

Micro and Macro Elasticities in a Life Cycle Model With Taxes

Richard Rogerson and Johanna Wallenius*

Arizona State University

Original Version: December 2006. This Version: June 2007.

Abstract

We build a life cycle model of labor supply that incorporates changes along both the intensive and extensive margin and use it to assess the consequences of changes in tax and transfer policies on equilibrium hours of work. We find that changes in taxes have large aggregate effects on hours of work. Moreover, we find that there is no inconsistency between this result and the empirical finding of small labor elasticities for prime age workers. In our model, micro and macro elasticities are effectively unrelated. Our model is also consistent with other cross-country patterns.

*We thank Ed Prescott, Martin Evans, Nobu Kiyotaki, Robert Hall, Fernando Alvarez and seminar participants at the University of Pennsylvania, University of Maryland, Georgetown University, ASU, Kansas City Fed, Philadelphia Fed and Atlanta Fed for useful comments. An earlier version of this paper was presented at the conference in honor of Robert Lucas held in Chicago, April 2007. Rogerson thanks the NSF for financial support, and Wallenius thanks the Yrjo Jahnsson Foundation for financial support.

1. Introduction

Time devoted to market work in continental Europe is currently only about 70% of the US level. Recent work by Prescott (2004), Rogerson (2005) and Ohanian et al (2006) argues that differences in tax and transfer policies can account for a large share of this difference. Following Lucas and Rapping (1969), these papers all use a stand-in household model which abstracts from the distinction between employment and hours per employee, and assume that the stand-in household has a relatively high labor supply elasticity. One critique of these exercises is that the assumed labor supply elasticity of the stand-in household is much larger than that implied by most estimates based on micro data. Specifically, if the labor supply elasticity of the stand-in household was instead set to standard estimates based on micro data, then it is no longer the case that taxes account for a large share of differences in market work between the US and continental Europe.¹ In this paper we argue that this critique is misplaced.

To make this point, we develop an overlapping generations model that replicates the salient features of life cycle labor supply, and then use this model to analyze how tax and transfer policies affect hours of work in the steady state. In this framework we can carry out both standard micro data estimation exercises based on life cycle variation for prime aged workers, as well as standard macro estimation based on variation in aggregates across steady states. Our main findings are twofold. First, macro elasticities are virtually unrelated to micro elasticities, and second, macro elasticities are large. In particular, for micro elasticities that

¹Alesina et al (2005) is a recent example where this critique is put forward.

vary by a factor 25, ranging from .05 to more than 1.25, the corresponding macro elasticities are in the range of 2.25 – 3.0.

Our model builds on the earlier work of Prescott et al (2006) on lifetime labor supply by imbedding it into a life cycle setting. Like them, we focus on the importance of non-linearities in the mapping from time at work to labor services provided. This feature gives rise to equilibrium allocations in which workers choose time allocations in which both the extensive and intensive margins are operative, i.e., a worker chooses both what fraction of his or her life to devote to employment, and what fraction of his or her period time endowment to devote to work while employed. By imbedding this analysis into a life cycle model we are able to generate standard life cycle profiles for hours of work—including the fact that hours drop discontinuously to zero at older ages. This allows us to reproduce micro estimates based on life cycle variation for prime aged workers.

In addition to reconciling micro and macro tax elasticities, our life cycle model delivers two additional predictions relative to earlier analyses based on stand-in household models, both of which are also corroborated by the cross-country data. First, our model predicts that increases in the size of tax and transfer policies imply less time devoted to work both in terms of employment to population ratios and hours of work per person in employment. Second, our model implies that differences in employment to population ratios are dominated by differences among young and old individuals.

Our results are related to those in a recent paper by Chang and Kim (2006). They study a model in which individuals are subject to idiosyncratic shocks to

wages, can self-insure through saving but have no access to insurance markets, and in which labor is indivisible, so that the supply decision consists solely of a decision at the extensive margin. Similar to our result, they find that in their model, micro and macro elasticities need not be the same, and that macro elasticities can be significantly larger. While we view our study as complementary to theirs, there are several important differences that distinguish the two studies. First, our model is a life cycle model and hence can explicitly connect to micro estimates based on life cycle variation. Second, our analysis allows for variation along both the intensive and extensive margin. Not only does this allow one to better match the cross-country differences in hours of work, but we show that there is an important interaction between intensive and extensive margins: less movement on the extensive margin necessarily implies more adjustment on the intensive margin. Third, and related to this point, Chang and Kim argue that increases in heterogeneity in the steady state cross-sectional distribution of wages imply a large reduction in implied macro elasticities. We find that having an operative intensive margin reduces the quantitative impact of this effect.

An outline of the paper follows. Section 2 presents the model, and Section 3 develops a simple characterization of steady-state life cycle labor supply. Section 4 considers the effects of tax and transfer programs on the equilibrium, and considers the relationship between micro and macro elasticities. Section 5 considers some extensions of the analysis, and section 6 concludes.

2. Model and Equilibrium

The model is purposefully specialized along several dimensions in order to best highlight those relationships that are the focus of our analysis. We consider a continuous time overlapping generations framework in which a unit mass of identical, finitely lived individuals is born at each instant of time t . Continuous time is convenient for our analysis because we will study the endogenous determination of the length of working life, and continuous time allows us to model this as a continuous choice variable. We normalize the length of the lifetime of each agent to one and assume that each individual is endowed with one unit of time at each instant. Letting a denote age, individuals have preferences over paths for consumption ($c(a)$) and hours worked ($h(a)$) given by:

$$\int_0^1 U(c(a), 1 - h(a)) da \quad (2.1)$$

where U is twice continuously differentiable, strictly increasing in both arguments and strictly concave. Note that we assume that individuals do not discount future utility.²

We assume that labor is the only factor of production, and write the aggregate production function as:

$$Y(t) = L(t) \quad (2.2)$$

where $L(t)$ is aggregate input of labor services. A key feature of our model is

²This is done purely for convenience to allow us to focus on a zero interest rate steady state. From the perspective of hours worked, interest rates and discount factors serve primarily to tilt the life cycle profile.

the mapping from hours of work into labor services, which is described by two functions, $\tilde{e}(a)$ and $g(h)$. In particular, we assume that if an individual of age a devotes h units of time to market work then it will yield $l = \tilde{e}(a)g(h)$ units of labor services. The function $\tilde{e}(a)$ is standard in the life cycle labor supply literature—it represents exogenous life cycle variation in individual productivity. This feature will be the driving force behind the variation in hours worked during that part of the life cycle in which an individual is employed. In a later section we will also consider variation in the disutility of work and show how it can play a similar role. We assume that the function $\tilde{e}(a)$ is single peaked, twice continuously differentiable, and has zero derivative only at its global maximum, i.e., that it has no flat spots.

Following Prescott et al (2006), the function $g(h)$ is somewhat nonstandard, but plays a critical role in the analysis. A standard assumption in the literature is that g is linear with slope equal to one, so that for a worker of a given age, labor services are linear in hours of work. We assume that g is continuous, and increasing, with $g(0) = 0$, but we depart from the assumption of linearity. In particular, we assume that there is a value $\bar{h}$ with $0 < \bar{h} < 1$ such that g is strictly concave over the region $[\bar{h}, 1]$, but not over the region $[0, 1]$. The function $g(h)$ may be convex over the region $[0, \bar{h}]$. One justification for the initial convex region is fixed costs associated with getting set up for a job and costs associated with being supervised. This is consistent with the observation that firms do not consider part-time workers for many positions.³ The justification for the strictly concave region

³Our analysis emphasizes non-linearities in the mapping from hours worked to labor services provided. Similarly, one could allow for the possibility of non-linearities in the mapping from

is fatigue. Figure 1 shows one possible g that satisfies these properties.




Figure 1: The $g(h)$ Function

It is important to note the significance of the assumed properties of the function $g(h)$. These properties serve to generate a nonconvexity in the aggregate technology set for this economy. In a static setting with homogeneous agents, this nonconvexity implies that it may be optimal to randomly select a fraction of workers to work positive hours and have the remaining workers work zero hours. Loosely speaking, in such a setting this feature of technology can serve to endogenize the length of working time in a model of indivisible labor. Generalizing this homogeneous worker model to a dynamic setting with no discounting and no life cycle effects, Prescott et al (2006) showed that if time is continuous, optimal allocations take the form of a constant working time for employed workers, and a constant fraction of individuals employed at each instant. Importantly, such an

leisure time to leisure services.

allocation can be implemented as an equilibrium without lotteries, since individuals can choose the fraction of their lifetime that they work and use asset markets to smooth consumption in the face of an uneven income stream.

In the next section we show that in our overlapping generations model with life cycle effects, this feature can give rise to equilibria in which there is a well-defined notion of a working life—individuals will begin work at a particular age, and work continuously until retirement, at which point hours of work drop to zero. Moreover, the events of entering and leaving the labor force are discontinuous events, in the sense that hours of work jump discontinuously at these two points. In particular, hours of work do not gradually decrease to zero prior to retirement.

