elasticity are significantly biased toward zero because of the neglect of the various factors noted above.

By virtue of using data on retirement to infer the value of the intertemporal elasticity of substitution, this paper is most related to the recent paper by Rogerson and Wallenius (2011). While both papers focus on properties of retirement as

a source of information, the underlying sources of identification are very different. In Rogerson and Wallenius (2011), 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 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. This paper shows that in a model with home production, the changing time allocation between leisure and home production also provides information on preference parameters. Nonetheless, although this paper derives a relationship that captures economic forces that are distinct from our earlier analysis, we are lead to a similar inference, namely, a relatively large values for the elasticity of intertemporal substitution.

An outline of the paper follows. In Section 2 we describe the basic life cycle model with home production and indivisible labor, and derive the key expression that links preference parameters to the changing allocations at retirement. Section 3 presents data from the ATUS and derives the quantitative implications of this expression. Section 4 concludes.

2. Retirement in a Life Cycle Model With Home Production

In this section we describe the life cycle model that we analyze, develop the 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 solves a simple life cycle maximization problem in continuous time. We adopt a continuous time formulation so that the retirement decision reflects a continuous choice. Lifetime utility is given by:

$$\int_0^1 [u(c(t)) + \frac{\alpha}{1 - \frac{1}{\gamma}} (1 - h(t))^{1 - \frac{1}{\gamma}}] dt. \quad (2.1)$$

where u has standard properties, i.e., it is strictly increasing, strictly concave and twice continuously differentiable, and $\gamma > 0$. 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. We abstract from discounting to simplify the analytics, but will also assume that the interest rate is zero, so that the interest rate and discount rate perfectly offset each other. In the spirit of Becker (1965) we include home production and assume that the consumption that individuals care about (c) is an aggregate of market purchased goods (g) and home production time (h_n). Following much of the literature, we assume a CES aggregator:

$$c(t) = [ag(t)^{1 - \frac{1}{\eta}} + (1 - a)h_n(t)^{1 - \frac{1}{\eta}}]^{\frac{\eta}{\eta - 1}}. \quad (2.2)$$

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 $h(t) = h_m(t) + h_n(t)$ is total time devoted to work, where $h_m(t)$ is time devoted to market

work.

Our focus is on how time allocations change at the time of retirement, where retirement takes the form of an individual moving from full time work to no work. We therefore assume that individuals are faced with a discrete choice problem in which the only two options for market work are full time work, denoted by $\bar{h}$, and no work, so that $h_m(t) \in \{0, \bar{h}\}$ for all t . In contrast, we assume that home production time can be varied continuously. The total time endowment of the individual is equal to one at all dates.

Following Ljungqvist and Sargent (2011) we assume that individuals accumulate human capital via learning by doing, so that wages increase with accumulated market work. Specifically, we assume that wages at time t are given by:

$$w(t) = AH_m(t)^\phi$$

where

$$H_m(t) = \int_0^t h_m(s) ds$$

We also assume that the individual is eligible for social security and/or pension benefits. In particular, we assume that starting at age R the individual will receive $b(t)$ for the remainder of life, where the value of b may depend on the (endogenous) values of $H_m(t)$, i.e., the amount of time devoted to work during the interval $[0, R]$, and $h_m(t)$, as well as the exogenous parameters A and ϕ . This specification allows for the possibility that benefits depend on age, current employment status, years of employment, and lifetime income. As we will see below, our main results are

robust to the exact specification of the benefit function, so we will not have to be more explicit. The key point that we want to stress is that our results are consistent with a large class of pension/social security systems.³

The individual faces a present value budget equation given by⁴:

$$\int_0^1 g(t)dt = \int_0^1 w(t)h_m(t)dt + \int_R^1 b(t)dt \quad (2.3)$$

2.2. Optimal Life Cycle Allocations

Consider the life cycle utility maximization problem solved by the individual in the case in which the retirement benefit is equal to zero. In this case the non-convexity associated with the discrete choice for market hours implies that the optimal solution may entail some dates at which the individual works $\bar{h}$ in the market and some dates at which he or she works 0 in the market. As noted in Ljungqvist and Sargent (2011), in this case the individual is indifferent about the timing of work. Given that there are many small perturbations of the model that would resolve this indeterminacy in favor of having work front-loaded (e.g., assuming that the parameter A decreases even slightly with age), it is natural to focus on the solution in which work is front-loaded. Given our rather general formulation of the retirement benefit function, the features of this function may

³We have abstracted from explicitly including taxes on income in this specification. In fact, for the results that we focus on we could allow for a very general function that maps total income into after tax income.

