FEDERAL RESERVE BANK OF MINNEAPOLIS QUARTERLY REVIEW

October 2012

Aggregate Labor Supply

Edward C. Prescott
Johanna Wallenius

FEDERAL RESERVE BANK OF MINNEAPOLIS

Quarterly Review Vol. 35, No.2

ISSN 0271-5287

This publication primarily presents economic research aimed at improving policymaking by the Federal Reserve System and other governmental authorities.

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

The *Quarterly Review* is published by the Research Department of the Federal Reserve Bank of Minneapolis. This has become an occasional publication; however, it continues to be known as the *Quarterly Review* for citation purposes. Subscriptions are available free of charge. To subscribe to the journal and be automatically notified whenever a new issue is published, please go to http://www.minneapolisfed.org/publications_papers/subscriptions/research/index.cfm to sign up. Readers will be notified by e-mail immediately upon publication.

Electronic files of past *Quarterly Review* articles are available through the home page of the Federal Reserve Bank of Minneapolis, http://www.minneapolisfed.org.

Comments, questions, and subscription requests about the *Quarterly Review* may be sent to the circulation assistant at mea@res.mpls.frb.fed.us.

Contents

A review of the long-held debate regarding the magnitude of the aggregate labor supply elasticity.

Aggregate Labor Supply

Edward C. Prescott
Senior Monetary Adviser
Research Department
Federal Reserve Bank of Minneapolis,
Department of Economics
Arizona State University,
and Center for the Advanced Study
in Economic Efficiency
Arizona State University

Johanna Wallenius
Department of Economics
Stockholm School of Economics

Fifty years ago, the labor supply decision was thought of as virtually irrelevant for macroeconomic analysis. The view was that in the aggregate, labor supply was not determined by the same factors that determined individual labor supply. Lucas and Rapping (1969) challenged this view. Tremendous advances have been made in macroeconomic analysis following the introduction of labor supply into the field. Key among these advances was endogenizing the labor supply decision in the neoclassical growth model, which allowed its use in studying business cycles. This framework can also be credited with introducing the stand-in aggregate household construct, which has proven to be a highly useful abstraction. Subsequently, the methodology originally used for studying business cycles has been used to advance learning in most areas of macroeconomics.

Today, we understand that labor supply matters for many key economic issues—not only for the effects of business cycle shocks but also, for example, for tax policy analysis. However, the extent to which labor supply matters for such questions depends on the labor supply elasticity. Although the importance of the labor supply elasticity is nowadays widely agreed upon among economists, the magnitude of the elasticity is not. Most labor economists argue that the elasticity is small, a view based on the low variations in hours worked and wages of prime-aged males. Macroeconomists, on the other hand, argue that the elasticity is big, a view based on differences in tax rates and aggregate hours across countries and time, as well as the fact that the neoclassical growth model displays the business cycle facts only if this elasticity is high. This apparent inconsistency is bothersome because it creates disagreement over the importance of labor supply for many important macroeconomic issues. What is needed is a theory consistent with both micro- and macroeconomic observations.

In this article, we demonstrate that such a theory now exists. We discuss the issues related to the apparent inconsistency between the micro- and macroeconomic observations and stress that much of the confusion stems from the notion that one can estimate the labor supply elasticity in one context and export it to another. This notion, we find, is ill advised.

Evidence from Business Cycles and Cross-Country Tax Analysis

The modern theory of economic growth evolved from the observation of striking similarity both over time and across countries. The success of the neoclassical growth model can be attributed to its ability to reproduce the growth facts (see Kaldor 1957). Similarly, economic fluctuations display remarkable empirical regularity, commonly referred to as the *business cycle facts*. These facts are as follows: (1) two-thirds of fluctuations are accounted for by variation in the labor input, whereas one-third of fluctuations are accounted for by variation in total factor productivity; (2) consumption moves procyclically; and (3) percentage-wise, investment is roughly 10 times as volatile as consumption. Regardless of this regularity, for a long time the study of short-term economic behavior—namely, fluctuations—was divorced from the study of long-term growth. The likely reason is that short-term movements in output are in large part accounted for by movements in the labor input, whereas long-term increases in living standards are mainly accounted for by increases in capital service inputs and total factor productivity. The premise of modern business cycle theory, however, is that growth and fluctuations are not distinct phenomena that should be studied with different tools.