2.1. Equilibrium

We consider the following market structure. We assume that at time zero there are markets for labor services and consumption at all future dates. Let $w(t)$ and $p(t)$ denote the paths for prices in these two markets. We assume competitive behavior in all markets. If a given individual is alive at two dates t and t' , this market structure implicitly allows an individual to borrow or lend resources across these two dates at the gross interest rate $p(t)/p(t')$. Given that the aggregate production function is linear in labor services, competitive equilibrium necessarily implies that $w(t) = p(t)$ at each t .

Our analysis will focus on steady state equilibria associated with this market structure. As is well known, overlapping generations models can give rise to multiple steady state equilibria. For our economy, one can show that there is

always one steady state equilibrium in which $p(t)$ is constant, i.e., a steady state with a zero interest rate. Given that we assumed no discounting, this steady state will dominate any other steady state, and thus it is natural to focus on it. One potential problem is that there might not be any equilibrium that converges to this steady state without government intervention. In particular, it may be necessary for the government to issue debt in order to achieve the zero interest rate steady state. In the analysis that follows we will assume that if necessary, the government follows a policy that results in this steady state equilibrium being reached, and will therefore focus on the zero interest rate steady state equilibrium. What matters for our subsequent analysis is not that the interest rate is equal to zero, but rather that the interest rate is constant across steady states in the face of the labor policies that we consider.⁴

Given that we focus on the steady state equilibrium with constant $p(t)$, we can normalize this price to one, which by our earlier remark implies that $w(t)$ will also be one for all dates. The lifetime utility maximization problem for a newborn individual in the steady state equilibrium can then be written as:

$$\begin{aligned} \max_{c(a), h(a)} \int_0^1 U(c(a), 1 - h(a)) da \\ s.t. \int_0^1 c(a) da = \int_0^1 \tilde{e}(a) g(h(a)) da \end{aligned}$$

⁴Alternatively, we could also assume that there is a storage technology that can turn one unit of output at any instant into one unit of output at any future instant. Either way, our later results should be interpreted as showing the implications of various policies for labor market outcomes abstracting from effects associated with changes in interest rates.

It is of interest to first consider the special case in which $\tilde{e}(a)$ is constant over an individual's life. This case was studied by Prescott et al (2006), and they show that the solution for $h(a)$ can take one of two forms. One possibility is that $h(a)$ is positive for all a , in which case the solution for $h(a)$ is unique and has $h(a)$ constant for all a . The other possibility is that $h(a)$ is equal to zero for some a (in a set with positive measure). In this case there is a continuum of solutions for $h(a)$, but each characterized by the same two values: f , the fraction of the individual's life spent in employment, and h , the time devoted to work at any instant in which the individual is employed. That is, the solution pins down hours of work when employed and total hours supplied over the lifetime, but the timing of work is indeterminate.⁵ Of course, in steady state equilibrium, it is necessary that the pattern of hours worked across individuals be such as to yield constant aggregate hours at each point in time.

We now return to the case in which $\tilde{e}(a)$ is not constant. The next proposition states a very simple property of the optimal solution to this problem.

Proposition 1: The optimal solution for $h(a)$ has a reservation property. In particular, there exists a value $\tilde{e}^*$ such that $h(a) > 0$ if $\tilde{e}(a) > \tilde{e}^*$ and $h(a) = 0$ if $\tilde{e}(a) < \tilde{e}^*$. Moreover, the solution for $h(a)$ is unique.⁶

Proof: Suppose the solution does not have the reservation property, i.e., suppose there are ages a_1 and a_2 such that $h(a_1) > 0, h(a_2) = 0$ and $\tilde{e}(a_2) > \tilde{e}(a_1)$. Consider the alternative solution in which the individual switches the hours of

⁵The idea that theory predicts only the total time spent in employment and not the timing of employment was first noted by Mincer (1962) in his study of labor supply by married women.

⁶Formally, this result can be violated on a set of measure zero. For simplicity, we will abstract from this issue in both the statement of propositions and our proofs.

work and consumptions at these two ages. The lifetime utility of work is identical under these two scenarios, but the alternative scenario generates higher lifetime income, implying that consumption can be increased, thereby leading to higher lifetime utility. It follows that there is a reservation value $\tilde{e}^*$. Given that the timing of work is pinned down by the reservation value, and that $\tilde{e}(a)$ has no flat spots aside from at its maximum, it follows from standard arguments that the optimal solution for $h(a)$ is unique.//

Relative to the case in which $\tilde{e}(a)$ is constant, we see that allowing this function to vary over the life cycle serves to eliminate the indeterminacy.⁷ Intuitively, making individual productivity vary over time breaks the indeterminacy regarding the timing of labor supply, since the individual prefers to work when productivity is relatively high. The above result does not rule out the possibility that $\tilde{e}^* = 0$, in which case all individuals will work positive hours in the market at all points during their lives. But independently of whether hours are not always positive, this result coupled with our assumption on the profile $\tilde{e}(a)$ implies that there is a unique starting time for employment and a unique stopping time for employment (though one or both of these could still be at a corner). In particular, if we define $A_1 = \min\{a \in [0, 1] : \tilde{e}(a) \geq \tilde{e}^*\}$ and $A_2 = \max\{a \in [0, 1] : \tilde{e}(a) \geq \tilde{e}^*\}$, it follows that the individual will begin work at age A_1 , and work continuously until reaching age A_2 , at which point the individual will retire and not devote any time to market work from this age on.

The same logic which implies that the individual should work when produc-

⁷Mulligan (2001) notes this same property in a model with indivisible labor.

tivity is highest also implies that conditional on working, hours of work should be increasing in $\tilde{e}(a)$. In particular, we have the following proposition.

Proposition 2: Let $h^*(a)$ be the optimal solution for hours of work over the life cycle. Let a_1 and a_2 be distinct ages for which $h(a) > 0$. Then $\tilde{e}(a_1) > \tilde{e}(a_2)$ implies $h(a_1) > h(a_2)$.

Proof: Let $c(a)$ be the optimal profile for consumption. Assume by way of contradiction that $\tilde{e}(a_1) < \tilde{e}(a_2)$ and $h(a_1) \geq h(a_2)$. The assumed profiles cannot be utility maximizing, since by switching the values of c and h at ages a_1 and a_2 , lifetime utility and expenditure are unchanged, while income increases, thereby allowing for higher utility. //

3. Solving the Individual's Problem

Although the analysis thus far has not made any specific assumptions about functional forms, for the analysis that follows we will impose some additional assumptions. In particular, we will assume that preferences are separable between consumption and leisure⁸:

$$U(c, 1 - h) = u(c) - v(h) \tag{3.1}$$

where the function $u(c)$ is twice continuously differentiable, strictly increasing and strictly concave, and $v(h)$ is assumed to be twice continuously differentiable, strictly increasing and strictly convex.

⁸Although Hall (2007) presents evidence against separability, we make this assumption for analytic convenience.

Given no discounting, a zero interest rate, separable utility and strict concavity of u , it follows that in the steady state each individual will choose a constant profile for c .⁹ Recalling our earlier result regarding a reservation property for employment, we can rewrite the individual's maximization problem as:

$$\max_{c, h(a), A_1, A_2} u(c) - \int_{A_1}^{A_2} v(h(a)) da \quad (3.2)$$

$$s.t. \ c = \int_{A_1}^{A_2} \tilde{e}(a) g(h(a)) da \quad (3.3)$$

In this formulation, we express all variables as a function of the individual's age. Because productivity is the key driving force behind variation in hours, and productivity is not monotone in age, hours are not monotone as a function of age.

It is analytically convenient to reformulate the problem with what is essentially a change of variable. The idea is to re-order time for a given individual from the highest productivity instants to the lowest productivity instants. In particular, we define a function $e(\lambda)$ for $\lambda \in [0, 1]$ by requiring that for each λ , $e(\lambda)$ solves:

$$\lambda = \int_0^1 I(\tilde{e}(a) \geq e(\lambda)) da \quad (3.4)$$

where $I(\tilde{e}(a) \geq e(\lambda))$ is the indicator function which takes the value 1 if the inequality holds and is zero otherwise. In words, $e(\lambda)$ is that level of productivity

⁹Our analysis will abstract from the life cycle consumption profile. We could add life cycle preference shifters for $u(c)$ as one way to match the consumption profile without altering any of our other results. Assuming non-separable preferences would also allow us to capture the consumption profile. See Heckman (1974) for an early discussion of this mechanism.

such that the individual has a higher productivity for exactly the fraction λ of their life. Note that by construction, $e(\lambda)$ is strictly decreasing. Put somewhat differently, $\tilde{e}(a)$ traces out what happens to productivity as we follow an individual chronologically, while $e(\lambda)$ traces out what happens to productivity as we follow an individual from highest productivity instants to lowest productivity instants. Figure 2 shows the $\tilde{e}(a)$ and $e(\lambda)$ functions for an $\tilde{e}(a)$ function that is piece-wise linear.