⁴While we assume the individual faces a single present value budget equation and hence assume complete markets in terms of borrowing and lending, our results would be unaffected if we assumed that young individuals were not able to borrow against future income and so were credit constrained. What matters for our results is that there is some point prior to retirement at which the individual is not credit constrained.

also create an incentive for the individual to not work at certain points of the life cycle. Additionally, assuming that benefits are strictly increasing in previous accumulated income, the presence of benefits would also give rise to an incentive to front load market work. In what follows we will assume that the optimal solution for market work involves front loading of work. Additionally, consistent with the evidence that the vast majority of individuals do retire and our desire to focus on how time allocations change at retirement, we will also assume that parameters of the problem are such that the optimal life cycle solution does entail an interior solution for the fraction of life that the individual spends in market employment. We will refer to the period of life in which the individual does not work as retirement.

The fact that hours of market work when employed are exogenously set to $\bar{h}$ independently of wages creates a symmetry to the individual's life cycle maximization problem that simplifies the nature of the solution. Specifically, the optimal solution will be described by five numbers: e , h_w , h_r , g_w , and g_r , where e is the fraction of life spent in (market) employment, g_w and g_r represent the consumption of market goods when working and retired, respectively, and h_w and h_r represent time spent in home production when working and retired, respectively. It is convenient to define c_w and c_r by:

$$c_w = [ag_w^{1-\frac{1}{\eta}} + (1-a)h_w^{1-\frac{1}{\eta}}]^{\frac{\eta}{\eta-1}} \quad (2.4)$$

$$c_r = [ag_r^{1-\frac{1}{\eta}} + (1-a)h_r^{1-\frac{1}{\eta}}]^{\frac{\eta}{\eta-1}} \quad (2.5)$$

The individual's life cycle optimization problem can now be written as:

$$\begin{aligned} \max \quad & e[u(c_w) + \frac{\alpha}{1 - \frac{1}{\gamma}}(1 - \bar{h} - h_w)^{1 - \frac{1}{\gamma}}] + (1 - e)[u(c_r) + \frac{\alpha}{1 - \frac{1}{\gamma}}(1 - h_r)^{1 - \frac{1}{\gamma}}] \\ \text{s.t.} \quad & eg_w + (1 - e)g_r = I \end{aligned} \quad (2.6)$$

$$I = \int_0^e A(t\bar{h})^\phi \bar{h} dt + \int_R^1 b(t) dt \quad (2.7)$$

Next we derive a relationship that must hold as part of the optimal solution to the life cycle maximization problem stated above. Interestingly, we will not make any use of the first order condition for e . As a result, the results that we derive are not contingent on the optimality of the timing of retirement, and so would continue to hold if retirement were exogenous. All that we require is that the value of e is interior. Given an interior value for e , and letting μ be the Lagrange multiplier on the budget equation (2.6), the first order conditions for g_w , g_r , h_w and h_r are given by:

$$g_w : u'(c_w)c_w^{\frac{1}{\eta}} ag_w^{-\frac{1}{\eta}} = \mu \quad (2.8)$$

$$g_r : u'(c_r)c_r^{\frac{1}{\eta}} ag_r^{-\frac{1}{\eta}} = \mu \quad (2.9)$$

$$h_w : u'(c_w)c_w^{\frac{1}{\eta}}(1 - a)h_w^{-\frac{1}{\eta}} = \alpha(1 - \bar{h} - h_w)^{-\frac{1}{\gamma}} \quad (2.10)$$

$$h_r : u'(c_r)c_r^{\frac{1}{\eta}}(1 - a)h_r^{-\frac{1}{\eta}} = \alpha(1 - h_r)^{-\frac{1}{\gamma}} \quad (2.11)$$

Note that the details of lifetime income determination included in equation (2.7) do not enter any of these four first order conditions. These details would

enter the first order condition for e , the fraction of life spent in employment, or equivalently, the age at which the individual retires. Although our analysis does assume that the solution for e is interior, the conditions that we derive for how time allocations change at retirement are independent of the age at which retirement takes place, implying that we can bypass explicit consideration of the first order condition for e , and hence the exact details of pension and/or social security provisions.