Kydland and Prescott (1982) use the neoclassical growth model to study business cycles. The framework introduced an aggregate or stand-in household construct, which has proven to be a most successful abstraction. The underlying aggregation theory is based on maximizing a weighted sum of individual utilities. The framework also endogenizes the labor supply decision. The growth facts state that consumption and investment shares of output are roughly constant and that variables other than labor supply and the return on capital grow over time. This pattern dictates a Cobb-Douglas production function. The growth facts also place restrictions on the utility function. They do not, however, pin down the aggregate labor supply elasticity, which turns out to be a key parameter for deriving the predictions of the growth model for business cycle fluctuations.

Kydland and Prescott (1982) show that the neoclassical growth model extended to allow for stochastic shocks to the rate of productivity growth generates real business cycles. However, the model displays the business cycle facts only if the aggregate labor supply elasticity is sufficiently large, around three. Many macroeconomists view this result as evidence of a highly elastic labor supply.

Prescott (2004) argues that differences in taxes and labor supply provide further macroeconomic support for the notion of a large aggregate labor supply elasticity. Striking differences can be seen in hours of market work both across countries and over time. To illustrate, aggregate hours worked are currently about 70 percent of the U.S. level in the continental European countries of Belgium, France, and Germany. Simultaneously, we observe large differences in marginal tax rates across countries. Prescott (2004) and Ohanian, Raffo, and Rogerson (2008) study the role of taxes in accounting for the differences in aggregate hours across countries and over time.

The premise for these studies is an aggregate household construct. Specifically, assume that the aggregate household has preferences over sequences of consumption (c) and hours worked (h) ordered by

$$\sum_{t=0}^{\infty} \beta^t \left[\log(c_t) - \alpha \frac{h_t^{1+\gamma}}{1+\gamma} \right],$$

where t denotes time, β is the discount factor, and α is the parameter governing the disutility from working. The key parameter is γ because it determines the aggregate intertemporal elasticity of substitution of labor. The per-period time endowment is normalized to one. The household owns the capital stock in the economy and rents it to the firm. The law of motion for the capital stock k is standard and given by

$$k_{t+1} = (1-\delta)k_t + i_t,$$

where δ is depreciation and i is investment. A Cobb-Douglas production function for the aggregate firm is assumed:

$$y_t = A k_t^{\theta} h_t^{1-\theta},$$

where θ is the capital share parameter and A is the total factor productivity parameter. The government imposes proportional taxes on income, the proceeds of which are rebated lump-sum back to the household. The period t budget constraint faced by the household is then

$$(1+\tau_c)c_t + (1+\tau_i)i_t = (1-\tau_h)w_t h_t + (1-\tau_k)(r_t - \delta)k_t + \delta k_t + T_t,$$

where τ_c is the tax on consumption, τ_i the tax on investment, τ_h the marginal tax rate on labor income, τ_k the tax on capital income, w_t the real wage, r_t the rental price of capital, and T_t the lump-sum transfers.

The labor and consumption taxes can be combined into one effective marginal tax rate on labor income. It is given by the fraction of additional labor income that is taken in the form of taxes:

$$\tau = \frac{\tau_h + \tau_c}{1 + \tau_c}.$$

The two key equations are the first-order conditions for the marginal rate of substitution between consumption and hours worked and the profit-maximizing condition that states that individuals are paid their marginal product:

$$\alpha c h^{\gamma} = (1-\tau) w$$
$$w = (1-\theta)\frac{y}{h}.$$

When combined, these equations determine the following equilibrium relation between aggregate labor supply, the consumption-output ratio, and the tax rate at time *t*:

$$h_t = \left(\frac{1-\theta}{\frac{c_t}{y_t} \frac{\alpha}{1-\tau_t}} \right)^{\frac{1}{1+\gamma}}.$$

The c/y term is a function of the distribution of future exogenous variables. The $(1-\tau_t)$ term captures the intratemporal distortion to the relative prices of consumption and leisure.

This equation can be used to predict the impact of taxes on labor supply. The conclusion is that in order for taxes to play an important role in accounting for the cross-country differences in aggregate hours, the labor supply elasticity—namely, $1/\gamma$—must be large.