Figure 2: The $\tilde{e}(a)$ and $e(\lambda)$ Functions

We can now reformulate the individual's lifetime maximization problem as:

$$\max_{c, h(\lambda), \lambda^*} u(c) - \int_0^{\lambda^*} v(h(\lambda)) d\lambda$$

$$s.t. \ c = \int_0^{\lambda^*} e(\lambda) g(h(\lambda)) d\lambda$$

where λ^* represents the fraction of lifetime spent in employment. Given a value for λ^* it is straightforward to back out the implied values for A_1 and A_2 . If solutions for both are interior, then they solve:

$$\tilde{e}(A_1) = \tilde{e}(A_2) = e(\lambda^*)$$

Using the budget equation to substitute for c in the objective function, this problem can be reduced to finding a value of λ^* and an hours profile $h(\lambda)$ for $0 \leq \lambda \leq \lambda^*$. Assuming that the solution for λ^* is interior¹⁰ and that $h(\lambda) < 1$ for all λ , we obtain the following first order conditions for λ^* and $h(\lambda)$ for $0 \leq \lambda \leq \lambda^*$:

$$\frac{v(h(\lambda^*))}{u'(\int_0^{\lambda^*} e(\lambda)g(h(\lambda))d\lambda)} = e(\lambda^*)g(h(\lambda^*)) \quad (3.5)$$

$$\frac{v'(h(\lambda))}{u'(\int_0^{\lambda^*} e(\lambda)g(h(\lambda))d\lambda)} = e(\lambda)g'(h(\lambda)) \quad (3.6)$$

Both of these equations have a standard interpretation: labor is adjusted on each margin until the marginal rate of substitution between leisure and consumption is equal to the marginal product of labor along each margin, though only the second equation corresponds to the standard textbook model of labor supply. For the case of hours worked while employed, the relevant margin is how many hours to work at a given instant, so that the marginal disutility associated with an additional hour is equal to $v'(h(\lambda))$, while the marginal product associated with an additional hour is $e(\lambda)g'(h(\lambda))$. For the case of adjusting the length of working

¹⁰Note that the solution for λ^* is interior as long as one of the A_i 's has an interior solution.

life, the marginal disutility associated with increasing the fraction of life spent in employment is $v(h(\lambda^*))$, while the marginal product associated with increasing the fraction of life spent in employment is $e(\lambda^*)g(h(\lambda^*))$.

Although solving the maximization problem requires finding the value λ^* and the entire $h(\lambda)$ profile, we will show that the problem can be reduced to finding two values: λ^* and $h(0)$. To see why this is the case, note that equation (3.6) implies that along the optimal profile $h(\lambda)$, we must have:

$$\frac{v'(h(\lambda))}{e(\lambda)g'(h(\lambda))} = \text{constant for all } \lambda \in [0, \lambda^*] \quad (3.7)$$

If $h(\lambda)$ lies in the region in which g is concave, then the strict convexity of v implies that the left hand side of this equation is strictly increasing in $h(\lambda)$. In this case it then follows that if $h(0)$ is known, the entire $h(\lambda)$ profile is uniquely determined. Additionally, given the monotonicity property just mentioned, an increase in $h(0)$ leads to an upward shift of the entire $h(\lambda)$ profile. A potential problem with this argument is that $h(\lambda)$ does not necessarily lie in the region where g is concave. However, although one cannot guarantee that $h(\lambda)$ lies in the region where g is concave, the second order conditions for the maximization problem nonetheless do require that at an optimum, the left-hand side of equation (3.7) is strictly increasing in h . Loosely speaking, although the optimum need not occur at a point at which g is concave, the convexity of v in hours must dominate the lack of concavity in g .¹¹

¹¹More formally, if one reformulates the problem as having the worker choose how many units of labor services to offer at each instant, the relevant disutility over labor services can be written as $v(g^{-1}(\frac{l}{e(\lambda)}))$ and the second order condition requires that v be convex at the optimum choice

Having established that solving the consumer's problem can be reduced to finding optimal values for $h(0)$ and λ^* , we next show that the equilibrium can be described as the intersection of two curves in $h(0) - \lambda^*$ space, one of which is (at least locally) upward sloping, and the other of which is (globally) downward sloping. To derive the upward sloping relationship, we note that dividing the two first order conditions by each other, evaluating at $\lambda = 0$, and rearranging, one obtains:

$$\frac{v'(h(0))}{e(\lambda)g'(h(0))} = \frac{v(h(\lambda^*))}{e(\lambda^*)g(h(\lambda^*))} \text{ for all } \lambda \in [0, \lambda^*] \quad (3.8)$$

To see that this expression implies a relationship between $h(0)$ and λ^* simply note that given a value of $h(0)$ we can infer the profile $h(\lambda)$ using equation (3.7). Given this profile, we then evaluate the right hand side at each point of the profile to find a value of λ such that equation (3.8) holds. We now show that in a neighborhood of the optimal solution, the relationship between $h(0)$ and λ^* is increasing.

Proposition 3: In a neighborhood of the optimal solution, equation (3.8) defines an increasing relationship between $h(0)$ and λ^* .

Proof: To prove this we establish three properties. First, the left-hand side of equation (3.8) is strictly increasing in $h(0)$. Second, given an optimal profile $h(\lambda)$, the right hand side of equation (3.8) is increasing in λ . Third, the marginal effect of an increase in $h(0)$ on the right hand side of equation (3.8) evaluated at the optimal λ^* is zero. Combining these three properties, the result necessarily

of labor services. This condition implies that the left hand side of equation (3.7) is strictly increasing in h .

follows. Since the first property has already been noted, it remains to establish the second and third properties.

To establish the second property we simply take the derivative of the right hand side of equation (3.8) with respect to λ and evaluate it at $\lambda = \lambda^*$. This yields:

$$-\frac{e'(\lambda^*)}{e(\lambda^*)^2} \frac{v(h(\lambda^*))}{g(h(\lambda^*))} + \frac{1}{e(\lambda^*)} \frac{g(h(\lambda^*))v'(h(\lambda^*)) - v(h(\lambda^*))g'(h(\lambda^*))}{g(h(\lambda^*))^2} h'(\lambda^*) \quad (3.9)$$

However, noting that equation (3.7) holds for $\lambda = \lambda^*$, it follows that the second term is necessarily zero, and that the sign of this expression is therefore simply the sign of $-e'(\lambda^*)$, which is necessarily positive by construction of the $e(\lambda)$ function.

To establish the third property, we simply differentiate the right hand side of equation (3.8) evaluated at λ^* with respect to $h(0)$, and evaluate at the optimal value of $h(0)$. This gives:

$$\frac{1}{e(\lambda^*)} \frac{g(h(\lambda^*))v'(h(\lambda^*)) - v(h(\lambda^*))g'(h(\lambda^*))}{g(h(\lambda^*))^2} \partial h(\lambda^*)/\partial h(0) \quad (3.10)$$

Since the middle term of this expression is equal to zero by equation (3.7), it follows that this derivative is zero.

Having established the three properties, the result follows.//

This upward sloping relationship is intuitive. Equation (3.8) has the interpretation that the marginal disutility per unit of additional output is equal across the two margins $h(0)$ and λ^* . Since this marginal disutility is increasing along both margins, if more time is devoted to market work along one margin, then more

time must be devoted to market work along the other margin as well, in order for the two disutilities to be equated. We will call the curve that captures this relationship the *optimal composition of work curve*.

Similarly, given a value for $h(0)$ and the implied profile for $h(\lambda)$, equation (3.6) evaluated for $\lambda = 0$ gives:

$$\frac{e(0)g'(h(0))}{v'(h(0))} = \frac{1}{u'(\int_0^{\lambda^*} e(\lambda)g(h(\lambda))d\lambda)} \quad (3.11)$$

We now show that this expression produces a negative relationship between λ^* and $h(0)$. To see this, note first that holding λ^* constant, the left hand side of equation (3.11) is decreasing in $h(0)$, while the right hand side is increasing in $h(0)$, since an increase in $h(0)$ shifts the entire $h(\lambda)$ profile upward and g is an increasing function. Second note that holding $h(0)$ constant, the right hand side is increasing in λ^* . It follows immediately that this equation describes a negative relationship between $h(0)$ and λ^* . This relationship is also intuitive. Equation (3.11) has the interpretation that the marginal rate of substitution between consumption and leisure be equated to the marginal product of time devoted to market work. But since an increase in time devoted to market work along one margin leads to higher consumption and therefore raises the marginal rate of substitution, it leads to less time devoted to market work along the other margin. Put somewhat differently, from the perspective of generating income, the two margins are substitutes. We will call the curve that captures this relationship the *optimal volume of work curve*.

