INDIVIDUAL AND AGGREGATE LABOR SUPPLY WITH COORDINATED WORKING TIMES

Richard Rogerson

Individual and Aggregate Labor Supply With Coordinated Working Times
Richard Rogerson
NBER Working Paper No. 16636
December 2010
JEL No. E24,J22

ABSTRACT

I analyze two extensions to the standard model of life cycle labor supply that feature operative choices along both the intensive and extensive margin. The first assumes that individuals face different continuous wage-hours schedules. The second assumes that all work must be coordinated across individuals. These models look similar qualitatively but have very different implications for how aggregate labor supply responds to changes in taxes. In the first model, curvature in the utility from leisure function plays relatively little role in determining the overall change in hours worked, whereas in the second model it is of first order importance. The second model has important implications for what data is best able to provide evidence on the extent of curvature in the utility from leisure function.

Richard Rogerson
Department of Economics
W.P. Carey School of Business
Arizona State University
Tempe, AZ 85287
and NBER
richard.rogerson@asu.edu

1. Background and Introduction

Starting with the contribution of Lucas and Rapping (1969) and continuing with the development of modern business cycle theory by Kydland and Prescott (1982), economists have sought to understand aggregate labor market outcomes using a framework in which individual economic agents solve explicit optimization problems and interact through explicitly specified market structures. An early controversy in this literature concerned the relationship between individual and aggregate labor supply elasticities. Specifically, in the early representative agent models that dominated the literature and featured solely an intensive margin of labor supply, an exercise of the sort pioneered by MaCurdy (1981) and Altonji (1986) would uncover the value of the key preference parameter that completely characterized both individual and aggregate labor supply responses. This raised an immediate problem, since the implied elasticities from micro data were much smaller than the implied elasticities from aggregate data. Some concluded that this inconsistency was evidence against the overall approach to aggregate economics emphasized by this research program. Heckman (1984) offered a different assessment in his discussion of Kydland (1984) and Ashenfelter (1984). He suggested that the key underlying issue was that the models being used both in micro and aggregate studies were abstracting from key features of individual labor supply problems, thereby limiting the usefulness of comparisons across these studies. Specifically, he argued that because adjustment at the extensive margin is so prevalent at the individual level, any compelling analysis that seeks to derive aggregate implications from individual choice problems would have to incorporate

an extensive margin into the analysis.[1] Absent such a model, it was not clear what the implications would be for aggregate labor supply, nor what the significance of the MaCurdy style estimates for prime aged males would be.

One interpretation of the comments in Heckman (1984) is that they issued a call to develop models which simultaneously capture the important margins of individual labor supply, allow for a rich structure of heterogeneity, and permit us to solve for aggregate outcomes. This would allow us to connect analysis of both individual and aggregate data in a consistent framework. A simple reality of economic analysis is that the quantitative implications of any particular model typically depend on the various features and parameter values that characterize the individual decision problems. Some features that are "realistic" may turn out to not be quantitatively important in terms of substantive economic implications. One of the key objective of economic research is to sort out the important from the not-so-important features in the context of specific issues of interest. To carry out this type of analysis in a consistent fashion requires exactly the sort of model just described.

Shortly after Heckman's (1984) comments, macroeconomists found a way to tractably introduce an extensive margin of labor supply into their models. Hansen (1985) introduced the indivisible labor assumption of Rogerson (1988) into an otherwise standard aggregate model. While in principle this could have facilitated a greater connection between these models and the micro data on labor supply, the immediate effect was almost the opposite. A key property of these represen-

---

[1]Heckman (1993) also emphasized the importance of heterogeneity and the extensive margin in connecting individual labor supply with aggregate labor supply.

tative household models was that the aggregate labor supply elasticity was large independently of the value of the elasticity parameter estimated by MaCurdy and others. As a result, many macroeconomists viewed the indivisible labor assumption as a justification to not look at the micro data on labor supply, since one prominent component of this literature was focused on a parameter which was no longer relevant. Browning et al (1999) pointed out this disconnect from the micro data. In particular, they suggested that while the equilibrium allocations studied by Hansen (1985) and others represented a significant improvement over the earlier analyses because of the presence of adjustment along the extensive margin, the individual employment histories in these models did not correspond at all to those found in the data.[2]

The last ten years has witnessed important extensions of the basic indivisible labor model, and puts us in a position to better address both micro level observations and implications for aggregates within a single model. Of particular interest is the work by Chang and Kim (2006, 2007) and An et al (2009). These authors study models that feature idiosyncratic shocks, incomplete markets and indivisible labor choice.[3] In the context of this model, one can carry out the same type of individual level estimation exercise that MaCurdy (1981) and others performed, at the same time that one can examine the implications of aggregate

---

[2]Browning et al (1998) claimed that the Hansen (1985) model implied equal employment probabilities next period for both the current employed and non-employed. In fact, the structure of the transitions is actually indeterminate. While this implies that the data are not inconsistent with the model, the model certainly does not help us understand why we see particular patterns in terms of movements between employment and non-employment.

[3]These models follow the important earlier contributions of Huggett (1993), Aiyagari (1994) and Krusell and Smith (1998).

shocks for aggregate hours of work in the spirit of Kydland and Prescott (1982).[4] In the context of these models, the old micro data estimates do not turn out to be of great importance in informing us about the size of aggregate fluctuations, while at the same time aggregate labor supply is somewhat less elastic than in the earlier representative agent models. More recently, Krusell et al (2010, 2011) add trading frictions to this framework and show that it can also account for the flows of individual workers between the employment, unemployment and out of the labor force states. French (2005) and Low et al (2010) show in partial equilibrium settings that a model of this sort is also consistent with many features of life cycle labor supply.

In summary, much has changed over the last twenty-five or so years. Viewed from the perspective of the current models that economists are using, the initial controversy about individual and aggregate labor supply elasticities based on importing the estimates of MaCurdy into aggregate representative agent models seems somewhat archaic. We can now explore which model features are important in accounting for various aspects of micro level data and assess the importance of these features for various issues that involve aggregate outcomes. While we now have a solid foundation, there are still many open questions, dealing with such issues as the importance of various types of shocks, the nature of human capital accumulation, the importance of different sources of heterogeneity, the role of

---

[4]An et al (2009) also revisits the contribution of Mankiw et al (1985) and shows that if one tries to interpret the aggregate data in this model as coming from a representative household that only adjusts labor along the intensive margin, one obtains the same types of problematic results that these authors find, i.e., parameter estimates of the wrong sign and violation of concavity.

market structure, trading frictions, family structure, etc....

A key point to take away from this brief overview of developments over the last twenty-five or so years is that choice along the extensive margin figures prominently in current thinking about individual and aggregate labor supply. In many existing models the extensive margin is introduced by assumption–in any given period the individual is assumed to have only two choices–work some pre-specified number of hours, or work zero hours. At a descriptive level, this assumption seems empirically reasonable, corresponding to the observation that there is a great deal of concentration in the distribution of work hours, at the weekly or even annual level. But while the assumption of indivisible labor is empirically descriptive, given its prominence in these models one might well ask what deeper forces lead to this concentration of working hours, and whether the aggregate properties of the model depend on the underlying cause of this concentration. In this paper I take a first look at this issue. In particular, I consider two different extensions of a standard life cycle labor supply model that involve explicit choice along the intensive and extensive margins, each of which represents ideas that have been explored quite a bit in the labor supply literature. In the canonical labor supply model, an individual can work any number of hours, and the wage per unit of time is fixed, leading to a linear budget equation. Researchers have argued that several factors are likely to influence the nature of the constraint set that individuals face when making labor supply decisions, and the two models that I consider incorporate some of these factors.

The first model that I consider assumes that workers face a (continuous) menu

of hours and wage options, with the property that the wage per unit of time is increasing in the volume of work performed. Examples of papers that have documented the presence of non-linear hours-wage menus include Moffitt (1984), Altonji and Paxson (1988), Biddle and Zarkin (1989), Dickens and Lundberg (1993), Keane and Wolpin (2001) and Aaronson and French (2004).[5] My analysis of this model largely summarizes recent work by Prescott et al (2009) and Rogerson and Wallenius (2007, 2009). A key feature of this first model is that workers are still free to choose their hours of work in any given period, although the budget set is no longer linear.

The second model assumes that the work schedule (i.e., intensive margin) is a collective choice in the economy, and that once the work schedule is chosen, the only choice that an individual worker faces is whether to work this fixed number of hours at the going wage rate. This assumption is meant to capture the desire for coordination. The motives for coordination may come from the need for workers within and across firms to work together, or for individuals to coordinate leisure time and/or family schedules. I do not model the underlying reason for coordination, and in a well defined sense focus only the best outcome given the need to coordinate. A key feature of this model is that given a particular collective choice for the work schedule, a worker is not free to work any number of hours, and so it is typically the case that the hours of work for an individual are not consistent with the hours that he or she would choose to work in that period if they faced

---

[5]While these papers assume that the non-linearity shows up in earnings, one could more generally assume that workers face fixed costs outside of the workplace, such as commuting costs. In this regard, the early work of Cogan (1981) on fixed costs is also very relevant.

a linear budget equation given the wage rate. A large literature documents that desired and actual hours of work typically diverge. See, for example, Kahn and Lang (1991), Bell and Freeman (2001) and Sousa-Poza and Henneberg (2003). From the perspective of the individual worker, one can think of this model as one in which the worker faces a discontinuous menu of wage and hours choices, with wages equal to zero for any hours choice other than the collectively decided value. But a key point is that the location of the discontinuity is determined by the collective choice rather than being a feature of technology.