Estimates of Individual Elasticity from Panel Data

Many labor economists argue that the aggregate labor supply elasticities used in the business cycle and cross-country tax studies are not in accordance with the microeconomic evidence. This disparity has led them to question the validity of the business cycle model and to argue that the effect of taxes on aggregate hours is overstated due to the large labor supply elasticity that is assumed.

For some economies, the labor supply elasticity of the aggregate household and the individual labor supply elasticity of the individuals being aggregated should be the same. This will, for example, be the case if preferences are convex, which means a concave utility function defined on a *convex* subset of the commodity space. For other economies, this will not be the case, and the utility function of the aggregate household will be markedly different from the one of the individuals being aggregated. Indeed, if the aggregate labor supply were *not* significantly higher than the individual labor supply, the micro labor statistics would lead to the rejection of the conclusions derived using the basic neoclassical growth model for business cycle fluctuations.

A microeconomic approach is used to identify individual labor supply elasticity from the variation of wages and hours over the life cycle. A simplified illustration of this approach is as follows. Consider a modified version of the formulation from the previous section, where the individual faces a present value budget equation:

$$\max \sum_{t=0}^{\infty} \beta^t \left[\log(c_t) - \alpha \frac{h_t^{1+\gamma}}{1+\gamma} \right]$$

$$\text{s.t.} \quad \sum_{t=0}^{\infty} \beta^t c_t = \sum_{t=0}^{\infty} \beta^t w_t h_t .$$

Taking first-order conditions, one gets

$$\frac{1}{c_t} = \lambda$$

$$\alpha h_t^{\gamma} = \lambda w_t .$$

The second equation has motivated people to run the following regression:

$$\ln h_t = B_0 + B_1 \ln w_t + \varepsilon_t .$$

Here, the coefficient B_1 is the estimate of $1/\gamma$. MaCurdy (1981), Altonji (1986), and Heckman and MaCurdy (1980) are early examples of studies that carry out this estimation on individual panel level data.[1] These studies typically find very small elasticities for prime-aged males, in the range of 0.3 or less, but much larger estimates for women. Intuitively, the underlying reason for the small elasticity estimates for men is that the hours profile is rather flat over the life cycle, whereas wages rise quite steeply, resulting in low covariation.

Mulligan (1995) argues that these traditional estimates are biased downward due to a failure to distinguish anticipated wage changes from those that are unanticipated or are artifacts of measurement error.[2] More recently, several authors have revised the original estimates in various ways (see, for example, Kimball and Shapiro 2003 and Pistaferri 2003) and found evidence of a labor supply elasticity in the range of 0.7–1.0 for men.

Domeij and Flodén (2006) argue that ignoring borrowing constraints can bias labor supply elasticity estimates downward. The intuition is that if an individual is credit constrained, the observation of high hours worked at a low wage does not provide evidence of the individual's willingness to intertemporally substitute labor supply. The authors find that the bias is on the order of 50 percent. Imai and Keane (2004), in turn, argue that the omission of endogenous human capital accumulation will bias labor supply elasticity estimates downward, as wages are not the correct measure of the opportunity cost of market time. Learning on the job provides an incentive to work when young at a low wage, as it leads to higher future wages. Thus, the opportunity cost of working is much flatter than the wage schedule. Imai and Keane (2004) find a labor supply elasticity in excess of three. Wallenius (2007), however, argues that this estimate is biased upward, and that adding skill accumulation does not lead to elasticity estimates that are much greater than one, an argument in line with the more recent literature.

Many economists proceed as if the estimate of γ from the microeconomic analysis is the value that should be used in aggregate models. In what follows, given micro observations and aggregation theory, we argue that the aggregate elasticity of labor supply should be much larger than the individual labor supply elasticity. In other words, micro observations support rather than cast doubt on the macro findings.

Indivisible Labor

Labor economists use important counterfactual predictions of the model to estimate what they call the *individual labor supply elasticity*. One such prediction is that everyone will make the same adjustment to hours worked in percentage terms. Empirically, however, this is not the case. Total hours worked is the multiple of employment and hours worked by those who are working. Over the business cycle, most of the adjustment in total hours arises from changes in employment, not hours worked by those who are employed. To be precise, Cho and Cooley (1994) document that three-quarters of the variation in total hours of work arises from movements in and out of the labor force. Many different factors impact employment—the fraction of lifetime worked, weeks of vacation, and holidays, to name a few.