Combining these results, it follows that one can represent the problem of solving for the optimal hours profile $h(\lambda)$ and the optimal time spent in employment,

λ^* , as the intersection of two curves in $h(0) - \lambda^*$ space, with one curve being upward sloping and the other curve being downward sloping. Figure 3 depicts the situation.




figure 3: Equilibrium Determination of $h(0)$ and λ^*

This diagrammatic representation will be useful when we consider the effect of tax policies.

4. Tax Policies and Labor Supply Elasticities

In this section we consider how a tax and transfer policy affects the equilibrium hours worked profiles for individuals, both analytically as well as quantitatively. In particular, similar to the tax and transfer policy studied by Prescott (2004), we assume that the government taxes all labor income at the constant rate of τ and uses the tax revenues to fund an equal lump-sum transfer T at each instant

to all individuals, subject to a balanced budget constraint at each instant of time. In steady state, tax revenues will be constant and hence the lump-sum transfer will be constant as well.

4.1. Analytic Results

Given this policy, the individual lifetime maximization problem becomes:

$$\max_{c, h(\lambda), \lambda^*} u(c) - \int_0^{\lambda^*} v(h(\lambda)) d\lambda$$

$$s.t. \ c = (1 - \tau) \int_0^{\lambda^*} e(\lambda) g(h(\lambda)) d\lambda + T$$

The first order conditions for this problem are:

$$\frac{v(h(\lambda^*))}{u'((1 - \tau) \int_0^{\lambda^*} e(\lambda) g(h(\lambda)) d\lambda + T)} = (1 - \tau) e(\lambda^*) g(h(\lambda^*)) \quad (4.1)$$

$$\frac{v'(h(\lambda))}{u'((1 - \tau) \int_0^{\lambda^*} e(\lambda) g(h(\lambda)) d\lambda + T)} = (1 - \tau) e(\lambda) g'(h(\lambda)) \quad (4.2)$$

Noting that the balanced budget rule for the government implies:

$$\tau \int_0^{\lambda^*} e(\lambda) g(h(\lambda)) d\lambda = T \quad (4.3)$$

the two previous equations can be written as:

$$\frac{(1 - \tau) e(\lambda^*) g(h(\lambda^*))}{u'(\int_0^{\lambda^*} e(\lambda) g(h(\lambda)) d\lambda)} = v(h(\lambda^*)) \quad (4.4)$$

$$\frac{(1 - \tau)e(\lambda)g'(h(\lambda))}{u'(\int_0^{\lambda^*} e(\lambda)g(h(\lambda))d\lambda)} = v'(h(\lambda)) \quad (4.5)$$

As before, equation (4.5) implies that along the optimal profile, $e(\lambda)g'(h(\lambda))/v'(h(\lambda))$ will be constant, implying that one can infer the entire $h(\lambda)$ profile from the value of $h(0)$. Moreover, using this fact and proceeding as before, one can show that equilibrium can be summarized by the following two relations, each of which describes a relationship between $h(0)$ and λ^* :

$$\frac{v'(h(0))}{e(0)g'(h(0))} = \frac{v(h(\lambda^*))}{e(\lambda^*)g(h(\lambda^*))} \quad (4.6)$$

$$(1 - \tau)\frac{e(0)g'(h(0))}{v'(h(0))} = \frac{1}{u'(\int_0^{\lambda^*} e(\lambda)g(h(\lambda))d\lambda)} \quad (4.7)$$

The first equation is the same expression as in the zero tax case, while the second expression now contains a term $(1 - \tau)$. It is straightforward to establish that an increase in τ implies a lower value of λ^* for a given value of $h(0)$. It follows that the optimal composition of work curve is unaffected by taxes, while an increase in taxes shifts the optimal volume of work curve downward. The following proposition is then immediate:

Proposition 4: An increase in τ leads to a decrease in λ^* and a downward shift in $h(\lambda)$ for all $\lambda \in [0, \lambda^*]$.

Proof: Follows immediately from Figure 3, and the previous result concerning the effect of τ on the two curves. //

It is interesting to contrast this result with that which obtains in the case where the $e(\lambda)$ profile is flat. In this case an increase in taxes leads to a lower

fraction of life spent in employment but no change in the hours profile.

4.2. Quantitative Results

In this subsection we report the results of some numerical simulations regarding the effects of taxes on hours of work.

4.2.1. Calibration

For these calculations we adopt the following functional forms:

$$\begin{aligned} u(c) &= \log(c), \quad v(h) = \alpha \frac{h^{1+\gamma}}{1+\gamma}, \\ g(h) &= (h - \bar{h}) \text{ for } h \geq \bar{h}, \quad 0 \text{ otherwise, } e(\lambda) = e_0 - (e_0 - e_1)\lambda. \end{aligned} \tag{4.8}$$

The function $u(c)$ is chosen in order to be consistent with balanced growth in a setting with technological progress, though we abstract from growth here. The choice of $v(h)$ is standard and is convenient since the parameter γ determines the elasticity of hours with respect to the tax rate in a standard labor supply model in which $g(h) = h$. Our choice for the function g is dictated by parsimony. The exact shape of g for $h \leq \bar{h}$ is not important as long as the individual chooses to work $h \geq \bar{h}$, and so we set $g = 0$ in this region. We also abstract from fatigue effects in this specification. The assumed functional form for $e(\lambda)$ implies a linear productivity profile, with initial productivity e_0 and final productivity e_1 . While the data suggests that linearity is not necessarily a good assumption for this profile, we adopt it because it permits a parsimonious way to investigate the importance of the slope of this profile in affecting how hours and employment

respond to changes in taxes. It is easy to show that with our assumption on preferences, specifically that utility from consumption is $\log(c)$, the solution for $h(\lambda)$ is unaffected by a proportional shift in the $e(\lambda)$ profile, so that we can normalize e_0 to one with no loss in generality.

Given these functional forms, we will investigate how the parameter γ matters for some properties of the life cycle profile and how this profile responds to tax and transfer policies. For each value of γ we will choose values for the three parameters α , $\bar{h}$, and e_1 so as to match three target values. The first target is the fraction of life spent in employment. If we interpret our model as representing an adult life span of 60 years, then a working life of approximately 40 years implies a target value for λ^* of .67. If the peak workweek for employed workers over the life cycle is around 45 hours per week and individuals have roughly 100 hours of discretionary time per week, then recalling that we normalized the time endowment to one at each instant, the target value for $h(0)$ is .45. Lastly, given a target value for λ^* , the value of e_1 will influence the range of productivities over the life cycle, and hence the range of hourly wages. We choose a value for e_1 so that hourly wages at their peak are twice as large as hourly wages at their lowest point. Because the results that we report below are very robust to changes in these values, we do not focus on justifying these exact values.

We note that the calibrated value of $\bar{h}$ is increasing in γ . This is intuitive. To see this note, that the greater the value of $\bar{h}$ the greater is the nonconvexity in the mapping from hours to labor services, and that it is this nonconvexity that induces retirement in the model. The higher the value of γ , the less the individual

likes to have hours change over the life cycle. Because retirement is necessarily associated with a large change in hours of work, it follows that it requires a greater nonconvexity to induce retirement for higher values of γ . Further discussion of this issue is postponed until the next subsection.

It is important to point out that some care needs to be taken in matching up wages in the model with wages in the data. In the model, the wage per unit of labor services, which we denoted by w , is equal to one at all points in time. But wages in the data are measured as labor earnings per hour of work, and so we compute this same measure in our model. We denote this wage rate as w^h , where the superscript h denotes that we are measuring wages per hour of work. If the function g were identically equal to one then earnings per hour of work in the model would be exactly equal to $e(\lambda)$ and hence the range of wages over the life cycle would be exactly equal to $e(0)/e(\lambda^*)$. But since the function g is non-linear, this no longer holds. The range of wages over the life cycle in our calibrated model is given by:

$$\frac{w^h(0)}{w^h(\lambda^*)} = \frac{e(0)g(h(0))/h(0)}{e(\lambda^*)g(h(\lambda^*))/h(\lambda^*)} \quad (4.9)$$

While the hourly wage ratio is influenced by $e(0)/e(\lambda^*)$, these values are no longer identical. Nonetheless, one can choose e_1 such that this ratio is equal to 2 in the benchmark calibration.