Divide (2.8) by (2.9) to get:

$$\left[\frac{g_w}{g_r}\right]^{\frac{1}{\eta}} = \frac{u'(c_w)}{u'(c_r)} \left[\frac{c_w}{c_r}\right]^{\frac{1}{\eta}} \quad (2.12)$$

Divide (2.10) by (2.11) to get:

$$\left[\frac{1 - h_r}{1 - \bar{h} - h_w}\right]^{\frac{1}{\gamma}} \left[\frac{h_w}{h_r}\right]^{\frac{1}{\eta}} = \frac{u'(c_w)}{u'(c_r)} \left[\frac{c_w}{c_r}\right]^{\frac{1}{\eta}} \quad (2.13)$$

Substituting (2.12) into (2.13), taking logs and rearranging gives:

$$\frac{\gamma}{\eta} = \frac{\log(1 - h_r) - \log(1 - \bar{h} - h_w)}{\log(g_w/g_r) - \log(h_w/h_r)} \quad (2.14)$$

Equation (2.14) 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 labor supply elasticities, 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 (2.14) 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 labor 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 how our method for imposing discipline on preference parameters with 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. Quantitative Implications for Preference Parameters

In this section we explore the quantitative implications of equation (2.14). In order to implement the analysis we need to have values for the objects on the right hand side of this equation: $\bar{h}$, h_w , h_r , and the ratio g_w/g_r . The first subsection presents empirical evidence on these values, and the second subsection presents the implications for preference parameters.

3.1. Empirical Evidence on Life Cycle Allocations

In the model we normalized the time endowment to equal one. We interpret this to represent the amount of discretionary time that an individual has, which we take to be 100 hours per week or 5200 hours per year. We calibrate $\bar{h}$ based on the evidence in Rogerson and Wallenius (2011). Based on the analysis of the PSID, they found that annual hours of work for males who worked full time at age 60 and retired over the subsequent ten years was slightly greater than 2000. Additionally, time use data suggests an average commuting time of roughly 200

hours per year. Combining these two numbers leads to $\bar{h} = .42$.

There is a sizeable literature that studies the drop in consumption at retirement. Aguiar and Hurst (2005) emphasize that a drop in consumption expenditure is not the same as a drop in true consumption. This distinction is captured by our model, since the drop in consumption expenditure corresponds to g_r/g_w , whereas the change in the flow of consumption corresponds to the ratio c_r/c_w . 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 represents fixed consumption costs associated with work that do not generate utility, the target in our model should be adjusted appropriately.⁵ Also, if we allowed for nonseparability between consumption and leisure, as in Laitner and Silverman (2005), it would induce an additional channel through which consumption of market goods would drop at retirement. From this we conclude that we should consider .85 to be a lower bound for g_r/g_w . For our benchmark specification we assume $g_r/g_w = .90$, but we also consider values of .85 and .80.

We set $h_w = .15$ in our benchmark specification, based on various estimates from time use data. Pooling all of the ATUS data from 2003 through 2010, we find that the average time devoted to household production by employed males aged 60 – 64 is 16.3.⁶ Using data from the 2001 Consumption and Activities Mail

⁵Specifically, if there are fixed costs $\bar{g}$ associated with working that do not generate utility, then our expressions would all involve $g_w - \bar{g}$ instead of g_w and we would be targeting $g_r/(g_w - \bar{g})$ rather than g_r/g_w .

⁶Also based on the ATUS, Ramey (2009) reports that time spent in home production for all

Survey (CAMS) supplement to the Health and Retirement Survey (HRS), Hurd and Rohwedder (2003) report that non-retired males between the ages of 60 and 64 devote 14.29 hours of time to home production.

The final 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. Given values for $\bar{h}$ and h_w , we will find it useful to parameterize h_r by thinking about the quantity $\varepsilon_{hp} = (h_r - h_w)/\bar{h}$, i.e., the fraction of the increased time available at retirement that is allocated toward home production. One way to estimate this value is to contrast the time devoted to home production between working and retired individuals at a given point in time. For example, the same survey reported earlier in Hurd and Rohwedder (2003) finds that retired males between the ages of 60 and 64 report 19.45 hours of time devoted to home production, implying $\varepsilon_{HP} = .0516/.42 = .12$. One concern with this calculation is the selection of individuals into the working and not working groups at a given age. In particular, one might plausibly think that individuals who prefer home production to market work are more likely to retire earlier, so that their home production time is an overestimate of the time that employed individuals will devote to home production when they retire.