Consistent with the data, both of these extensions to the standard life cycle model generate life cycle profiles for hours at the individual level in which adjustment at the extensive margin plays a key role. I provide analytic and graphical characterizations of optimal lifetime labor supply along the intensive and extensive margins for both models.

I then consider how aggregate labor supply in the two models responds to an increase in the scale of a simple tax and transfer program, and in particular, I contrast them with the outcome that emerges in the standard life cycle model with only an intensive margin of adjustment. The striking finding is that the two models with operative extensive margins for life cycle labor supply generate dramatically different aggregate outcomes. In the standard model in which the only relevant margin is the intensive margin, the aggregate response is tightly connected to the preference parameter that dictates curvature in the utility from leisure. Whether the response is large or small depends critically on this parameter. Consistent with the results of Rogerson and Wallenius (2009), the model with

7

a continuous hours/wage menu generates relatively large responses independently of the preference parameter that dictates curvature in the utility from leisure. In sharp contrast, the results in the work schedule model closely mirror those in the standard model. The important result that follows is that the mere presence of an important role for the extensive margin in terms of life cycle labor supply does not necessarily generate large aggregate elasticities.

While the work schedule model has very different implications for what factors shape the response of aggregate hours to a change in tax and transfer programs, it also has important implications for empirical work that aims to uncover these factors, in particular the curvature parameter in the utility from leisure function. Specifically, this model implies that individual labor supply responds differently to idiosyncratic variation in driving forces than it does to aggregate variation in driving forces.[6] If the work schedule is a collective choice, then it will be invariant to purely idiosyncratic variation in the economic environment, but not to changes in aggregate or common factors. The implication in the stark model studied in this paper is that analysis of micro data may not be sufficient in estimating this key preference parameter. In particular, aggregate data may have a significant role to play in determining the value of this parameter.[7]

An outline of the paper follows. Section 2 describes the three models and characterizes their implications for life cycle labor supply. Section 3 considers the

---

[6]This point is also made by Chetty et al (2011). In their model, the coordination problem only exists within a given firm, but aggregate conditions influence the hours choice made by each firm and create more bunching in hours that would otherwise be observed.

[7]This issue has been recognized in the literature. See for example, Ham (1982), Biddle (1988) and Khan and Lang (1991). See also the discussion in Fehr and Goette (2007) regarding the desirability of experimental data.

implications of each model for changes in the scale of a simple tax and transfer program, both qualitatively and in some simple numerical examples. Section 4 discusses the results of the analysis and Section 5 concludes.

2. Three Models of Life Cycle Labor Supply

In this section I describe three different models of life cycle labor supply. The first model is the canonical model of life cycle labor supply in which the individual can work any number of hours and faces a period budget equation that is linear in hours worked. As we will see below, this model can generate a corner solution for hours worked in which hours worked are zero at some ages. The key property from our perspective is that hours worked are a continuous function of underlying state variables, so there is never a discrete jump from positive hours to zero hours assuming that life cycle productivity varies continuously. The other two models both feature the possibility that hours will jump discontinuously from a positive amount to zero even if underlying productivity is changing continuously. Following Prescott et al (2009), one of these will assume a nonconvex mapping from time devoted to work to labor services. As noted in the introduction, this assumption implies that workers face a menu of wage/hours combinations when making labor supply choices. The other model will constitute the analysis that is novel to this paper, and assumes that an individual is forced to choose a fixed work length for all dates at which labor supply is positive. This is meant to capture the notion that due to coordination issues, all production must be carried out with a fixed working schedule. If this work schedule were exogenously given, this would

9

amount to the standard indivisible labor model of Rogerson and Hansen. But the novel feature here is that the work schedule is chosen by the worker at the beginning of life.

2.1. The Standard Life Cycle Model

Because the extensions that I consider next will have implications for the fraction of lifetime that an individual spends in employment, it is convenient to formulate the model in continuous time, implying that the labor supply choice along the extensive margin is a continuous choice variable and thereby allowing one to characterize optimal choices using standard methods. Consider an individual with length of life normalized to one with preferences defined by:

$$\int_0^1 [u(c(a)) + v(1 - h(a))]da$$

where $c(a)$ is consumption at age a, $h(a)$ is time devoted to market work at age a, $u(\cdot)$ gives the utility flow from consumption and $v(\cdot)$ gives the utility flow from leisure. We assume that these two function are twice continuously differentiable, strictly increasing, strictly concave and satisfy:

$$\lim_{c \to 0} u'(c) = +\infty$$

$$\lim_{h \to 1} v'(1 - h) = +\infty$$

I assume that utility is separable between consumption and leisure. While this has counterfactual implications for the behavior of consumption over the life cycle

10

in the analysis that follows, it serves to simplify the analytic presentation of the results and so is convenient for purposes of exposition. I have also chosen to assume that the individual does not discount future utility flows. As will become clear shortly, this also serves to simplify the analytic characterization of the solution to the individual's maximization problem. I will focus on the case where the interest rate is also zero, so that these two factors will be offsetting as is standard in many aggregate models with infinitely lived agents.

Following much of the life cycle labor supply literature, I assume that the productivity of an individual's time varies systematically over the life cycle. In particular, if an individual of age a devotes h units of time to market production, I assume that this yields $\tilde{e}(a)h$ units of labor services. In the numerical work that follows I will assume that the age profile for this productivity process follows the shape shown in Figure 1.

A few remarks are in order regarding the assumed shape of this productivity profile. First, for reasons of analytic tractability, I am assuming that this productivity process is exogenous, and so in particular, I am abstracting from human capital accumulation decisions that may lie behind this profile. In the data, wages are not symmetric over the life cycle, in the sense that wages at the end of the life cycle are much higher than wages at the beginning of the life cycle. If one takes wages as exogenous and assumes complete markets for borrowing and lending, this can create a problem for a model that includes an endogenous retirement decision. The reason for this is that there is an incentive for individuals to avoid working in the early part of life in order to avoid the low wages during the period,
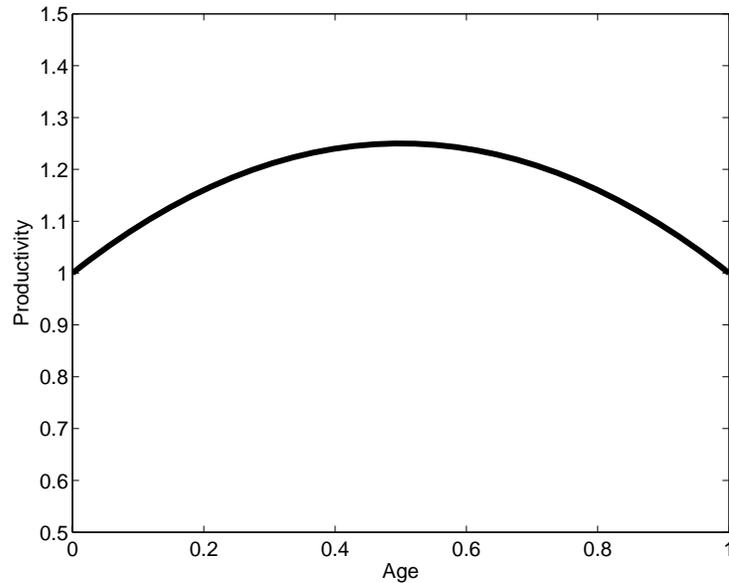
Figure 1: Productivity Over the Life Cycle

and to instead work more at the later part of the life cycle when wages are higher. Wallenius (2009) develops a life cycle labor supply model with an operative extensive margin and human capital accumulation, and shows how it can match the life cycle profile for both wages and hours.[8] To maintain tractability, rather than include a human capital accumulation decision, I choose to abstract from trying to match the actual profile of wages over the life cycle.

Alternatively, one can assume that productivity is constant over the life cycle, but that the utility from leisure varies systematically over the life cycle. In

---

[8]Ljungqvist and Sargent (2010) also develop a tractable model of life cycle labor supply in the presence of human capital accumulation assuming an exogenously given restriction of indivisible labor.
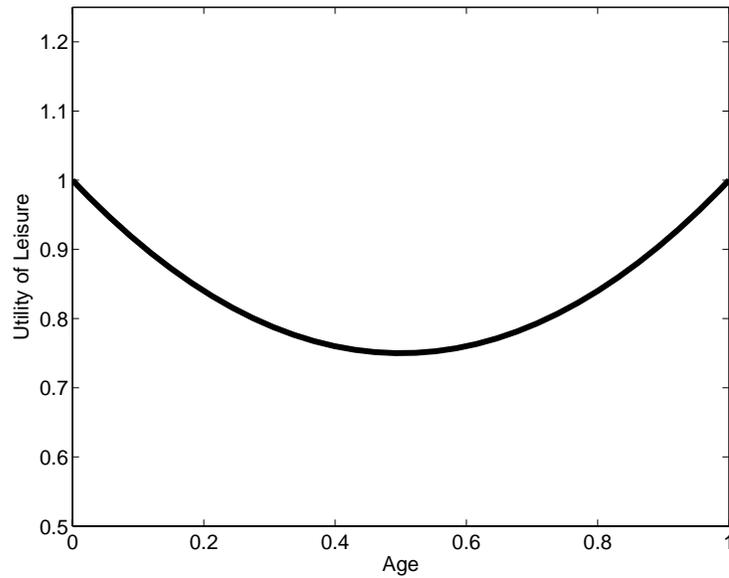
12

Figure 2: Value of Leisure Over the Life Cycle

particular, assume that preferences are given by:

$$\int_0^1 [\log(c(a) + \tilde{\alpha}(a)v(1 - h(a))]da$$

where $\tilde{\alpha}(e)$ has the shape given by Figure 2 below:

The key issue from a modelling perspective is to have something in the model that gives rise to a systematic change in the static net return to working over the life cycle. In what follows I will only consider the specification shown in Figure 1.