In calibrating the model, we also assume a tax rate of .3, which corresponds to the average effective tax on labor income in the US in recent years. Having calibrated the model, we will then examine what happens to equilibrium hours if the tax rate were increased to .5, which corresponds to the average effective tax

on labor income in several economies in continental Europe in recent years.¹²

4.2.2. Micro Elasticities

Before reporting the results of the change in tax and transfer policies, it is of interest to examine some features of the calibrated benchmark economies. Given a value of γ and the calibration procedure just described, the model will generate a life cycle profile for hours worked, $h(\lambda)$, and hourly wages, $w^h(\lambda)$. We generate a panel life cycle data set for hourly wages and hours worked by choosing 67 equally spaced values of λ , running from 0 to .66 and evaluating the two functions $h(\lambda)$ and $w^h(\lambda)$ at these points. Note that all of the data points in the sample are times at which individuals are employed. Using t to index the data points for a given individual, as is standard in the labor supply literature we take this data and run the regression:

$$\log(h_t) = b_0 + b_1 \log(w_t^h) + \varepsilon_t \quad (4.10)$$

The resulting parameter estimate b_1 is the so-called micro Frisch labor supply elasticity.

Table 1 shows the estimated values of b_1 for our benchmark calibrated model for four different values of γ : .5, 1, 2, and 10.

¹²Several authors have produced estimates of effective tax rates for various countries, including Mendoza et al (1994), Prescott (2004) and McDaniel (2006). While there are small differences in methodology across studies, the 20% differences between the US and countries such as Belgium, France, Germany and Italy is a robust finding.

Table 1

Estimated Micro Frisch Elasticities			
$\gamma = .5$	$\gamma = 1$	$\gamma = 2$	$\gamma = 10$
1.29	.59	.28	.05

The table shows that higher values of γ are associated with lower Frisch elasticities, though note that the nonlinearity of the g function implies that the Frisch elasticity is not equal to $1/\gamma$. In particular, the nonlinearity of g implies that higher hours imply higher hourly wage rates, thereby lowering the estimated elasticity relative to a standard model. The “bias” induced by the nonlinearity of the g is substantial. The values of γ are only about 50 – 60% of the values that one would infer based on a linear specification.

There is a voluminous literature that has estimated Frisch elasticities using variation in hours and wages over the life cycle. Early examples include Ghez and Becker (1975), MaCurdy (1981), and Heckman and MaCurdy (1980). The early literature found relatively small estimates for males, on the order of .3 or less, but much larger values for women. Subsequent work, including recent papers by Kimball and Shapiro (2003), Pistaferri (2003) and Domeij and Floden (2006) have refined these estimates in various ways, and found larger estimates, in the range of .7 – 1.0 for males. (See Hall (2007) for a critical survey of the recent literature.)

Before proceeding, we revisit the issue of the size of the value of nonconvexity necessary to induce retirement in each of the specifications. When $\gamma = .5$, the calibrated value of $\bar{h}$ is .06, while when $\gamma = 10$ the calibrated value is .39. When $\gamma = 1$ the calibrated value for $\bar{h}$ is .15. In interpreting these values for $\bar{h}$ it

is important to keep in mind that these calculations assume a single source of nonconvexity. As noted earlier, it is reasonable to also consider nonconvexities in the mapping from leisure time to leisure services, and if one included this factor the required values of $\bar{h}$ would obviously be lower. With this in mind, we would suggest that for values of γ less than 1 the required nonconvexities do not seem unreasonably large. In contrast, the values of $\bar{h}$ required when γ takes on a value such as 10 seem beyond reasonable. Given that recent micro estimates summarized above suggest values of γ that are less than 1, one could simply view this as an additional piece of evidence that such high values of γ are hard to reconcile with the data. In particular, any argument in support of values of γ that are as high as 10 would have to face the challenge of explaining how such a value is consistent with observed retirement patterns. Having noted this qualification regarding the cases of high γ values, in what follows we will report results for the full range of values considered above.

4.2.3. Tax and Transfer Policies

We now turn to the evaluation of tax and transfer policies. For each of the four different calibrated economies (one for each of the four values for γ in Table 1), we consider what happens to the steady state hours profile if we increase the tax rate on labor income from .3 to .5, assuming that the proceeds fund a uniform lump-sum transfer to all individuals subject to a balanced budget constraint at each point in time. With our functional forms, one can show that such a tax causes a proportional shift in the hours profile, conditional on being employed. It

follows that one can summarize the shift in the hours profile by simply reporting the shift in $h(0)$. For each economy we compute the values of aggregate hours (H), fraction of life spent in employment (λ^*), and peak hours worked over the life cycle ($h(0)$), all relative to the values in the benchmark calibrated economy with $\tau = .3$. Table 2 reports the results.

Table 2			
Relative Outcomes for $\tau = .5$			
γ	H	λ^*	$h(0)$
.50	.777	.857	.856
1.00	.784	.825	.918
2.00	.788	.808	.956
10.00	.790	.794	.991

Several features are worth noting. First, note that the implied change in aggregate hours worked is large in all four cases—more than 20%. Second, despite the dramatic differences in estimated Frisch elasticities in the four economies—a factor 25 difference between the highest and lowest—the changes in aggregate hours worked are essentially constant across the four different economies. Third, although the value of γ has virtually no effect on the change in aggregate hours worked, it has very significant effects on how the change in aggregate hours is broken down into changes in working life versus changes in hours worked while employed. In analyzing this decomposition, it is important to note that the relative change in $h(0)$ is a measure of the change in total hours due to changes in the h profile holding λ^* constant, since as noted earlier, the h profile shifts propor-

tionately, and for a given λ^* , a proportionate shift in the profile shifts aggregate hours by the same amount. However, it is not true that a shift in λ^* leads to a proportionate shift in aggregate hours, since as λ^* decreases the marginal employment episodes that are lost represent fewer hours of work. In any case, when $\gamma = .50$ the downward shift in the hours profile accounts for over 60% of the total decrease in hours, while when $\gamma = 10$ this downward shift accounts for less than 5% of the shift.

There are two additional implications of the model not reported in Table 2 that are also of interest. Because time devoted to work and labor services are not proportional in our model, it is of interest to contrast the effects of taxes on aggregate labor services with the effect on aggregate time devoted to work. Additionally, although changes in taxes do not affect technology in our analysis, they can affect productivity measures such as output per hour because of the difference between labor services and time devoted to market work. Note that there are two opposing effects of higher taxes on productivity per hour in our model. On the one hand, the decrease in hours is concentrated among lower productivity workers since this is where the extensive margin is operative, which would lead to higher output per hour in the high tax economy. On the other hand, higher taxes shift the hours profile down, thereby lowering the ratio of labor services to hours, and leading to lower productivity. The importance of these two effects is influenced by the relative size of adjustment along the intensive and extensive margin, and hence by the value of γ . However, it turns out that these effects are relatively small in our numerical examples. For all four economies the

increase in taxes is associated with a drop in output per hour, but the decrease is less than 1%, and ranges between .9% and .6% as γ is varied from 10 to .5.

While Table 2 contrasted outcomes for just two different tax rates for a range of values of γ , we note that the effects are very close to linear in the tax rate. For completeness, Table 3 presents results for a range of tax rates for the specification that corresponds to $\gamma = 2$.

Table 3						
Tax Effects ($\gamma = 2$, all values relative to $\tau = .30$)						
	$\tau = .35$	$\tau = .40$	$\tau = .45$	$\tau = .50$	$\tau = .55$	$\tau = .60$
H	.950	.899	.844	.788	.729	.665
λ^*	.956	.910	.860	.808	.752	.692
$h(0)$	.990	.979	.968	.956	.945	.933

4.2.4. Comparison With a Stand-In Household Economy

The model economy that we have studied is not a single agent economy, in the sense that at any point in time there are many different types of individuals alive. However, it is interesting to ask what one might infer about labor supply if one were to interpret the outcomes generated by the tax changes in our model by using a standard static stand-in household model. In particular, consider a static economy with a single agent, with preferences given by:

$$\log(c) - \mu \frac{h^{1+\theta}}{1+\theta} \tag{4.11}$$

There is a linear technology in this economy that can turn one unit of time into one unit of consumption:

$$c = h$$

and there is a government that taxes labor at the constant proportional rate of τ and uses the proceeds to fund a lump-sum transfer to the representative agent.