In a model with a standard labor-leisure decision where labor is divisible and the household decides what fraction of the time endowment to devote to work each period, the labor supply elasticity depends on the utility function. Specifically, the parameter governing the curvature of the disutility from working, γ, is the key preference parameter. Rogerson (1984, 1988) proposes a framework with indivisible labor, where people either work some fixed workweek or do not work at all. In such a framework, the elasticity of substitution of labor across periods for the aggregate economy is independent of the elasticity of substitution implied by the individuals' utility functions. Moreover, the aggregate labor supply is much more elastic than when labor is divisible. This is true up to the point where all are employed.

Consider a static economy that is populated by a continuum of identical agents of measure one. Each agent is endowed with one unit of time. Time is indivisible, implying that the agent supplies either the entire unit of time to the market or none at all. Agents have an identical utility function given by

$$u(c) - v(h),$$

where c is consumption, h is labor, function u is increasing and concave, and function v is increasing and convex. With labor assumed indivisible, the only values of the $v(h)$ function that matter are $v(0)$ and $v(1)$. Assume that $v(0) = 0$ and that $v(1) = b,$ where b is a positive constant.

The individual agent's decision problem is then given by

$$\max u(c) - bh$$
$$\text{s.t. } c = wh,\ c \geq 0,\ h \in \{0,1\}.$$

A decreasing returns to scale production function uses only labor to produce output $F(H)$.

In several studies, Rogerson (1984, 1988) introduces lotteries where a social planner chooses a fraction φ of the population to work. Let c_w and c_n denote consumption for someone who is working and someone who is not working, respectively. The problem now becomes one of choosing φ, c_w, and c_n in the following problem:

$$\max \varphi[u(c_w) - b] + (1-\varphi)u(c_n)$$
$$\text{s.t. } \varphi c_w + (1-\varphi)c_n = F(\varphi),\ c_w \geq 0,\ c_n \geq 0,\ 0 \leq \varphi \leq 1.$$

For a given individual, the probability of working is φ. The first-order conditions for c_w and c_n imply $c_w = c_n = c$. This in turn implies that the social planner's problem can be rewritten as

$$\max u(c) - \varphi b$$
$$\text{s.t. } c = F(\varphi).$$

Since $\varphi = H$, this is simply a special case of the representative agent, divisible labor model with linear disutility from working. The implication is that an economy populated by individuals with identical preferences behaves as if populated by a single agent with preferences unlike those of any individual. In the presence of nonconvexities (resulting from indivisible labor supply), the aggregate is very different from the individual entities that are being aggregated. This case has a well-known parallel on the production function side.

In mapping this specification to the more standard ones in the previous two sections, one notes that assuming indivisible labor amounts to assuming $\gamma = 0$. If one takes indivisible labor as the starting point, estimating γ from micro data is irrelevant.

Hansen (1985) extends this analysis to the business cycle setting. He finds that the economy with indivisible labor displays larger fluctuations than the one with divisible labor.

Labor Supply Elasticity Function of Preference and Technology Parameters

The amount of labor supplied by an individual over his or her lifetime is effectively characterized by the fraction of lifetime spent working and hours worked when employed. Instead of thinking in terms of a lottery that determines who works and who does not, the problem can be recast as one in which the individual chooses the fraction of his or her lifetime to devote to work. Prescott, Rogerson, and Wallenius (2009) develop a simple, tractable framework that delivers this characterization in equilibrium (here, we follow Ljungqvist and Sargent 2007).

A key feature of their model is a nonconvex mapping from hours supplied to the market to labor services. In particular, they assume that if an individual supplies h units of time to the market, this yields l units of labor services, where

$$l = g(h).$$

The function g is initially convex and later concave. The former is intended to capture the fixed costs associated with getting set up in a job and being supervised, whereas the latter is included to allow for fatigue.[3,4] With this mapping, people will choose to work some fraction of their lifetime instead of spreading work evenly throughout their lifetime. In fact, the individual choice problem can be formulated as choosing a fraction e of his or her lifetime to work and the hours of work h to be supplied when working. Each individual, therefore, solves

$$\max \log(c) - e\tilde{v}(h)$$

$$\text{s.t. } c = (1-\tau)eg(h) + T,\ 0 \le e \le 1, 0 \le h \le 1.$$

The assumption is that the government taxes all labor income at the constant rate of τ and uses the tax revenues to fund a lump-sum transfer T. The authors also assume that the government balances the budget, implying that