The wage rate per unit of labor services is assumed to be constant and equal to w. The individual faces complete markets for borrowing and lending, so assuming as noted above that the interest rate on borrowing and lending is equal to zero,

the present value budget equation is given by:

$$\int_0^1 c(a)da = w \int_0^1 h(a)\tilde{e}(a)da$$

To this point I have only described a single agent decision problem. In the subsequent analysis I will view this single agent problem in the context of a steady state equilibrium in an overlapping generations model. At the risk of trivializing the general equilibrium considerations, but with the gain of transparency, I will assume that we are considering a small open economy in which the real interest rate is exogenously fixed at zero, and that there is an aggregate production function that is linear in labor services with marginal product equal to A.[9] It follows that if the price of output is normalized to one, the equilibrium wage rate w must be equal to A. Assuming a new generation of identical individuals with total mass equal to one is born at each instant, in steady state a new-born household will solve the decision problem depicted above.

Characterizing the solution to the individual's maximization problem is standard. The individual solves:

$$\max_{c(a),h(a)} \int_0^1 [u(c(a)) + v(1 - h(a))]da$$

$$s.t. \int_0^1 c(a)da = w \int_0^1 h(a)\tilde{e}(a)da$$

---

[9]The small open economy assumption is not essential. In this model one can always specify a government debt policy that will support a steady state equilibrium with a zero interest rate.

$$c(a) \geq 0, 0 \leq h(a) \leq 1$$

The solution to this problem will entail a constant flow of consumption, and hence the problem can be rewritten as:

$$\max_{c,h(a)} u(c) + \int_0^1 v(1 - h(a))da$$

$$s.t. c = w \int_0^1 h(a)\tilde{e}(a)da, 0 \leq h(a) \leq 1$$

Substituting the budget equation into the objective function, we obtain the following condition for an interior solution for $h(a)$:

$$\frac{\tilde{e}(a)}{\int_0^1 h(a)\tilde{e}(a)da} = v'(1 - h(a)) \tag{2.1}$$

which can also be written as:

$$v'(1 - h(a)) = \mu w \tilde{e}(a) \tag{2.2}$$

where μ is the marginal utility of consumption.

There are two different forms that the solution may take. One possibility is that the entire profile for $h(a)$ is positive (other than possibly for a set of measure zero, such as the two endpoints), as shown in Figure 3.

The other possibility is that the solution has zero hours of work for an interval at the beginning and end of life, as shown in Figure 4.[10]

---

[10]In their analysis of female labor supply Heckman and MaCurdy (1980) generate corner
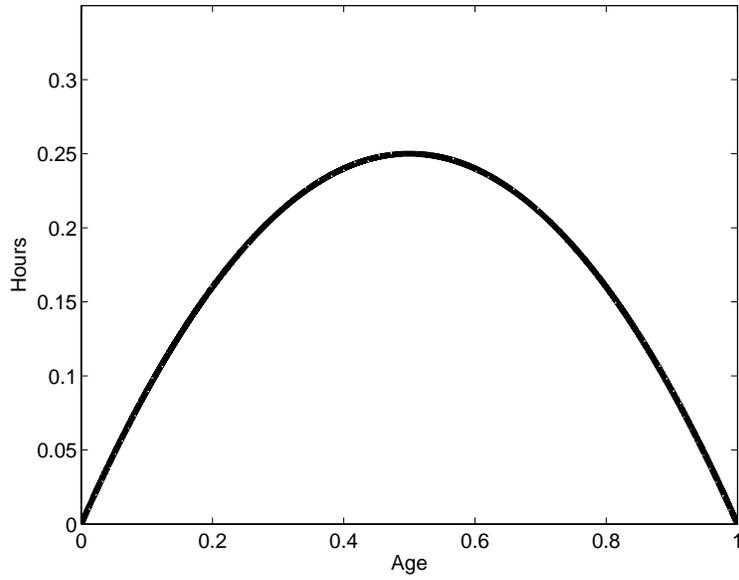
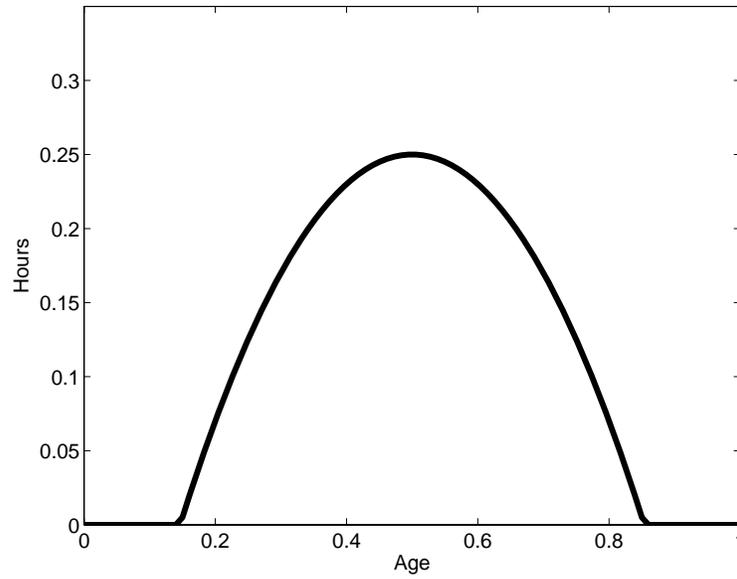Figure 3: Interior Solution for Life Cycle Hours

Figure 4: Corner Solution for Life Cycle Hours

16
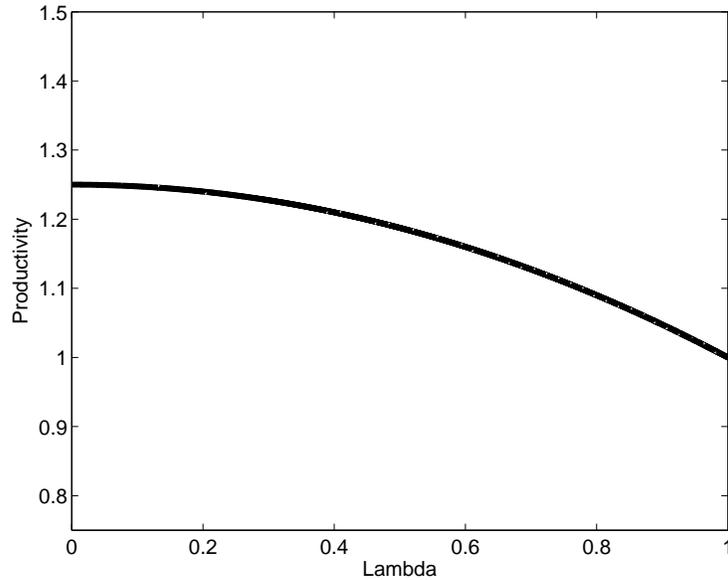
Figure 5: Transformed Productivity Profile

For reasons that will become clear subsequently, it is convenient to transform the maximization problem via a simple change of variables. In particular, instead of examining the optimal labor supply decision as a function of age, it will be convenient to reorder time so as to create a monotone profile for productivity and to instead focus on the mapping from productivity to hours without reference to chronological time. I will use λ as the new index, and will denote the new (monotone decreasing profile) for productivity as $e(\lambda)$. Given $\tilde{e}(a)$ as in Figure 1, the corresponding figure for $e(\lambda)$ is shown in Figure 5.

Recasting the two different solutions denoted above as functions of λ we get Figures 6 and 7.

---

solutions in this manner, though they do not restrict attention to wage profiles of the shape that I consider.

17

Figure 6: Interior Hours Solution With Transformed Productivity

Figure 7: Corner Hours Solution with Transformed Productivity

18

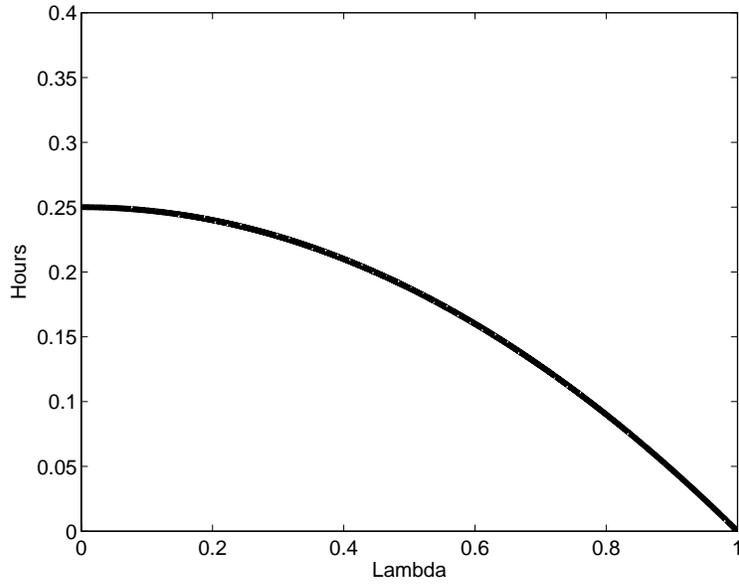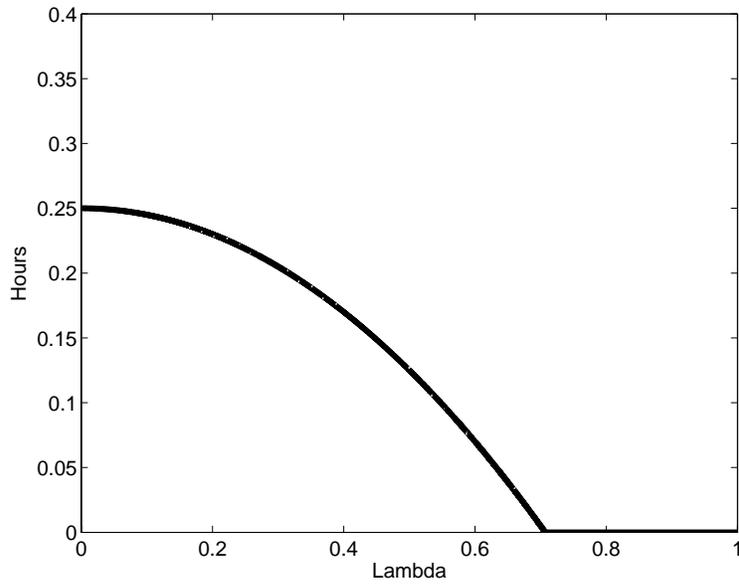For future reference it is of interest to note that if we were to assume that $e(\lambda)$ were constant, the solution would be that h is constant.