Faced with the information in Table 2, we ask what an economist using this model to interpret the hours differences would conclude about the parameter θ that dictates the labor supply elasticity for the stand-in household in this economy. Standard calculations for this economy lead to the following expression for hours of work in terms of taxes:

$$h = \left[\frac{1 - \tau}{\mu} \right]^{1/(\theta+1)} \quad (4.12)$$

If we let h_i denote the hours that correspond to a country with tax rate τ_i , for $i = 1, 2$, then using the above expression to interpret data on taxes and hours of work leads to the following expression for θ :

$$\theta = \frac{\log(1 - \tau_1) - \log(1 - \tau_2)}{\log(h_1) - \log(h_2)} - 1 \quad (4.13)$$

Applying this expression to the four calibrated economies, we obtain the results shown in Table 4:

Table 4

Implied Values for θ			
$\gamma = .5$	$\gamma = 1$	$\gamma = 2$	$\gamma = 10$
.33	.38	.41	.43

The associated Frisch elasticities, given by $1/\theta$, range from 2.3 to 3, despite the fact that the Frisch elasticities inferred from micro data range from .05 to 1.25.¹³

The above calculation shows that a static stand-in household model with a fairly high labor supply elasticity can reproduce the steady state effects of taxes on aggregate hours found in the life cycle model studied earlier. It is also of interest to ask whether the welfare implications of tax changes are similar across the two specifications. Our measure of welfare is the percent increase in lifetime consumption required to make households living in the high tax economy indifferent to living in the low tax economy. It turns out that the answers are remarkably similar in the life cycle and stand-in household models. For example, in the $\gamma = 1$ economy the welfare cost of the higher tax system is 10.7% of consumption, while in the corresponding stand-in household model the welfare cost is 10.4% of consumption.

4.2.5. The Role of $g(h)$

The above results indicate that in our life cycle economy, micro labor supply elasticities are not particularly relevant in predicting the aggregate effects of per-

¹³One can show that our steady equilibrium corresponds to the allocation that maximizes an equal weighted sum of individual utilities. It follows that one can derive an analytic expression for preferences of the stand-in household.

manent changes in taxes. It is important to emphasize the feature of the economy that is responsible for this result. In particular, the mere fact that our economy is an overlapping generations model is not important in generating this result. Rather, the key feature of our economy is the nonlinear mapping from time spent working to labor services, and the fact that this feature generates a life cycle profile for hours worked with hours equal to zero for some parts of the lifecycle. To understand this, consider an economy that is identical to the one that we have studied except assume that the function g is identically equal to one. Figure 4 illustrates how this will influence the findings.




Figure 4: Changes in Life Cycle Hours

In this figure, the top line shows the life cycle productivity profile. The two solid lines indicate the life cycle profile for hours worked in the case of linear and nonlinear g . As the picture shows, if g is nonlinear then we can generate outcomes in which hours worked are concentrated in the period of life in which productivity

is highest. In particular, hours worked are not continuous in productivity. In contrast, if g is linear, it is optimal for the individual to smooth hours worked across time, although hours of work will be higher when productivity is higher. But in this case hours vary continuously with productivity. The two dashed lines indicate the effects of higher taxes on hours of work in the two cases. If g is nonlinear, then the hours worked profile shifts down and the reservation productivity level shifts up, while in the case of a linear g function, the only effect is a downward shift in the hours profile. In both cases the extent of the downward shift of the hours profile is very strongly related to the micro labor supply elasticity. Because this downward shift is the only effect when g is linear, it turns out that there is a strong relationship between micro and macro elasticities in this case.

4.2.6. The Importance of Heterogeneity

Previous work on the implications of labor indivisibilities for aggregate labor supply elasticities has stressed that heterogeneity may have a large influence on the implied aggregate labor supply elasticity. This argument appears in different contexts in both Mulligan (2001) and Chang and Kim (2006a, 2006b). Mulligan argues that increasing heterogeneity in preferences for consumption versus leisure serves to decrease the implied aggregate labor supply elasticity in a model with indivisible labor. Chang and Kim consider a model in which all individuals have identical preferences, but face idiosyncratic wage shocks and do not have access to any insurance markets. As a result, individuals accumulate assets in order to

self-insure against these shocks. This model yields a non-degenerate distribution of assets across individuals in steady state. In this setting, Chang and Kim show that there is a mapping from assets into reservation wages, i.e., for any asset position there is a wage such that if the individual faces a wage higher than that level, then they will choose to be employed. Chang and Kim go on to show that the distribution of reservation wages plays a key role in determining the implied aggregate labor supply elasticity in their model.

Given these findings, it is of interest to examine the importance of heterogeneity for the aggregate response of hours to taxes in our model. The key dimension of heterogeneity in the cross-section of our steady state is the distribution of individual productivities across consumers. This heterogeneity in the cross-section is exactly the same as the heterogeneity that a given individual faces over their lifetime, and hence is characterized by the function $e(\lambda)$. In our calibrated examples considered above, we chose the value of e_1 so as to achieve a given degree of heterogeneity in wages both in the cross-section and over the life cycle. In fact, the variation in e_1 across the four different cases was quite small, ranging from .46 for $\gamma = .5$ to .522 for $\gamma = 10$. To explore the importance of heterogeneity for our results, in this subsection we continue to choose values for $\bar{h}$ and α so as to match values for λ^* and $h(0)$, but will no longer calibrate the value of e_1 to target a particular range of wages. Instead, we simply consider a range of values for e_1 and report the results for the effect of an increase in taxes from .3 to .5 in each case. We will also report the implications for the amount of cross-sectional heterogeneity across employed workers.

Table 5 shows the implied effects of differences in e_1 from .1 to .9 for the case of $\gamma = 2$.¹⁴

Table 5

Heterogeneity and Relative Outcomes ($\gamma = 2$)					
e_1	H	λ^*	$h(0)$	$w(\lambda^*)/w(0)$	b_1
.90	.728	.732	.991	.875	.26
.70	.758	.770	.972	.664	.27
.50	.789	.810	.956	.493	.28
.30	.820	.852	.941	.351	.30
.10	.849	.893	.928	.238	.32

One issue to note up front is that although we are holding γ constant in this exercise, it does not follow that the estimate of the Frisch elasticity from the life cycle profile is necessarily constant. The last column of Table 5 reports the estimated Frisch elasticities for the different values of e_1 . Although the values are influenced by e_1 , the range of estimates is not very large. Turning to the results, the second through fourth columns report the same information that we have focused on before—the relative values for aggregate hours, time spent in employment and peak hours of work. As e_1 is decreased, the amount of heterogeneity, measured either as the range of e values across individuals or the range of hourly wages across individuals, increases. Consistent with the findings of previous researchers, this increase in heterogeneity does reduce the aggregate response for a

¹⁴Recall that a proportional shift in the $e(\lambda)$ profile has no effect on hours of work in our model, so that it is the ratio e_1/e_0 that matters and not $e_0 - e_1$.

given increase in tax rates and a given value of γ . However, while the effect is significant, it should be emphasized that relative to our benchmark calibration, even if we were to consider a factor 5 increase in the range of e values in the cross-section (or a doubling of the range of cross-sectional hourly wages), the aggregate consequences are still very large—a 20% increase in taxes still leads to a decrease in hours of work of more than 15%.

It is important to note that the other studies that we referred to were based on indivisible labor models, i.e., they assumed that hours of work conditional upon employment were exogenous and equal for all workers. In our model, there is a nonconvexity that leads to outcomes that are similar to what happens with indivisible labor, but hours of work are endogenous and respond to changes in the environment. This is significant, since a comparison of the results for relative values of λ^* and $h(0)$ shows that as e_1 decreases, the drop in λ^* becomes smaller, but the drop in $h(0)$ actually increases. When $e_1 = .9$, the drop in $h(0)$ is basically one percent, while when $e_1 = .1$ the drop in $h(0)$ exceeds seven percent. Precisely because of this opposing effect on $h(0)$, a comparison of the second and third columns indicates that changes in e_1 have a much larger effect on λ^* than on H . Two important messages follow from these results. First, an indivisible labor model may preclude an important margin of adjustment in some contexts. Second, the effect of taxes on hours of work for employed individuals is not purely a function of the parameter γ —this table illustrates that the response in $h(0)$ is jointly determined with the response in λ^* .

There is one additional point of interest to note concerning heterogeneity. In

our model we assumed that there is no heterogeneity within a cohort. One may conjecture that adding heterogeneity within a cohort will further diminish our aggregate effects. However, at least one form of within cohort heterogeneity will have no impact on our findings. Specifically, assume that within each cohort there is a distribution of permanent productivities, represented as proportional shifts of the productivity profile $e(\lambda)$. As noted earlier, with $u(c) = \log(c)$, proportional shifts of the productivity profile have no impact on the lifetime hours profile, and so this form of heterogeneity would have no impact on our findings.

4.3. An Alternative Source of Life Cycle Variation

In the preceding analysis we assumed that the driving force for life cycle variation in hours of work was exogenous variation in individual productivity over the life cycle. Though this assumption underlies much of the literature on estimating labor supply elasticities using micro data, there is probably good reason to question the reasonableness of the assumption that productivity will double from young to middle age independently of labor supply decisions over that period. Recent work by Imai and Keane (2004) departs from this assumption by considering the possibility that future wages are influenced by human capital that is accumulated via a learning by doing process. While it is of interest to examine human capital accumulation in the context of our model, we leave this possibility for future work. But in this subsection we briefly describe the results that emerge from considering a different driving force for life cycle variation in hours worked. Specifically, rather than assuming that productivity varies with age we assume that the disutility of

work varies exogenously with age, so that preferences are now given by:

$$\int_0^1 [u(c(a)) - \tilde{e}(a)v(h(a))]da \quad (4.14)$$

where we assume that the profile $\tilde{e}(a)$ is U -shaped. The mapping from time spent at work to labor services is now given by:

$$l = g(h) \quad (4.15)$$

where g is assumed to have the same shape as before.