Ideally, one would like to have micro panel data on time use in order to estimate ε_{HP} . Unfortunately, the best available source of time use data is the American Time use Survey (ATUS), but it does not have a panel component. Here we carry out an alternative procedure that attempts to overcome some of the possible selection effects despite not having access to panel data. Specifically, we pool all employed males aged 18-64 has averaged close to 16 hours per week since 2000.

of the cross-section data available in the ATUS and construct a synthetic panel of time use allocations for individuals between the ages of 60 and 70. We then examine how average time devoted to home production responds to changes in average time devoted to market work in the synthetic panel.⁷

Before presenting the results, given that our goal is to characterize what happens to home production time at retirement, we think it is important to emphasize that retirement from full time work is in fact the dominant source of variation in market hours in our synthetic panel. To see this, Table 1 presents pooled cross-section data for annual hours worked in the previous year from the CPS for the years 2002-2004, and shows the fraction of individuals in each of the various bins.

⁷While creation of a synthetic panel does avoid the selection problems inherent in a single cross-section, it can only accomodate a limited set of cohort effects. Given that we focus on a relatively narrow age range and that the surveys come from a ten year period, we suspect that cohort effects may not be too serious.

Table 1

Distribution of Male Annual Hours by Age, CPS 2002-2004

Age	Annual Hours					
	0	(0, 250)	[250, 750)	[750, 1250)	[1250, 1750)	≥ 1750
60	.26	.02	.03	.05	.05	.59
61	.29	.02	.03	.05	.06	.55
62	.36	.02	.04	.06	.05	.47
63	.43	.02	.04	.06	.05	.39
64	.48	.03	.06	.06	.04	.34
65	.54	.03	.05	.07	.04	.27
66	.59	.03	.05	.06	.04	.22
67	.63	.04	.04	.06	.05	.19
68	.65	.03	.05	.06	.04	.16
69	.69	.03	.04	.05	.04	.15
70	.75	.03	.03	.04	.03	.12

The key message from this Table is that as males age from 60 to 70 the key dynamic is the movement from full time work (more than 1750 hours per year) to no work.⁸

Table 2 presents the synthetic panel for time use based on the ATUS for males in 5 categories: market work (MW), home production (HP), eating and drinking (E&D), leisure (L) and personal care (PC).⁹

⁸This dynamic is not specific to men. In the appendix we present the same information for women and men and show that the same pattern holds.

⁹We aggregate activities in the ATUS into these categories as follows. Market work includes

Table 2

Time Use By Age: Men

Age	MW	HP	L	E&D	PC
60	23.12	17.90	41.15	9.11	63.84
61	19.27	17.58	44.33	8.47	65.24
62	17.65	18.62	44.24	9.27	64.39
63	14.36	19.10	46.73	9.39	65.25
64	14.16	18.88	47.75	9.51	65.30
65	10.12	19.80	51.75	9.07	64.98
66	9.62	19.64	51.13	9.20	66.03
67	9.70	19.29	50.83	9.41	65.86
68	8.06	19.57	52.73	9.71	64.97
69	7.30	17.56	53.20	9.45	66.93
70	6.92	17.83	52.96	9.56	66.66

We now ask how the decrease in market hours is related to changes in the other categories, in particular home production and leisure. To do this we run

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. Home production includes 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 includes consumer purchases, professional and personal care purchases, purchasing household services, and purchasing government services. Leisure includes socializing and communicating, attending and hosting social events, relaxing and leisure, arts and entertainment, and waiting associated with the above. Personal care consists of sleeping, grooming, health related self-care, personal activities and personal care emergencies. Note that the categories in Table 2 are not exhaustive; there is a small set of other activities that account for a roughly constant number of hours over this age range.

simple regressions of the form:

$$cat_{ia} = \beta_0 + \beta_i mw_a + \varepsilon_{it}$$

where cat_{ia} is time use in category i at age a and mw_a is time devoted to market work at age a . In terms of our earlier notation, we have $\varepsilon_i = -\beta_i$. Table 3 reports results.