2.2. A Nonconvex Mapping from Hours to Labor Services

Following Prescott et al (2009) and Rogerson and Wallenius (2009), this specification modifies the worker's problem by adding a feature to the mapping between time devoted to work and labor services. In particular, we assume that this mapping features a nonconvexity. For simplicity, I focus on the special case in which when a worker of age a devotes h units of time to market work, the resulting supply of labor services is given by

$$\max\{h - \bar{h}, 0\}\tilde{e}(a)h$$

While one interpretation of $\bar{h}$ is that it represents commuting costs, Prescott et al interpret it to reflect set-up costs on the job. From the perspective of equilibrium allocations it does not matter which interpretation is adopted. But the interpretation does matter for observed wages. Specifically, if $\bar{h}$ reflects set-up costs at work, then the model delivers tied hours-wage bundles in equilibrium, whereas with the pure commuting cost interpretation the wage per unit of time will be independent of how many hours are worked.

Taking into account that the worker will choose a constant profile for consumption, the worker's maximization problem can be written as:

$$\max_{c,h(\lambda)} u(c) + \int_0^1 v(1 - h(\lambda))d\lambda$$

19

$$s.t. c = w \int_0^1 \max(h(\lambda) - \bar{h}, 0) e(\lambda) d\lambda, 0 \leq h(\lambda) \leq 1$$

Note that the budget equation implicitly takes into account the fact that although compensation per unit of labor services is constant and equal to w, compensation per unit of time is non-linear in the number of hours devoted to market work. The importance of this feature is that it creates a force for concentration of working time as opposed to smoothing of working time. The analytics of this case are contained in the somewhat more general analysis of Rogerson and Wallenius (2007), but I sketch the details here for completeness. Once again it is convenient to work in λ space instead of age space. First note that if $\bar{h}$ is sufficiently large, the key qualitative result is that instead of having a continuous solution for $h(\lambda)$, we get a solution that drops discontinuously to zero at some point, implying a solution for $h(\lambda)$ as depicted in Figure 8.[11]

If λ is such that the solution for $h(\lambda)$ is interior, one gets the following first order condition:

$$we(\lambda)u'(c) = v'(1 - h(\lambda)) \tag{2.3}$$

We know that $h(0)$ will be positive since this corresponds to the highest productivity for the individual, and the marginal utility of consumption at zero consumption is infinite. Given a value for $h(0)$, it follows from equation (2.3) that for any other

---

[11]See Rogerson and Wallenius (2010) for an analysis of how large $\bar{h}$ must be in order to generate an interior solution along the extensive margin.
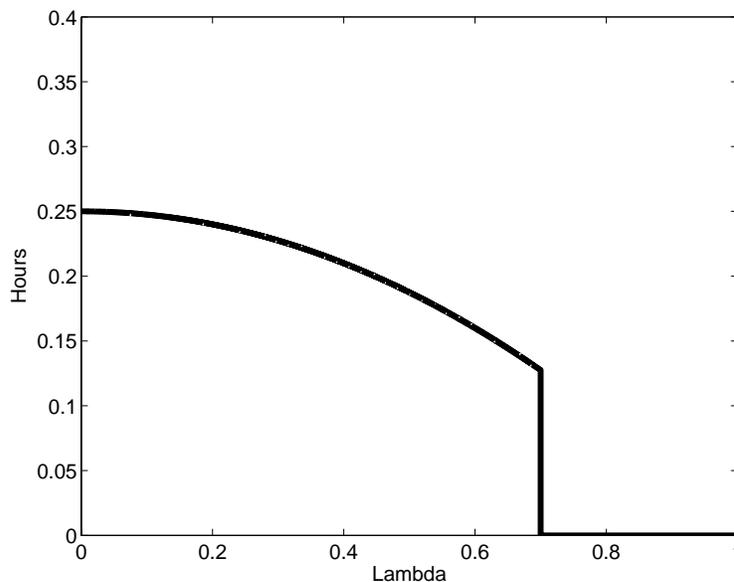
Figure 8: Life Cycle Hours with Fixed Time Costs

value of λ with $h(\lambda) > 0$, the value of $h(\lambda)$ must satisfy:

$$v'(1 - h(\lambda)) = \frac{e(\lambda)}{e(0)} v'(1 - h(0)) \tag{2.4}$$

Given a value for $h(0)$ one can solve for the entire profile for $\tilde{h}(\lambda; h(0))$ that solves equation (2.4). Since $e(\lambda)$ is decreasing, it follows trivially that $\tilde{h}(\lambda, h(0))$ is also decreasing in λ. If the optimal $h(\lambda)$ profile were interior at all values, we would be done at this point except for the determination of $h(0)$. If it is not interior for all λ, the fact that productivity is decreasing in λ, implies that a simple reservation property holds:

$$h(\lambda) = 0 \text{ for all } \lambda \geq \lambda^*$$

21

Given a value for $h(0)$ and the solution for $\tilde{h}(\lambda, h(0))$ implied by equation (2.4), the optimal value of λ^{*} is found by solving:

$$\max_{\lambda^*} u(c) + \int_0^{\lambda^*} v(1 - \tilde{h}(\lambda, h(0)))d\lambda + (1 - \lambda^*)v(1)$$

$$s.t. c = w \int_0^{\lambda^*} \max(\tilde{h}(\lambda, h(0)) - \bar{h}, 0)e(\lambda)d\lambda$$

Assuming an interior solution, the first order condition for λ^* is:

$$w(\tilde{h}(\lambda^*, h(0)) - \bar{h})e(\lambda^*)u'(c) = v(1) - v(1 - \tilde{h}(\lambda^*, h(0))) \qquad (2.5)$$

Combining this with equation (2.3) evaluated at $\lambda = 0$ gives:

$$\frac{v(1) - v(1 - \tilde{h}(\lambda^*, h(0)))}{(\tilde{h}(\lambda^*, h(0)) - \bar{h})e(\lambda^*)} = \frac{v'(1 - h(0))}{e(0)} \qquad (2.6)$$

As shown in Rogerson and Wallenius (2007), equation (2.6) represents an upward sloping relationship between $h(0)$ and λ^* in the neighborhood of the optimal solution. The economic intuition behind this relationship is that an optimal time allocation must have the property that disutility per unit of income should be equated along all margins. So just as equation (2.4) implies that an increase in $h(0)$ implies an increase in $h(\lambda)$ for all λ, it is also true that an increase in $h(0)$ implies that the individual should work deeper into the productivity distribution.

Given $h(0)$ and the solution for $\tilde{h}(\lambda, h(0))$ from equation (2.4), we can represent

c as a function of $h(0)$ and λ^*:

$$c = w \int_0^{\lambda^*} \max(\tilde{h}(\lambda, h(0)) - \bar{h}, 0)e(\lambda)d\lambda \qquad (2.7)$$

Since c is increasing in λ^* it follows that equation (2.7) represents a downward sloping relation between $h(0)$ and λ^*. This downward sloping relation represents the simple fact that in terms of generating consumption, the intensive and extensive choices are substitutes. That is, if an individual is working longer hours, the value of consumption at the margin from working more along the extensive margin is lower.

Assuming that the solution for λ^* is interior, the intersection of these two curves is the solution to the life cycle labor supply problem for this individual. For future reference, I note that if the productivity profile were constant, then the optimal solution would be that hours are constant and then drop discontinuously to zero at some point. This highlights the sense in which this model has a force that opposes the desire of the individual to have smooth hours of work.

2.3. Coordinated Working Times

In this section I add a different constraint to the standard life-cycle problem. In particular, in the spirit of the need to coordinate working schedules, I assume that the worker must choose a fixed work schedule that applies to all periods in which labor supply is positive. Once the work schedule is fixed, the worker faces a simple choice between working and not working at the fixed schedule. In a setting in which workers are heterogeneous, there is a nontrivial issue associated

23

with how to determine the standard work schedule, since different workers may prefer different values. Also, in a changing environment there is an issue about how the work schedule may be altered. I am purposefully abstracting from these potentially important issues in order to focus on some basic implications of this feature that are present even in the absence of these other issues. A more extensive discussion of these simplifications is postponed until later. The worker's problem can now be written as:

$$\max_{c(\lambda),h,I(\lambda)} \int_0^1 [u(c(\lambda)) + v(1 - hI(\lambda))]d\lambda$$

$$s.t. \int_0^1 c(\lambda)d\lambda = wh \int_0^1 I(\lambda)e(\lambda)da$$

$$c(\lambda) \geq 0, 0 \leq h \leq 1, I(\lambda) \in \{0,1\}$$

where in this problem h is the work schedule choice, and then $I(\lambda) \in \{0,1\}$ is an indicator function that represents the choice of whether to work given the fixed work schedule. As before, consumption will be constant over time, and as in the previous subsection, the employment decision will be characterized by a reservation rule: work when $\lambda \leq \lambda^*$. Of course, it is possible that the value of λ^* is equal to one, implying that the individual works in all periods.