Analogous to before, the optimal labor supply decision entails a reservation disutility level, and hours worked when employed will be decreasing in the level of disutility. Although there is no exogenous variation in productivity over the life cycle, one can show that there will still be a positive relationship between hours and wages per unit of time because of the properties of the g function. It follows that this specification can still account for the standard life cycle patterns in hours of work and wages. Moreover, at a qualitative level equilibrium in this model is determined just as before—one can again represent the equilibrium as the intersection of two curves in $h(0) - \lambda^*$ space.

When we carry out the numerical analysis analogous to that carried out above, we get virtually identical results, and so in the interest of space we do not go into them in any detail. Specifically, we find that the aggregate effect of the increase in taxes is both large and virtually independent of the Frisch elasticity estimated from life cycle data for prime aged individuals. One difference between

the two specifications is that it takes considerably greater variation in $e(\lambda)$ in the variable disutility case than it does in the variable productivity case to generate the factor-two variation in wage rates, owing to the fact that variable disutility affects wages only indirectly via the $g(h)$ function, while in the case of variable productivity there is also a direct effect. In line with our earlier results about the role of increased heterogeneity, we find that subject to matching a wage range of factor 2 over the life cycle, the aggregate effects are slightly smaller than in the variable productivity case, yielding relative hours of about .82 rather than .78. The point that we want to stress is simply that the previous results are robust to this alternative driving force for life cycle variation in hours worked.

5. Taxes and Market Work: Continental Europe and the US

In this section we use the life cycle model developed earlier to discuss the ability of tax and transfer policies to account for observed differences in hours worked between the US and several countries in continental Europe.

5.1. Hours and Employment in Europe and the US

It is well known that hours of market work per person of working age are much lower in continental Europe than in the US. In this section we present data to establish two further properties. First, the large differences in total hours is the result of important differences along two margins: the employment to population ratio and annual hours worked per person in employment. Second, the differences

in employment to population ratios are due almost exclusively to differences in this ratio for young and old workers. We deal with each of these in turn.

Cross-country data sets on hours of work allow one to decompose total annual hours of work into two components: the number of people employed, and the annual hours worked per person in employment. In making cross-country comparisons it is necessary to normalize employment relative to some measure of population. In what follows we use the size of the population aged 15-64, though the empirical findings are not sensitive to this choice. Table 6 shows the relative values for aggregate hours per person aged 15-64 (Hours/Pop), the employment to working age population ratio (Emp/Pop), and annual hours of work per person in employment (Hours/Emp) for four economies in continental Europe relative to the US. These data are for the year 2003.

Table 6

Market Work in Europe Relative to the US				
	Belgium	France	Germany	Italy
Hours/Pop	.71	.68	.73	.69
Emp/Pop	.83	.88	.91	.79
Hours/Emp	.86	.77	.80	.87

The first row shows the well-known fact that market work in these economies is only about 70% of the US level. The next two rows show that this large difference in market work is the combination of large differences both along the employment and the hours per employee margin. On average across the four economies, the hours per person in employment margin accounts for slightly more than half of the

aggregate difference. While some of the differences both among these countries and between these countries and the US is due to differences in the volume of part time employment, this is not the dominant difference in differences in hours per person in employment.

Next we examine the differences in employment rates by age. Figure 4 shows the employment rate by age relative to the US for each of the four economies in Table 6.




Figure 5: Life Cycle Employment Profiles Relative to the US

The important point to note here is that for Belgium, France and Germany, the employment rates for prime aged workers are nearly the same as those in the US. But for younger (15-24) and older (55-64) individuals, the differences are dramatic. While relative employment rates for prime age individuals are somewhat lower for Italy, it remains true in this case that the differences are most dramatic for young and old individuals. We conclude that the employment differences are heavily concentrated among younger and older individuals.¹⁵

¹⁵While it would be of interest to examine data for annual hours of work per person in employment by age, this data is not available in the cross-country data sets.

5.2. Implications of the Life Cycle Model

We start this section by reminding the reader of the results that have been generated in the literature concerning taxes and differences in hours of work between Europe and the US in the context of a standard (i.e., infinitely lived agent and divisible labor) model. As shown in Prescott (2004), differences in tax and transfer policies of the type considered in the previous section can generate large differences in hours of work across economies. In particular, assuming that preferences were log linear in consumption and leisure, Prescott shows that differences in tax rates on the order of 20% can account for most of the differences in aggregate hours worked between Europe and the US. Such a model necessarily abstracts from the issue of how the hours differences are decomposed into differences in employment rates and differences in hours per employee, and to what extent the differences vary over the life cycle. Given that the data reveals some important patterns across these other dimensions, it seems of interest to seek to understand these differences in addition to those at the aggregate level. In principle, these patterns may contain additional information that helps us evaluate the mechanism at work in the standard model.

Based on the analysis in the previous section, it should be clear that our life cycle model is able to generate not only large differences in aggregate hours in response to differences in tax and transfer policies, but also that it qualitatively matches the patterns found in the more disaggregated data as well. In particular, we showed in the previous section that an increase in taxes will both lower the hours profile and decrease the fraction of life spent in employment, implying that

there is a response along both margins. Moreover, given that the age profile for λ is single peaked, this necessarily implies that the employment differences are concentrated among young and old workers. It follows that at a qualitative level, the extension of the analysis to a life cycle setting with a nonconvexity present in the relationship between the provision of time and labor services is successful in helping us understand the additional patterns in the data.

However, although the extension is successful at a qualitative level, the quantitative analysis in the previous section raises some concern about the ability of the model to match both standard micro elasticity estimates as well as generate changes on the employment and hours per employee margins that are consistent with the cross-country differences. In particular, the cross-country differences suggest that the employment and hours per employee margins are of roughly equal importance. If we look at the results in Table 2, one sees that in response to higher tax rates, our model will generate roughly equal responses in the two margins only if the value of γ is relatively small, on the order of .5. Referring to Table 1, this value of γ results in a Frisch elasticity based on life cycle data that exceeds 1.25. Such a value is certainly large relative to estimates based on male labor supply.

The preceding discussion raises the issue that the model may have difficulty in reconciling the relative variation in time devoted to market work along the two margins. In the remainder of this section we propose a modification that can qualitatively address this issue.

5.3. Retirement Policies

The tax and transfer policy that we analyzed in the previous section assumed that workers receive a transfer at each instant of their life. It should be apparent, however, that it is the present value of the transfer that matters, and not the timing of the transfer. As a result, our previous analysis can also be used to assess the consequences of a retirement program in which all workers are taxed at the proportional rate τ while working, and then receive a constant transfer payment from some age $\bar{a}$ onward, assuming that the retirement benefit is independent of the labor supply decision. Our analysis can also be applied to the case in which both types of transfers exist—one component that the individual receives at each instant, and another component that is received as a retirement benefit. Since retirement programs are one of the largest tax and transfer programs run by most governments, it is important to know that our framework can be applied to this type of program.

An important next step in the research program that analyzes the effects of tax and transfer schemes on cross-country differences in market work is to perform a thorough quantitative analysis of how various features of retirement contributions and benefits influence lifetime labor supply decisions. Our goal here is much more modest—we show how one particular feature of retirement programs may be important for the relative division of hours differences along the employment and hours per employee margin. The feature that we focus on is the fact that many retirement programs contain a provision that imposes a minimum number of years of full-time work in order to qualify for full benefits. Of course, if such

a provision is present, it does not follow that it will necessarily bind—this will depend upon the costs and benefits to the individual of choosing a working life that does not meet this requirement. In what follows here, we will simply ask what the consequences of a binding constraint are and not deal with the issue of under what circumstances the constraint might bind.

In order to understand how a tax and transfer program with a restriction on required years of employment impacts on lifetime labor supply, it is instructive to first consider the effect of a policy that simply stipulates the length of working life in the absence of any taxes or transfers. To this end, consider a policy which stipulates the length of working life to be $\bar{\lambda}$. The next proposition describes how this policy effects the hours profile.

Proposition 6: A change in $\bar{\lambda}$ that is binding leads to a shift in the opposite direction for $h(0)$, and hence the entire hours profile.