$$T = \tau e g(h).$$

Using the first-order conditions to derive expressions for the optimal length of the workweek and the fraction of time spent in employment yields

$$\frac{\tilde{v}'(h)}{\tilde{v}(h)} = \frac{g'(h)}{g(h)}$$

$$e = \frac{1-\tau}{\tilde{v}(h)}.$$

From these expressions, it becomes apparent that the model implies a large aggregate labor supply elasticity in response to changes in tax and transfer programs. In fact, the elasticity of *eh* with respect to $1-\tau$ is equal to 1. At the same time, the model predicts zero elasticity for hours of work of continuously employed individuals. In this respect, the model mimics the indivisible labor model discussed previously. *A key message of the study by Prescott, Rogerson, and Wallenius (2009) is that the aggregate labor supply elasticity with respect to changes in taxes is a function of both preference and technology parameters*. In particular, the mapping from hours supplied to the market to labor services is critical in determining the aggregate labor supply elasticity.

Life Cycle Model with Extensive and Intensive Margins of Labor Supply

Rogerson and Wallenius (2009) embed the Prescott, Rogerson, and Wallenius (2009) framework into a life cycle setting. Nonconvexities in the mapping from time devoted to market work to labor services again give rise to allocations where individuals choose both the fraction of lifetime to devote to employment (extensive margin) and hours worked when employed (intensive margin). Embedding the analysis in a life cycle model enables them to generate standard life cycle profiles for hours of work, most notably that hours of work drop discontinuously to zero at older ages. Note that in this life cycle framework, the timing of work is no longer indeterminate, as was the case in the Prescott, Rogerson, and Wallenius (2009) framework.

Consider a continuous time overlapping-generations framework in which a unit mass of identical, finitely lived individuals is born at each instant of time. Letting *a* denote age, individuals have preferences over paths for consumption $(c(a))$ and hours worked $(h(a))$:

$$\int_0^1 \left[\log c(a) - \alpha \frac{h(a)^{1+\gamma}}{1+\gamma} \right].$$

An individual who devotes $h(a)$ hours to market work produces $l(a)$ units of labor services, where $l(a) = e(a)g(h(a))$. The $e(a)$ function denotes an exogenous, age-varying productivity profile, which results in hours worked varying over the life cycle. For simplicity, the profile is assumed to be piecewise linear. The $g(h)$ function is again a nonconvex mapping from hours worked to labor services, which serves to endogenize the length of the working life. Hours worked exhibit a reservation property, with people choosing to work above a certain productivity and not to work below it.

Given a value of γ, the size of the nonconvexity, the productivity profile, and the disutility from working parameter are chosen to match three target values: a working life of two-thirds, peak hours of 45 hours per week, and a doubling of wages over the life cycle.

Rogerson and Wallenius (2009) are interested in studying how the value of γ affects the life cycle profile for hours and how it in turn responds to changes in labor tax rates. The value of γ is therefore varied over a wide range. Given a value of γ and the calibrated parameters, the model generates life cycle profiles for hours and wages. The framework therefore allows Rogerson and Wallenius (2009) to reproduce micro estimates of the labor supply elasticity based on life cycle variation for prime-aged workers. More importantly, they are able to simultaneously carry out standard macro estimation based on variation in aggregates across steady states as tax rates are altered. They find that macro elasticities are virtually unrelated to micro elasticities, and moreover that macro elasticities are large. Although the micro elasticity is virtually irrelevant for the aggregate elasticity with respect to taxes, it does matter for how the tax response is broken down between the extensive and intensive margins of labor supply. Specifically, the smaller the micro elasticity is, the larger the share of the action on the extensive margin.

There has long been a need for a theory that is consistent with the above micro- and macroeconomic observations. This article presents such a framework. On a related note, Chang and Kim (2006) construct a framework in which the aggregate labor supply elasticity depends on the heterogeneity of the cross-sectional wage distribution. They also find that macro and micro elasticities can be significantly different, with macro elasticities considerably larger than micro elasticities. In their framework, however, all adjustment takes place along the extensive margin of labor supply. In this respect, the analysis is similar to that of Rogerson (1984, 1988) and Hansen (1985).