Taking this information into account we can write the maximization problem as:

$$\max_{h,\lambda^*} u(wh \int_0^{\lambda^*} e(\lambda)d\lambda) + \lambda^* v(1 - h) + (1 - \lambda^*)v(1)$$

$$s.t. 0 \leq \lambda^* \leq 1, 0 \leq h \leq 1$$

where h is the choice of work schedule that will hold throughout the individual's lifetime, and λ^* is the fraction of life spent in employment, which necessarily consists of the fraction λ^* that has the highest productivity. It is useful to introduce the function $G(\lambda^*)$ defined by:

$$G(\lambda^*) = \int_0^{\lambda^*} e(\lambda)d\lambda \tag{2.8}$$

Assuming interior solutions for both λ^* and h we obtain the following first order conditions:

$$\frac{1}{h} = \lambda^* v'(1-h) \tag{2.9}$$

$$e(\lambda^*)u'(G(\lambda^*)) = v(1) - v(1-h) \tag{2.10}$$

The first equation depicts a downward sloping relationship between λ^* and h. The second equation depicts an upward sloping relationship between λ^* and h. The intuition behind these two relations is identical to that offered in the previous subsection. It follows that one can depict the solution to this problem as the intersection of two curves, one of which is downward sloping and the other of which is upward sloping.

2.4. Comparisons

To contrast the three different solutions we begin by considering the extreme case in which the productivity profile is flat. In the benchmark case this necessarily leads to a flat profile for hours over the entire interval $[0,1]$. In the case of coordinated working times the solution will be the same as in the benchmark model,

since the coordinated working time problem is simply the benchmark problem with some additional constraints. But since the optimal solution satisfies the additional constraints, it follows that this solution must also be optimal in the presence of the additional constraint. In the case of the nonconvexity in the provision of labor services there are two separate cases to consider, depending on whether the non-convexity is sufficiently large to create an interior solution for the fraction of life spent in employment. If this solution is not interior, then hours will necessarily be constant. In this case the nonconvexity acts just like a reduction in the wage rate. Depending upon the function $u(\cdot)$, this may shift the hours profile up or down relative to the other cases, but it will necessarily be constant. In the second scenario, the result will be constant hours when working, but the individual will only work for a fraction of his or her lifetime.

There are some interesting differences to note here. Although the fixed work schedule model looks in general like an indivisible labor model, in the absence of any life-cycle heterogeneity it is identical to the model with perfectly divisible labor. We will return to this later. In contrast, the model with the nonconvexity in the supply of labor services can look like an indivisible labor model even in the absence of life-cycle heterogeneity.

Now consider the models when the $e(\lambda)$ profile does change over time, i.e., in the presence of life-cycle heterogeneity. Assume also that the heterogeneity is sufficient to generate interior solutions for employment in all three cases. Figure 9 shows hours profiles for the three cases, imposing that lifetime labor supply is the same in all cases and that hours are positive for 70% of the individual's life.[12]

---

[12]Note that for given preference and productivity parameters one would not expect the three

26
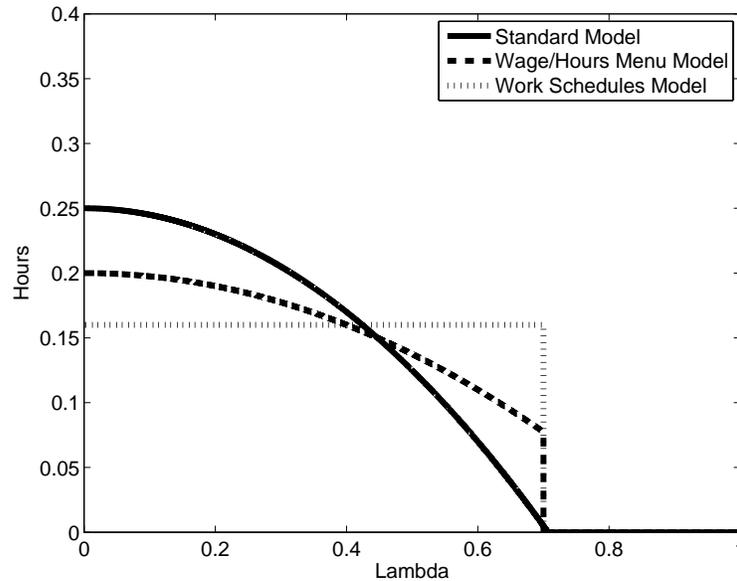
Figure 9: Life Cycle Hours in Three Models

While all three models produce outcomes in which the individual works for only a fraction of his or her life, a key distinction is that in the benchmark model the hours profile drops to zero continuously, while in the other two models the hours profile drops to zero discontinuously. This discontinuity in the life cycle profile for hours worked is a key distinguishing feature of the second and third models relative to the benchmark model. Another interesting distinction has to do with the profile for hours worked while working. In the work schedules model, the hours profile is flat for the region in which hours are positive, but in the other two models there is a positive relationship between hours and productivity in the region with positive hours. In this regard the first two models are similar, while the third is different. In terms of life cycle variation, the work schedule model

_____

solutions to have these properties.

looks like a pure indivisible labor model.

3. Tax and Transfer Programs

In this section I contrast the implications of the three models for how a simple tax and transfer program affects individual life cycle and aggregate labor supply outcomes. In particular, I will consider a policy that levies a proportional tax τ on labor income and uses the proceeds to fund a lump-sum transfer to all individuals currently alive, subject to a balanced budget constraint. This simple policy has been studied extensively in the literature on cross-country differences in hours of work and serves as a useful benchmark for contrasting the implications of these three models.[13]

3.1. Analytic Results

It is easy to derive the implications of such a tax policy on the optimal labor supply choices of individuals. Given the specification of the model we can interpret the changes in the individual choices as reflecting changes in the steady state equilibrium. In the case of the benchmark model, the first order condition that determines the optimal choice of $h(0)$ becomes:

$$(1-\tau)e(0)wu'((1-\tau)w \int_0^1 h(\lambda)e(\lambda)d\lambda + T) = v'(1-h(0)) \qquad (3.1)$$

---

[13]This type of policy has been studied by Prescott (2004), Rogerson (2008), Ohanian et al (2008), Rogerson and Wallenius (2009), and McDaniel (2009). See Rogerson (2006) for a presentation of differences in hours of work across countries.

The condition that relates $h(\lambda)$ to $h(0)$ is unchanged, since the tax rate affects these choices in the same fashion:

$$\frac{v'(1 - h(\lambda))}{v'(1 - h(0))} = \frac{e(\lambda)}{e(0)} \tag{3.2}$$

The government budget constraint implies that:

$$\tau w \int_0^1 h(\lambda)e(\lambda)d\lambda = T \tag{3.3}$$

Combining the government budget constraint into the first order condition for $h(0)$ gives:

$$(1 - \tau)e(0)wu'(w \int_0^1 h(\lambda)e(\lambda)d\lambda) = v'(1 - h(0)) \tag{3.4}$$

Because the solution for $h(\lambda)$ as a function of $h(0)$ is unchanged, it follows that the solution for $h(0)$ is decreasing in τ.

Similar calculations can be done for the other two models. In both cases, taxes do not distort the conditions that relate optimal choices along the intensive and extensive margins. The reason is that both of these margins are distorted by taxes in the same fashion, so that the distortions cancel. In terms of a diagrammatic exposition, this implies that the upward sloping curve that relates optimal choices of intensive and extensive margins does not shift in response to a change in the tax and transfer system. But the tax and transfer scheme does distort the condition that relates total amount of time spent working to the marginal utility of consumption. This leads to a downward shift in the downward sloping relation. It follows that in both cases an increase in the scale of the tax

29

and transfer system leads to a decrease in hours worked along both the intensive and extensive margins. It is important to note that although the two models have similar qualitative predictions, this does not imply that the quantitative predictions are similar. The reason for this is that the upward and downward sloping relationships in the diagrammatic expositions are different across the two models.

It is also instructive to contrast the implications of the three models for the case in which there is no life cycle heterogeneity. In this case, it is easy to show that the benchmark model and the work schedule model will imply that all adjustment takes place along the intensive margin, whereas the nonconvex labor services model implies that all of the adjustment takes place along the extensive margin.

3.2. Numerical Examples

In this subsection I solve some numerical examples to further explore the responses studied analytically in the previous subsection. In order to do this one needs to choose functional forms and parameter values. I choose the standard form for the choice of the utility from leisure function:

$$v(1 - h) = \frac{\alpha}{1 - \frac{1}{\gamma}}(1 - h)^{1 - \frac{1}{\gamma}}$$

I also assume that $u(c)$ is given by $\log c$. This corresponds to the standard assumption in most aggregate analyses, that preferences are consistent with balanced growth. The appropriateness of this assumption in terms of matching individual level trends in hours worked is somewhat of an open question. While this as-

30

sumption is important for the level of the effects that I report, it is not likely to be of first order importance in terms of the relative findings across the three specifications.