Table 3

Time Use Responses to Changes in Market Work			
Age Range	ε_{HP}	ε_L	ε_{PC}
60 – 70	.05 (.06)	.77 (.04)	.13 (.07)
60 – 65	.16 (.05)	.78 (.11)	.09 (.04)
60 – 68	.14 (.02)	.77 (.05)	.09 (.04)
61 – 66	.20 (.05)	.82 (.12)	.07 (.08)
61 – 68	.16 (.04)	.81 (.07)	.05 (.06)
61 – 70	.04 (.08)	.79 (.05)	.12 (.06)

While our main focus is on the value of ε_{HP} , we briefly discuss the other estimates first. The range of estimates for ε_L is reasonably tight, varying only from .77 to .82. The point estimates for personal care are often not significant at the 5% level. To the extent that additional sleep at the margin can be considered a component of leisure, one might interpret the increases in personal care as effectively representing increases in leisure.

The range of estimates for ε_{HP} is .04 to .20. Most of the estimates are significant at the 5% level, and the tightest estimate comes for the age range 60 to 68

and has a point estimate of .14. The estimate is not significant for either regression in which the upper age is set to 70. This may reflect that fact that health deteriorates at older ages, making both market and home work more onerous. As we will see below when we present implications for preference parameters, a conservative approach corresponds to choosing higher estimates of ε_{HP} . With this in mind we think a plausible range for ε_{HP} is .15 to .20, and in our benchmark specification will choose a value of .20.¹⁰

Our estimates are slightly larger than the estimate based on the data in Hurd and Rohwedder (2004), though it should be noted that the two data sets are different. Hamermesh (2007) carries 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. He finds that 40 minutes of the roughly 170 minute decrease in market work is allocated to home production, implying an estimate of $\varepsilon_{HP} = .23$, slightly above our estimate.

In a different but related exercise, Aguiar et al (2011) use the ATUS to ask how the decrease in market work during the current recession has been allocated to other time use categories. They conclude that between 40 and 50 percent of the decrease is allocated to home production. While this estimate is much larger than ours, two key differences should be noted. First, there is no reason to think that the response to a business cycle shock should apply to the retirement decision. Second, their estimate is for the entire population rather than just those

¹⁰We also ran these regressions using data for both men and women. In this case the range of estimates was .05 to .15. So while it is the case that on average women devote more time to home production than men, the change in home production in response to a change in market work seems to be roughly similar for both groups.

of retirement age. Nonetheless, we will consider values of ε_{HP} as large as .45 in our analysis below.

3.2. Results

Given values for $\bar{h}$, g_r/g_w , h_w , and h_r , equation (2.14) tells us the value of γ/η that is required to support those values and therefore implicitly defines a one dimensional curve in $\gamma - \eta$ space. In presenting results we will consider a range of estimates for η and then map out the implications for values of γ . The essence of the idea of home production is that home production is a substitute for market goods. This suggests that one is a lower bound for η . We will see below that even this weak restriction is somewhat informative. Beyond this weak restriction, there are several estimates of η in the literature. 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. Using micro data, Rupert, Rogerson and Wright (1995) find an estimate in the range 1.67 to 1.8, while Aguiar and Hurst (2007) report an estimate for their benchmark specification in the range of 2.0 to 2.5. Based on this, we consider values for η that range from 1.0 to 2.5.

For our benchmark specification, equation (2.14) implies that $\gamma = 1.05\eta$, so that to a first order, these two preference parameters are roughly equal. Table 4 shows the implied values for γ for specific values of η .

Table 4

Implied Values of γ for Benchmark Specification					
$\eta = 1.00$	$\eta = 1.25$	$\eta = 1.50$	$\eta = 1.75$	$\eta = 2.00$	$\eta = 2.50$
1.05	1.31	1.57	1.84	2.10	2.62

Even at the lower bound of $\eta = 1.00$ we see that the benchmark specification requires a value of γ equal to 1.05. But, given that most estimates of η lie above 1.5, our benchmark specification yields values of γ that are well above those obtained in much of the literature based on examining changes in wages and hours for working individuals. However, estimates based on models that assume human capital accumulation, such as Imai and Keane (2004) and Wallenius (2011) do deliver estimates that are within this range.

If we consider smaller values for g_r/g_w the lower bound on γ is decreased, but the main message remains. For example, if we use $g_r/g_w = .85$ instead of .90 as in the benchmark, the lower bound on γ is reduced to .95, and if we use $g_r/g_w = .80$, the lower bound for γ is reduced to .86.