The key message from these analyses is that we should not estimate parameter values in one setting and apply them to a different one. Rather, we should work with frameworks in which the choice problem of an individual is explicitly formulated and try to identify the underlying structural parameters of that problem. This message is similar in spirit to that of Browning, Hansen, and Heckman (1998).

Relating the Life Cycle Model to the Representative Household Model

We have seen that in a life cycle model with an extensive and intensive margin of labor supply, micro and macro elasticities are virtually unrelated. Given that the stand-in household model has proven to be a useful abstraction in many settings, suppose one wanted to mimic a life cycle model with a single agent model with no intensive and extensive margin. What is the labor supply elasticity that should be used in such a model? Rogerson and Wallenius (2009) show that a stand-in household model with a relatively high labor supply elasticity can reproduce the steady-state effects of taxes on aggregate hours that they find in their life cycle model. It is worth mentioning that the elasticity of the stand-in agent model is not the labor supply elasticity of any given individual; rather, it is capturing the heterogeneity in the data.

Connection between Retirement and the Intertemporal Elasticity of Substitution

The typical retirement pattern is a transition from full-time work directly into little or no work. In a recent paper, Rogerson and Wallenius (2010) argue that this transition contains important information on the value of the intertemporal elasticity of substitution. The intuition underlying their argument is that since retirement represents a very large change in leisure, the fact that individuals willingly incur such a significant change in leisure should provide information about their willingness to intertemporally substitute.

Rogerson and Wallenius (2010) consider models where retirement is an optimal property of life cycle labor supply and, moreover, where nonconvexities are the key feature generating retirement. In other words, in the presence of nonconvexities, people find it optimal to concentrate work in some fraction of their lifetime, as opposed to spreading it evenly throughout their lifetime. In their paper, the authors consider different sources of nonconvexities, namely, the fixed time and consumption costs associated with work, and nonlinear wage-hours schedules. They show that although nonconvexities in production can generate retirement, the size of the nonconvexities needed to do so increases sharply as the intertemporal elasticity of substitution for labor decreases. It is, therefore, very difficult to rationalize values of the intertemporal elasticity of labor supply that are below 0.75, given empirically reasonable values for the extent of nonconvexities.

Fraction of Lifetime Worked

Prescott, Rogerson, and Wallenius (2009) and Rogerson and Wallenius (2009) define the aggregate labor supply elasticity as the responsiveness of aggregate hours to a change in tax rates. Both studies model tax and transfer programs simply as a proportional tax accompanied by a lump-sum transfer. A natural extension is to model tax and transfer programs in greater detail. In particular, it is of interest to study whether modeling the earnings dependence of certain transfers, such as Social Security, greatly affects the results.[5]

Despite the success of the stand-in household construct in addressing many questions, it is not a good abstraction for thinking about retirement and Social Security reform. For these questions, one needs a life cycle model. We have already established that the extensive margin of labor supply is a very important margin for understanding business cycles as well as differences in aggregate labor supply across countries and time. When we look at the data, it is apparent that differences along the extensive margin are dominated by the young and the old. This finding naturally points to Social Security as a potential source of differences in the labor supply behavior of older workers.

Wallenius (2009) builds a general equilibrium model of life cycle labor supply that features endogenous retirement and human capital accumulation, which is parameterized to match U.S. data on life cycle profiles for hours worked and wages. The model is used to study the extent to which differences in Social Security, and more generally tax and transfer programs, can account for the cross-country differences in aggregate hours worked between the United States and continental Europe. Wallenius (2009) finds that differences in Social Security account for 35 percent to 40 percent of the cross-country differences in aggregate hours between the United States and Belgium, France, and Germany. Once other differences in labor taxation in addition to Social Security are included in the analysis, the model implies that tax and transfer programs account for roughly 60 percent of the difference in aggregate hours worked between the United States and continental Europe. Similar to Rogerson and Wallenius (2009), the aggregate responses are not sensitive to the micro labor supply elasticity.

On a related note, İmrohoroğlu and Kitao (2009) show that the effects of various forms of Social Security reform are invariant to reasonable values of the labor supply elasticity.

Note that the extensive margin is important not only at the individual level in determining the fraction of lifetime spent in employment, but also at the household level. In particular, the effect of changes in tax policy can have large implications for the secondary wage earner in the household (see Guner, Kaygusuz, and Ventura 2010).