For the life cycle productivity profile I choose the quadratic specification:

$$e(\lambda) = 1 - p_0 \lambda^2$$

Note that with log utility over consumption, the labor supply profile is invariant to proportional shifts in the productivity profile, so that there is no loss in generality in assuming $e(0) = 1$.

My main objective here is to examine the relationship between the value of γ and the resulting effects of a change in taxes. As a result I will consider several different values of γ. Because the analysis of the work schedules model is the novel contribution of this paper, I will focus on this model when choosing parameter values. As noted earlier, the solution to this model will involve two values: the fraction of time devoted to work when working, and the fraction of life spent in employment. For each value of γ considered, I choose values of the two parameters α and p_0 so that the solution to this model gives $h = .45$ and $\lambda^* = .74$, implying that aggregate hours worked, denoted by H, satisfy $H = 1/3$. When choosing these values I assume that $\tau = .30$. This procedure would also fully parameterize the benchmark model, since it has no additional parameters. However, I recalibrate α for this model so as to ensure that aggregate hours of work also satisfies $H = 1/3$. The non-convex model has one additional parameter $\bar{h}$. For this model I fix p_0 at the level corresponding to $\gamma = 1.00$ in the work schedule model and then

choose $\bar{h}$ and α so as to achieve the values $\lambda^* = .74$ and $H = \int h(\lambda)d\lambda = 1/3$. Table 1 shows the calibrated parameter values for the three models.

Table 1

Calibration of the Three Models

	$\gamma = 0.10$	$\gamma = 0.25$	$\gamma = 0.50$	$\gamma = 1.00$
		Benchmark Model		
α	.0057	.2065	.6632	1.1814
p_0	1.66	1.35	1.00	.65
		Work Schedule Model		
α	.0053	.1466	.4462	.8210
p_0	1.66	1.35	1.00	.65
		Non-Convex Model		
α	.0297	.4584	.9947	1.5713
p_0	.65	.65	.65	.65
$\bar{h}$	.37	.25	.13	.037

A few patterns emerge in Table 1. The fact that p_0 is decreasing in γ is intuitive. The smaller the value of γ, the more curvature there is in the utility from leisure function, and the greater is the desire to have a smooth profile for leisure. In order for the individual to choose a path for hours with a given change in hours as the worker moves between employment and nonemployment, one needs a greater incentive in terms of productivity differences. That is, in order to encourage the worker to go from $h = .45$ to $h = 0$ at $\lambda = .74$, a smaller value of γ requires a lower

value of productivity at $\lambda = .74$. The required adjustment is very large. One way to gauge this is to compute the implied range of wages during the working part of the life cycle. The ratio of the lowest to the highest wage, i.e., $e(\lambda^*)/e(0)$ is .09, .26, .45 and .64 as γ increases from .10 to 1.00. The corresponding ratio in the data is about .50. One interpretation of this is that it is difficult for this model to reconcile low values of γ with observed hours and wage patterns over the life cycle. A similar issue arises in the non-convex model. The greater the curvature in the utility from leisure function, the greater the nonconvexity must be in order to generate a solution in which the extensive margin is operative in the non-convex model. As one can see from the table, for small values of γ, almost all of a worker's time is devoted to set-up costs. If one were to compute the implied penalty for part-time work, associated with say $h = .20$ versus $h = .40$, then the penalty would be 100% for $\gamma = .10$ or .25, and about 15% when $\gamma = 1.00$. Estimates in the data suggest a penalty no greater than about 20%.[14] Once again, it appears difficult to reconcile low values of γ with observations on hours and wages in this model.[15]

I now consider the consequences of increasing τ to .50. This difference in tax and tax transfer programs approximates the scale of differences that one finds in average effective labor tax rates between the US and several countries in continental Europe, including Belgium, France, Germany and Italy.[16] Table 2 presents the results for the effects of increases in the scale of the tax and transfer program

---

[14]See, for example, Moffitt (1984), Keane and Wolpin (2001), and Aaronson and French (2004).

[15]Rogerson and Wallenius (2010) pursue this issue in greater detail.

[16]See Mendoza et al (1994) and McDaniel (2006) for details concerning cross-country differences in effective tax rates on labor.

in the benchmark model. In this table H and E represent aggregate hours and the aggregate employment rate, respectively, while $h(0)$ refers to hours worked at peak productivity, and asterisks refer to these values in the $\tau = .30$ equilibrium.

Table 2

Tax Effects in the Benchmark Model

	H/H^*	E/E^*	$h(0)/h^*(0)$	Static H/H^*
$\gamma = 0.10$	.96	1.00	.96	.94
$\gamma = 0.25$	.90	.99	.93	.90
$\gamma = 0.50$	.85	.96	.89	.87
$\gamma = 1.00$	.79	.93	.87	.81

The first column in this table shows the value of steady state aggregate hours when $\tau = .50$ relative to the case with $\tau = .30$. Equivalently, this is the lifetime labor supply for an individual when $\tau = .50$ relative to the case where $\tau = .30$. The next two columns show the breakdown of this change in aggregate hours into adjustment along the extensive and intensive margins. The last column shows the implication for aggregate hours in a model that assumed a constant value of productivity over the life cycle. With log utility, hours worked are independent of the level of productivity, so one does not need to specify the actual level of productivity as long as it is constant. In this case the employment rate is always equal to one and there is no adjustment along the extensive margin. Since the individual makes the same choice at each instant of time in this case, it turns out that this specification is equivalent to what one would obtain in a static labor supply problem. The results in this table should not come as a surprise to anyone

familiar with this type of exercise. There is a strong positive relationship between the effect of a tax/transfer on steady state hours and the value of γ. When $\gamma = 1.00$ the effect of an increase in the labor tax rate of twenty percentage points leads to an decrease in hours of work of roughly 20 percent. In contrast, when $\gamma = .10$, the effect is only about one-fifth as large. The last column shows that the effect of life cycle heterogeneity on this result is relatively small. One point of interest in this table is that even in the context of this model, one can obtain a response along both the intensive and extensive margins. As the second column shows, when γ is large, there is a significant change along the extensive margin. But when γ is small, this response disappears.

Next we consider the case of the work schedule model. Results are shown in Table 3.

Table 3

Tax Effects in the Work Schedule Model

	H/H^*	E/E^*	$h(0)/h^*(0)$
$\gamma = 0.10$	.96	1.00	.96
$\gamma = 0.25$	.90	.98	.92
$\gamma = 0.50$	.85	.95	.89
$\gamma = 1.00$	.79	.92	.86

For notational consistency I continue to label the adjustment along the intensive margin by $h(0)/h^*(0)$, though it should be noted that the hours profile is flat in this model. The key result that emerges from this table is that both the aggregate response and the breakdown of this response into intensive and

extensive margins is virtually identical to that in the benchmark model. We postpone further discussion until we have presented results for the model with the nonconvexity in the supply of labor services. Table 4 reports these results.

Table 4

Tax Effects in the Non-Convex Labor Services Model

	H/H^*	E/E^*	$h(0)/h^*(0)$
$\gamma = 0.10$	.84	.85	.98
$\gamma = 0.25$	.83	.85	.96
$\gamma = 0.50$	.81	.86	.93
$\gamma = 1.00$	.79	.88	.89

Because hours are not constant over the life cycle in this model, I report the change along the intensive margin as the change in hours worked at peak productivity. As noted in the analytic derivations, the hours profile shifts down proportionately, so this same drop is observed at all other points at which the individual has positive hours. The results here effectively mirror those presented in Rogerson and Wallenius (2009). First, and most important, the effect of γ on the aggregate response is relatively small.[17] For values of γ that are below .50, to a first approximation the effect of an increase in τ from .30 to .50 is for aggregate hours to decrease by about 17%. There is however, a large impact of γ on the response of hours worked at peak productivity. As detailed in the third column,

---

[17]The effect of γ on the aggregate response in hours is somewhat larger in this table than in the corresponding results in Rogerson and Wallenius (2009). The main reason for this is they recalibrated the slope of the productivity profile as γ changed in order to keep the relative wage ratio constant over the lifecycle, taking into account the effect of the nonconvexity on compensation per unit of time.

when $\gamma = 1.00$ the decrease is 11%, while when $\gamma = .10$ the response is only 2%. This factor five difference is similar to what was found in the two previous models, but what is different in this model is that at the same time that decreasing γ leads to a decrease in the response along the intensive margin, it leads to an increase in the response along the extensive margin.

4. Discussion

The main result that I want to focus on is the very dramatic difference between the work schedule model and the nonconvex labor services model. First and foremost, the work schedule model implies that the aggregate response is very sensitive to the value of the preference parameter γ, while the nonconvex labor services model implies much less sensitivity. For the specifications considered, the two models generate similar responses in aggregate hours when $\gamma = 1$, but when $\gamma = .10$ the response in the nonconvex labor services model is roughly four times as large. However, what is striking about these two different cases is that if one looks only at the life cycle profiles for the two cases, then one finds in both cases that in terms of life cycle labor supply almost all of the variation comes from the extensive margin. In particular, consider the previous results for the case in which $\gamma = .10$. In the case of the work schedules model, by assumption there is no change along the intensive margin for an individual who remains in employment over the life cycle. In the case of the nonconvex labor services model, the range of work along the intensive margin for an individual over their life cycle is from .433 to .457. So both of these models are consistent with the observation in the

37

US data that annual hours of work vary relatively little for full time employed males over the life cycle, with most of the variation in total hours for this group due to changes from either full time to nonemployment, or full time to part time. (See, e.g., Prescott et al (2009) for evidence on this point.) However, despite this fact, the two models have dramatically different implications for the magnitude of the response in aggregate hours to a change in the scale of the tax and transfer program.