Finally, we explore how the implications are affected by using higher values for ε_{HP} , including values as high as those reported in Aguiar et al (2011) for time use during the current recession. We note that the evidence in Aguiar et al refers to the entire population and so also entails a higher base level of time devoted to home production. Consistent with this we assume $h_w = .20$ in these calculations. Results are shown in Table 5.

Table 5

Effect of ε_{HP} on γ/η : $g_r/g_w = .90$, $h_w = .20$								
$\varepsilon_{HP} :$	.10	.15	.20	.25	.30	.35	.40	.45
$\gamma/\eta :$	2.33	1.75	1.39	1.14	.96	.82	.71	.62

If $\eta = 1$ is a lower bound, then the values in Table 5 are also the lower bounds on γ . If $\varepsilon_{HP} = .45$ then the lower bound on γ is .62, a value that is still on the high side relative to much of the earlier literature, but perhaps not too far removed from the value of .45 that Chetty (2011) argues is an appropriate estimate if one reinterprets previous empirical evidence in light of optimization frictions. However, Aguiar et al (2011) also argue that η should be at least as high as 2.0 to rationalize the data on time use during the recent recession, in which case the implied estimate for γ is 1.24 even assuming $\varepsilon_{HP} = .45$.¹¹

4. Conclusion

In a life cycle model that features home production and an endogenous retirement decision, we show that the change in time allocations between home production and leisure at retirement imposes a tight relationship between two key preference parameters that 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,

¹¹If we had maintained $h_w = .15$ in these calculations the lower bound for γ would be .47 and assuming $\eta = 2.0$ would imply $\gamma = .94$.

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. Even assuming a lower bound of one for the elasticity of substitution between time and goods, our model implies a value of the intertemporal elasticity of substitution that exceeds one. If we instead assume an elasticity of substitution between time and goods that is in the vicinity of 2, as suggested by empirical work, our model implies that the intertemporal elasticity of substitution for leisure will exceed two.

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. 2011. "Time Use During Recessions." NBER Working Paper #17259.
- [4] Becker, Gary, "A Theory of the Allocation of Time." *The Economic Journal* 75 (1965): 493-517.

- [5] Blau, D., and Shvydko, T. 2007. "The Labor Market Rigidities and Employment Behavior of Older Workers." Mimeo, Ohio State University.
- [6] 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.
- [7] Chang, Y., and Schorfheide F. 2003. "Labor Supply Shifts and Economic Fluctuations." *Journal of Monetary Economics* 50, 1751-1768.
- [8] Domeij, D., and Floden M. 2006. "The Labor Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased." *Review of Economic Dynamics* 9, 242-262.
- [9] Gronau, Reuben. 1977. "Leisure, Home Production and Work—The Theory of the Allocation of Time Revisited." *Journal of Political Economy* 75, 1099-1124.
- [10] Hamermesh, D. 2007. "The Time and Timing Costs of Market Work." NBER Working Paper #13127.
- [11] 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.
- [12] Hurd, M., and Rohwedder, S. 2003. "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement." NBER Working Paper #9586.

- [13] Imai, S., and Keane M. 2004. "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review* 45, 601-641.
- [14] Laitner, J., and Silverman, D. 2005. "Estimating Life-Cycle Parameters from Consumption Behavior at Retirement." NBER Working Paper #11163.
- [15] Ljungqvist, L., and Sargent, T. 2011. "Curvature of Earnings Profile and Career Length." Mimeo, New York University.
- [16] MaCurdy, T. 1981. "An Empirical Model of Labor Supply in a Life Cycle Setting." *Journal of Political Economy* 89, 1059-1085.
- [17] 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.
- [18] Ramey, V. 2009. "Time Spent in Home Production in the 20th Century United States: New Estimates from Old Data." *Journal of Economic History* 69, 1-47.
- [19] Rogerson, R.(2011. "Individual and Aggregate Labor Supply in a Model With Coordinated Working Times." Forthcoming, *Journal of Money, Credit and Banking*.
- [20] Rogerson, R., and Wallenius, J. 2011. "Nonconvexities, Retirement and the Elasticity of Labor Supply." Mimeo, Princeton University.
- [21] Rupert, P., Rogerson R., and Wright R. 1995. "Estimating Substitution Elasticities in Household Production Models." *Economic Theory* 6, 179-193.

- [22] Rust, J., and Phelan, C. 1997. “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets.” *Econometrica* 65, 781-831.
- [23] 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.