Proof: The impact of such a policy can be derived by setting $\lambda^* = \bar{\lambda}$ and using the first order condition for the $h(\lambda)$ profile to determine the effect on hours. It remains true that the entire $h(\lambda)$ profile can be inferred from the value of $h(0)$, and that a higher value of $h(0)$ implies a higher value for all other $h(\lambda)$ as well. To determine whether $h(0)$ shifts up or down, we note that following expression must hold:

$$1 = \left[\int_0^{\bar{\lambda}} e(\lambda)g(h(\lambda))d\lambda \right] \frac{v'(h(\lambda))}{e(\lambda)g'(h(\lambda))} \quad (5.1)$$

Holding the h profile fixed, an increase in $\bar{\lambda}$ leads to an increase in the right hand side. Since the right hand side is increasing in $h(0)$ for a given $\bar{\lambda}$ it follows that the hours profile must shift downward.//

This result is intuitive—if you force the individual to spend a greater fraction of their life in employment than is optimal, they respond by decreasing the time devoted to work while employed. A change in $\bar{\lambda}$ is effectively a move along the volume of work curve in Figure 3. The significance of this result is that it shifts work between the two components.

If one combines a tax and transfer program as studied earlier, with a binding restriction on fraction of life spent in employment, then one can reduce the reduction along employment margin and increase the reduction along the hours per employee margin. In the remainder of this section we report the results of a simple example based on our earlier computations to illustrate the effects of this combination of policies. In particular, we consider the calibration in the previous section that corresponds to $\gamma = 1$, which reflects a moderate value for Frisch elasticity estimated from individual level data. In terms of the previous discussion, in the case of only a tax and transfer policy, this specification implies that the reduction in hours per employee relative to employment is too small relative to the cross country data. In our previous calculations, we found that increasing the tax rate from .3 to .5 lead to a decrease in λ^* from .67 to .55. Table 7 shows the effects of having a tax rate of $\tau = .5$ coupled with a policy which leads to a binding restriction on λ^* that ranges from .57 to .61. As before, all values are relative to the corresponding values in the $\tau = .3$ case.

Table 7

Effect of Working Life Constraints			
$\bar{\lambda}$	H	λ^*	$h(0)$
.55	.784	.825	.918
.57	.802	.855	.911
.59	.819	.885	.904
.61	.836	.915	.898

The first row in this table simply repeats the results from the case in which there is no constraint on working life. As $\bar{\lambda}$ is increased from this value, we see that the reduction in λ decreases, while the reduction in $h(0)$ increases, as implied by our analytic results. It is of interest to note that the reduction in H is also reduced by the increase in $\bar{\lambda}$, so that while this policy does serve to produce changes in λ and $h(0)$ that are more in line with cross country observations, one of the consequences is to imply somewhat smaller responses at the aggregate level.

6. Conclusion

In this paper we develop a general equilibrium life cycle model of labor supply that incorporates both intensive and extensive margins of labor supply. In the equilibrium of our model, individuals have well-defined working lives, in the sense that they enter the workforce at some point in their life and then work continuously until some later point, at which time they withdraw from employment and do not work again. We then use this model to analyze the implications for observed

differences in tax and transfer programs between the US and several countries in continental Europe. In the context of this exercise we can use our model to compute micro labor elasticities using life cycle variation in hours and wages for prime age workers, as well as macro labor elasticities using variation in aggregate hours across economies. Our analysis produces five main findings. First, macro elasticities and micro elasticities are virtually unrelated: a factor 25 difference in micro elasticities is associated with only a thirty percent change in the associated macro elasticities. Second, macro elasticities are large—in the range of 2.3 – 3.0. Third, in our model with variation in either productivity or disutility of work over the life cycle, tax and transfer programs necessarily imply that higher taxes lead to less work on both the extensive and intensive margin. Fourth, the employment differences generated by differences in tax and transfer programs are necessarily concentrated among young and old workers. Fifth, the assumed nonlinearity of labor services in work hours implies a significant bias in the mapping from estimated Frisch elasticities to the preference parameter governing curvature over disutility of work.

There are many natural extensions of interest. In terms of understanding life cycle variation in wages and hours of work it is of interest to consider alternatives which stress endogenous accumulation of human capital.¹⁶ In terms of assessing the implications of tax and transfer programs for hours of work it is of interest to consider a richer description of how tax and transfer programs interact with age and productivity. Lastly, while our analysis has addressed the issue of reconciling

¹⁶Two recent examples are Imai and Keane (2004) and Ljungqvist and Sargent (2007). The latter also examines the implications for tax and transfer programs.

micro and macro elasticities in the context of permanent differences in tax and transfer programs, there is a related issue of reconciling micro and macro elasticities in the context of business cycle fluctuations which our analysis does not specifically address.¹⁷

References

- [1] Alesina, A., Glaeser E., and Sacerdote B., “Work and Leisure in the US and Europe: Why So Different?,” NBER Macroeconomics Annual, 2005.
- [2] Chang, Y., and Kim S., “Heterogeneity and Aggregation in the Labor Market: Implications for Aggregate Preference Shifts,” mimeo, Seoul National University, 2006b.
- [3] Chang, Y., and Kim S., “From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy,” International Economic Review 47(2006a), 1-27.
- [4] Conesa, J., and Kehoe T., “Productivity, Taxes and Hours Worked in Spain, 1970-2003,” , mimeo, 2005.
- [5] Davis, S., and Henreksson M., “Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich Country Comparisons”, NBER Working Paper #10509, 2004.

¹⁷Chang and Kim (2006) pursue this issue in a model with incomplete markets, idiosyncratic shocks and indivisible labor.

- [6] Domeij, D., and Floden M., “The Labor Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased,” *Review of Economic Dynamics* 9 (2006), 242-262.
- [7] Ghez, G. and Becker, G., *The Allocation of Time and Goods Over the Life Cycle*, New York, NBER, 1975.
- [8] Hall, R., “Sources and Mechanisms of Cyclical Fluctuations in the Labor Market,”, working paper, Stanford University, 2007.
- [9] Heckman, J., “Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption Over the Life Cycle,” *American Economic Review* 64 (1974), 188-194.
- [10] Heckman, J., “What Has Been Learned About Labor Supply in the Past Twenty Years?,” *American Economic Review* 83 (1993), 116-121.
- [11] Heckman J., and MaCurdy, T., “A Life Cycle Model of Female Labour Supply,” *Review of Economic Studies* 47 (1980), 47-74.
- [12] Imai, S., and Keane M., “Intertemporal Labor Supply and Human Capital Accumulation,” *International Economic Review* 45 (2004), 601-641.
- [13] Kimball, M., and Shapiro M., “Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?,” working paper, University of Michigan, 2003.

- [14] Ljungqvist, L., and Sargent, T., “Do Taxes Explain European Employment? Indivisible Labor, Human Capital, Lotteries and Savings”, forthcoming, NBER Macroeconomics Annual.
- [15] Lucas, R.E., and L. Rapping, “Real Wages, Employment and Inflation,” *Journal of Political Economy* 77 (1969), 721-754.
- [16] MaCurdy, T., “An Empirical Model of Labor Supply in a Life Cycle Setting,” *Journal of Political Economy* 89 (1981), 1059-1085.
- [17] McDaniel, C. , “Effective Tax Rates for 15 OECD Countries: 1950-2003,” mimeo, 2006.
- [18] Mendoza, E., Razin A., and Tesar L., “Effective Tax Rates in Macroeconomics: Cross-Country Estimates of Tax Rates on Factor Incomes and Consumption,” *Journal of Monetary Economics* 34 (1994), 297-323.
- [19] Mincer, J., “Labor Force Participation of Married Women: A Study of Labor Supply,” in H.G. LJ: Princeton University Press 1962, 63-97.
- [20] Mulligan, C., “Aggregate Implications of Indivisible Labor,” *Advances in Macroeconomics Berkeley Electronic Press* 1 (2001), 1007-1007.
- [21] Ohanian, L., Raffo, A., and Rogerson, R., “Long-Term Changes in Labor Supply and Taxes: Evidence from OECD Countries, 1956-2004,” NBER Working Paper, 2006.

- [22] Pistqaferri, L., “Aniticipated and Unanticipated Wage Changes, Wage Risk and Intertemporal Labor Supply,” *Journal of Labor Economics* 21 (2003), 729-754.
- [23] Pissarides, C., “Unemployment and Hours of Work: The North Atlantic Divide Revisited,” forthcoming in the *International Economic Review*, 2006.
- [24] Prescott, E., “Prosperity and Depression,” *American Economic Review* 92 (2002), 1-15.
- [25] Prescott, E., “Why Do Americans Work So Much More than Europeans?”, *Quarterly Review of the Federal Reserve Bank of Minneapolis*, July 2004, 2-13.
- [26] Prescott, E., “The Transformation of Macroeconomic Policy and Research, 2004 Nobel Prize Address, forthcoming, *Journal of Political Economy*, 2006.
- [27] Prescott, E., Rogerson, R., and Wallenius, J., “Lifetime Aggregate Labor Supply with Endogenous Workweek Length,” mimeo, 2006.
- [28] Rogerson, R., “Structural Transformation and the Deterioration of European Labor Market Outcomes,” mimeo, 2005.