Loosely speaking, one might view these two models as two different models that give rise to something that looks like a pure indivisible labor model, where by a pure indivisible labor model I have in mind a model that takes the intensive margin as fixed exogenously. Such a model has two striking implications. First, the aggregate labor supply response to changes in the scale of tax and transfer programs is large, and second, the aggregate response in hours worked is independent of the curvature parameter on the utility from leisure function. One might conjecture that any underlying model that generates something that looks like the pure indivisible labor model in steady state might have these same implications. The nonconvex labor services model analyzed by Prescott et al (2009) and Rogerson and Wallenius (2009) is one example for which this conjecture is true.

However, the very simple work schedule model analyzed in this paper shows that this conjecture is not generally true. That is, the mere fact that the extensive margin is the key margin of adjustment in terms of labor supply over the life cycle need not imply that aggregate responses are large or that the preference parameter dictating curvature in the utility from leisure function is irrelevant. In fact, despite

the fact that the work schedules model implies that all adjustment over the life cycle takes place along the extensive margin, we found that aggregate hours in this model behaved virtually identical to how they behave in a model that has all life cycle adjustment take place along the intensive margin. One simple conclusion from the above discussion is that the implications of indivisible labor for aggregate labor supply depend very much on what the underlying source of the apparent indivisibility is.

This discussion cautions us that just because the extensive margin is dominant in terms of individual labor supply changes over the life cycle, we should not conclude that the curvature parameter in the utility from leisure function is irrelevant for understanding the response of aggregate hours to various changes in the economic environment. But the work schedule model also has important implications regarding what information is needed to obtain estimates of the curvature parameter γ. In particular, one might argue that micro data provides the best opportunity to learn about preference parameters because individuals are subjected to many large idiosyncratic changes in the factors that shape their economic environment. Coupled with the fact that micro data sets give us observations on a relatively large number of individuals, it follows that this source provides us with lots of independent observations on how individuals respond to large changes in their economic environment. But the simple model of work schedules that I have sketched out implies that idiosyncratic variation is irrelevant to the determination of the work schedules.[18] That is, if we were to simply change the tax rate that a

[18]Chetty et al (2011) makes a similar point but more general point. In their model each firm chooses its own work schedule, but takes into account the distribution of desired work hours

39

given individual faces, as opposed to the tax rate that all individuals face, then nothing would happen to the economy wide choice of work schedules and the individual decision problem would look just like a pure indivisible labor problem. Even observing the entire lifetime labor supply response of a given individual to a particular change in their economic situation would not allow us to uncover the value of γ.

To pursue this a bit further, consider an individual who is near the point of peak life cycle productivity, and consider how this individual responds to an unanticipated change in the return to work at this point in time. If the unanticipated change leads to higher returns to work, there is no option for the individual to increase his or her hours of work, so we will necessarily not observe any change in hours of work in response to this event. One might falsely conclude that individuals are not very willing to substitute leisure, either across time or in return for additional consumption. Alternatively, suppose that the effect of the change is to reduce the return to work. If the period being considered is close to peak productivity, then in order for the change to bring about a change in hours worked at that point would require a sufficiently large change in the return to work to lower the return below that of the reservation productivity level $e(\lambda^*)$. It follows that except for very large changes we would again observe no change for this individual. In either case, looking at contemporaneous responses at the individual level to idiosyncratic changes in the return to working would lead us to find effectively

in the workforce. The implication is that tax changes which affect only a small fraction of the population are less likely to lead to changes in work schedules than are changes which affect a large fraction of the population.

40

no response in hours worked. Yet the same change in the return to working, when relevant for all individuals might lead to a large change in aggregate hours worked. The key point is that despite the volume of changes that take place at the micro level, it may be that the responses of individuals to aggregate changes in the economic environment may play a key role in helping us uncover the key parameters of individual preferences. Put somewhat differently, the results in Chang and Kim (2006) show us that given a fixed working length per period, micro data on hours and wages do not provide information on the parameter γ.

At this point it is useful to remark on some of the features of the very simple work schedule model considered here. The work schedule model studied here is really nothing more than an example that is (hopefully) useful in illustrating some basic points. But as a serious model that might be used to provide compelling guidance to either data analysis or the response of hours worked to policy changes it undoubtedly raises some basic questions. First, taking as given that one of the most robust patterns in the micro data is the increasing profiles for wages and annual hours worked over the life cycle, an important limitation of the work schedule model studied here is that it does not account for the increase in annual hours worked over the life cycle. One generalization of the model considered here that could address this issue would be to allow for the possibility that as individuals accumulate experience they perform different roles within an organization, and that some of these roles might require a different number of hours. For example, when an individual gets promoted from being a regular worker to being a supervisor, he or she may have to show up for work earlier and stay later in order to facilitate the

41

opening and closing of the establishment. The key point is that a work schedules model may incorporate the reality that different positions may be associated with different hours, with these differences reflecting considerations from the production side. That is, even if wages and hours move together, the variation might be determined solely by features of production, and so not provide any information about preferences aside from the obvious revealed preference implication, that is, that the individual is willing to work the additional hours associated with the increase in pay and responsibility.

A second issues concerns the economy wide nature of the work schedule. While one might accept the fact that a given establishment must choose a work schedule that coordinates the working hours of its employees, it is somewhat less clear that this needs to be done across establishments. One might then expect to see different establishments with different work schedules, each one reflecting the desired work hours for different subsets of the population. In the simple model that I studied, this might seem a reasonable solution to the problem. But there are three issues to raise. First, it seems plausible that for many establishments, an important attribute is its hours of operation, and that it is important to be available to deal with customers during what are perceived to be "usual" or "normal" business hours. Second, to the extent that establishments care about turnover, it is important to incorporate the potential productivity losses associated with outcomes that would involve individuals moving across establishments whenever there was a change in their desired hours of work. Third, to the extent that there are some frictions in labor markets, establishments may need to take into account

the work preferences of the "average" worker when creating a position.

Having raised these issues, it is of course also true that there are some differences in work schedules across establishments within industries, as well as across occupations or industries. Teachers have annual work schedules that are far different from many other occupations. It is plausible that this plays a role in the decision of some individuals whether to enter a particular occupation. But it is more likely that the selection relevant for this choice has to do with permanent differences in preferences as opposed to life cycle changes in the return to working. There are also many jobs in which coordination of work schedules is viewed as less important. Many restaurants and retail establishments hire a mix of full and part time workers. For certain subsets of the population these opportunities are likely to be quite important and effectively create a situation in which the worker faces a flexible hours choice at a given wage. Finally, even in a world in which many jobs have fixed work schedules, an individual can always augment his or her hours of work by taking on an additional job from a sector that offers opportunities with flexible hours. The simple model that I considered assumed that such possibilities do not exist. Understanding both the extent of the availability of such opportunities and how they influence our inference about labor supply is an important issue for future research.

5. Conclusion

I analyze three different models of life cycle labor supply. The first is the standard model in which hours of work change continuously in response to continuous

changes in underlying primitive, such as productivity or preferences. The other two both feature operative choices along both the intensive and extensive margin, in the sense that hours of work can change discontinuously to zero in the face of continuous changes in underlying primitives, but they differ in terms of the underlying economic reason for the operative choice along the extensive margin. The first assumes that individuals face different wage rates per unit of time depending upon the amount of time devoted to work. The second assumes that all work must be coordinated across individuals, implying that all individuals must work the same amount of time. Qualitatively these two models look similar, in that both can generate life cycle labor supply profiles in which adjustment along the extensive margin is the key margin of adjustment. However, these two models have very different implications for how aggregate labor supply responds to changes in the scale of a simple tax and transfer scheme. In the first model, curvature in the utility from leisure function plays relatively little role in determining the overall change in aggregate hours worked, whereas in the second case this parameter is of first order importance. It follows that an operative extensive margin need not imply that this curvature is irrelevant for aggregate labor supply responses. However, I also argue that the second model has important implications for what data is most able to provide evidence on the extent of curvature in the utility from leisure function. In particular, panel micro data that is dominated by idiosyncratic variation in the key forcing factors may not of much use. Instead, one may want to find forcing variables which feature aggregate variation.

The analysis presented here has focused on some very simple settings in order

to best communicate the messages just repeated above. An important task for future work is to examine the extent to which these messages are robust to settings that are more closely connected to the data. The two models analyzed here represent extensions of the canonical labor supply model in which the constraint set facing the individual is altered. In the standard model, the one-period budget constraint is linear in hours worked. In the model with a wage-hours menu, the one period budget constraint is non-linear. In the work schedules model it is represented by two disjoint points. A basic message of the analysis is that the nature of the constraint set matters both for inferring parameter values of individual utility functions and for the aggregate response to a given change in the overall economic environment. An important issue is to consider more general versions of these features and to assess the quantitative significance of each of them. For example, even if hours at a particular employer are fixed from the perspective of an individual worker, there are other employers that may offer different hours of work. Important issues then involve the extent to which this individual can locate these other opportunities (i.e., are there search frictions) and the extent to which the worker's skills are transferable across employers (i.e., the nature of human capital). Chetty et al (2011) is a recent paper that pushes in this direction. There may also be opportunities for working on a second job if the worker wishes to increase his or her hours of work. One would then need to incorporate the form that these opportunities take. The effect of market imperfections, e.g., credit market imperfections, on constraint sets is also relevant, as is the nature and extent of human capital accumulation. Finally, in the context of models with

uncertainty, there are important issues about the nature of uncertainty and the information sets of individual agents. While some work along these lines has been carried out, there are still many issues to be resolved, specifically in assessing the importance of these various features for the properties of aggregate labor supply.

References

[1] Aaronson, Daniel, and Eric French. (2004) "The Effect of Part-Time Work on Wages: Evidence from the Social Security Rules." *Journal of Labor Economics,* 22:2, 329-352.

[2] Aiyagari, Rao. (1994) "Uninsured Idiosyncratic Risk and Aggregate Savings." *Quarterly Journal of Economics,* 109:3, 659-683.

[3] Altonji, Joseph G. (1986) "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *Journal of Political Economy,* 94:3, S176-S215.

[4] Altonji, Joseph G., and Christina H. Paxson. (1988) "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs." *Journal of Labor Economics,* 6:2, 254-276.

[5] _____. (1992) "Labor Supply, Hours Constraints and Job Mobility." *Journal of Human Resources,* 27:2, 256-278.

[6] An, Sungbae, Yongsung Chang, and Sun-Bin Kim. (2009) "Can a Representative Agent Model Represent a Heterogeneous Agent Economy?" *American Economic Journal Macroeconomics*, 1:2, 29-54.

[7] Ashenfelter, Orley. (1984) "Macroeconomic Analyses and Macroeconomic Analyses of Labor Supply." *Carnegie Rochester Series on Public Policy,* 21, 117-156.

[8] Bell, Linda A., and Richard B. Freeman. (2001) "The Incentive for Working Hard: Explaining Hours Worked Differences in the US and Germany." *Labour Economics,* 8:2, 181-202.

[9] Biddle, Jeff E.. (1988) "Intertemporal Substitution and Hours Restrictions." *Review of Economics and Statistics* ,70:2, 347-351.

[10] Biddle, Jeff E., and Gary A. Zarkin (1989). "Choice Among Wage-Hours Packages: An Empirical Investigation of Male Labor Supply." *Journal of Labor Economics,* 7:4, 415-437.

[11] Browning, Martin, Lars P. Hansen and James J. Heckman. (1999) "Micro Data and General Equilibrium Models." in *Handbook of Macroeconomics,* edited by John B. Taylor and Michael Woodford, 543-633.

[12] Chang, Yongsung, and Sun-Bin Kim. (2006) "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy." *International Economic Review,* 47:1, 1-27.

[13] _____. (2007) "Heterogeneity and Aggregation in the Labor Market: Implications for Aggregate Preference Shifts." *American Economic Review,* 97:5, 1939-1956..

[14] Chetty, Raj, John N. Friedman, Tore Olsen and Luigi Pistaferri. (2011) "Adjustment Costs, Firm Responses and Labor Supply Elasticities: Evidence from Danish Tax Records." Forthcoming in *Quarterly Journal of Economics.*

[15] Cogan, John F. (1981) "Fixed Cost and Labor Supply." *Econometrica,* 49:4, 945-963.

[16] Dickens, William T. and Shelly J. Lundberg. (1993) "Hours Restrictions and Labor Supply." *International Economic Review,* 34:1, 169-192.

[17] Domeij, David, and Martin Floden. (2006) "The Labor Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased." *Review of Economic Dynamics,* 9:2, 242-262.

[18] Fehr, Ernst, and Lorenz Goette. (2007) "Do Workers Work More If Wages Are Higher? Evidence From a Randomized Field Experiment." *American Economic Review,* 97:1, 298-317.

[19] French, Eric. (2005) "The Effect of Health, Wealth and Wages on Labour Supply and Retirement Behavior." *Review of Economic Studies,* 72:2, 395-427.

[20] Ham, John C. (1982) "Estimation of a labor Supply Model with Censoring Due to Unemployment and Underemployment." *Review of Economic Studies,* 49:3, 335-354.

[21] Hansen, Gary D. (1988) "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics,* 16:3, 309-337.

[22] Heckman, James J. (1984) "Comments on the Ashenfelter and Kydland Papers." *Cargnegie-Rochester Conference Series on Public Policy,* 21, 209-224.

[23] _____. (1993) "What Has Been Learned About Labor Supply in the Past Twenty Years?" *American Economic Review,* 83:2, 116-121

[24] Heckman James, and Thomas E. MaCurdy. (1980) "A Life Cycle Model of Female Labour Supply." *Review of Economic Studies,* 47:1, 47-74.

[25] Huggett, Mark. (1993) "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17:5-6, 953-969.

[26] Imai, Susumu, and Michael P. Keane. (2004) "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review,* 45:2, 601-641.

[27] Kahn, Shulamit, and Kevin Lang. (1991) "The Effect of Hours Constraints on Labor Supply Estimates." *Review of Economics and Statistics,* 73:4, 605-611.

[28] Keane, Michael P., and Kenneth I. Wolpin. (2001) "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment." *International Economic Review,* 42:4, 1051-1103.

[29] Krusell, Per, and Anthony A. Smith. (1998) "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy,* 106:5, 867-896.

[30] Krusell, Per, Toshi Mukoyama, Richard Rogerson, and Aysegul Sahin. (2010) "Aggregate Labor Market Outcomes: The Roles of Choice and Chance," *Quantitative Economics*, 1:1, 97-127.

[31] _____. (2011) "A Three State Model of Worker Flows in General Equilibrium." Forthcoming in *Journal of Economic Theory*.

[32] Kydland, Finn. (1984) "Labor Force Heterogeneity and the Business Cycle." *Carnegie Rochester Conference Series on Public Policy,* 21, 173-208.

[33] Kydland, Finn., and Edward C. Prescott. (1982) "Time to Build and Aggregate Fluctuations." *Econometrica,* 50:6, 1345-1370.

[34] Ljungqvist, Lars, and Thomas Sargent. (2010) "Curvature of Earnings Profile and Career Length." mimeo, New York University.

[35] Low, Hamish, Costas Meghir and Luigi Pistaferri. (2010) "Employment Risk and Wage Risk over the Life Cycle." *American Economic Review,* 100:4, 1432-1467.

[36] Lundberg, Shelly. (1985) "Tied Wage-Hours Offers and the Endogeneity of Wages." *Review of Economics and Statistics,* 67:3, 405-410.

[37] Lucas, Robert E., and Leonard Rapping. (1969) "Real Wages, Employment, and Inflation." *Journal of Political Economy,* 77:5, 721-754.

[38] MaCurdy, Thomas E. (1981) "An Empirical Model of Labor Supply in a Life Cycle Setting." *Journal of Political Economy,* 89:6, 1059-1085.

[39] Mankiw, N. Gregory, Julio J. Rotemberg and Lawrence H. Summers. (1985) "Intertemporal Substitution in Macroeconomics." *Quarterly Journal of Economics,* 100:1, 225-251.

[40] McDaniel, Cara. (2006) "Effective Tax Rates for 15 OECD Countries: 1950-2003." mimeo, Arizona State University.

[41] _ _ _ _ _ _ _ _. (2009) "Hours of Work in the OECD 1960-2003: Driving Forces and Propagation Mechanisms." manuscript, Arizona State University.

[42] Mendoza, Enrique G., Asaf Razin and Linda L. Tesar. (1994) "Effective Tax Rates in Macroeconomics: Cross-Country Estimates of Tax Rates on Factor Incomes and Consumption." *Journal of Monetary Economics,* 34:3, 297-323.

[43] Moffitt, Robert. (1984) "The Estimation of a Joint Wage-Hours Labor Supply Model." *Journal of Labor Economics,* 2:4, 550-566.

[44] Ohanian, Lee, Andrea Raffo and Richard Rogerson. (2008) "Long-Term Changes in Labor Supply and Taxes: Evidence from OECD Countries, 1956-2004." *Journal of Monetary Economics,* 55:8, 1353-1362.

[45] Prescott, Edward C. (2004) "Why Do Americans Work So Much More than Europeans?" *Quarterly Review of the Federal Reserve Bank of Minneapolis,* 28:1, 2-13.

[46] Prescott, Edward C., Richard Rogerson and Johanna Wallenius. (2009) "Lifetime Aggregate Labor Supply with Endogenous Workweek Length." *Review of Economic Dynamics,* 12:1, 23-36.

[47] Rogerson, Richard. (1988) "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics,* 21:1, 3-16.

[48] _____. (2006) "Understanding Differences in Hours Worked." *Review of Economic Dynamics,* 9:3, 365-409.

[49] _____. (2008) "Structural Transformation and the Deterioration of European Labor Market Outcomes." *Journal of Political Economy,* 116:2, 235-259.

[50] Rogerson, Richard, and Johanna Wallenius. (2007) "Micro and Macro Elasticities in a Life Cycle Model with Taxes." NBER Working Paper #13017.

[51] _____. (2009) "Micro and Macro Elasticities in a Life Cycle Model with Taxes." *Journal of Economic Theory,* 144:6, 2277-2293.

[52] _____. (2010 "Fixed Costs, Retirement and the Elasticity of Labor Supply." mimeo, Arizona State University.

[53] Sousa-Poza, Alfonzo, and Fred Henneberg. (2003) "Work Attitudes, Work Conditions and Hours Constraints, An Explorative, Cross-National Analysis." *Labour,* 14:3, 351-372.

[54] Wallenius, J. (2009). "Social Security and Cross-Country Differences in Hours: A General Equilibrium Analysis." mimeo, Stockholm School of Economics.