Conclusions and Directions for Future Research

Our goal in writing this article has been to highlight the importance of the labor supply elasticity for policy questions, such as business cycles and tax policy analysis, and to persuade the reader that the micro elasticity and the aggregate elasticity of labor supply are indeed distinct. We have striven to do this in a simple, easily accessible way.

The micro elasticity is the preference parameter governing the curvature of the individual's disutility from working. This is a well-defined concept. Traditional estimates of this elasticity, based on the covariation of hours worked and wages over the life cycle for prime-aged males, are small. However, the micro evidence along with aggregation theory predicts that the aggregate labor supply elasticity will be much higher than the elasticity of the individuals being aggregated. The labor economists' estimates would be good estimates of the aggregate elasticity of labor supply only in empirically uninteresting worlds such as a Robinson Crusoe world with limited ability to transform current consumption into future consumption.

The starkest example of the failure of aggregation is the indivisible labor supply model. Although the indivisible labor supply model can be credited with being the first to clearly highlight the distinction between the micro and macro elasticities, it does not allow one to reconcile aggregate observations with life cycle properties of labor supply. For this purpose, a life cycle model with an intensive and extensive model of labor supply is needed. Such a model allows one to carry out standard microeconomic analysis of the labor supply elasticity based on life cycle variation for prime-aged workers, as well as macro estimation based on variation in aggregates across steady states as tax rates are altered. Here, the aggregate labor supply elasticity is defined as the responsiveness of aggregate hours to changes in the effective wage.

Having stipulated that the micro and macro labor supply elasticities differ, one can ask when it makes sense to talk about the micro elasticity. In the stark case of indivisible labor supply, the micro elasticity has no real meaning. In a model with both intensive and extensive margins, the micro elasticity governs responses along the intensive margin. These responses are certainly of interest in their own right.

Given the resounding success of the stand-in household abstraction for many questions, it is important to know how the life cycle model maps into the stand-in household construct. A stand-in household model with a relatively high labor supply elasticity can reproduce the steady-state effects of taxes on aggregate hours from the life cycle model. It should be noted that the elasticity of the stand-in agent model is not the labor supply elasticity of any given individual; rather, it is capturing the heterogeneity in the data.

What we have gleaned from the recent advances in the literature is that models with flexible labor supply should match aggregate observations as well as life cycle properties of labor supply. The key facts that these models should strive to match are as follows: (1) a micro labor elasticity of around one, which is in accordance with more recent estimates of this parameter; (2) an aggregate labor supply elasticity of around three, which is consistent with both business cycle theory and cross-country tax analysis; and (3) hours worked varying with age or productivity (or both) over the life cycle and dropping discontinuously to zero at the time of retirement.

Although tremendous advances have been made in the field of macroeconomic analysis, many interesting questions remain. We have already established that the extensive margin of labor supply is the most important margin for understanding business cycles as well as differences in aggregate labor supply across countries and time. When we look at the data, it is apparent that differences along the extensive margin are dominated by the young and the old. This finding naturally points to Social Security as a potential source of differences in the labor supply behavior of older workers. Despite the success of the stand-in household construct in addressing many questions, it is not a good abstraction for thinking about retirement and Social Security reform. For these questions, one needs a life cycle model. Wallenius (2009) takes a step in this direction, but much further work is still warranted.

Recall that the extensive margin is important at the individual level in determining the fraction of lifetime spent in employment, but it is also important at the household level. The effect of changes in tax policy can have large implications for the secondary wage earner in the household (see Guner, Kaygusuz, and Ventura 2010). Work remains to be done in successfully modeling household labor supply decisions.

Notes

[1]For a more complete survey of this literature, see, for example, the study by Pencavel (1986).

[2]Mulligan (1995) also notes that the approach of MaCurdy (1981), Altonji (1986), and others ignores certain key features of the micro data, such as seasonal variation. Accounting for seasonal variation, he estimates a large labor supply elasticity. Another likely source of downward bias is that much production is club production, which gives rise to all working the same number of hours at an establishment. When someone is promoted to a higher-paying supervisory position, that person does not have the option to work longer hours.

[3]Note that in Rogerson (1984, 1988), the nonconvexity was due to a discrete choice in hours, whereas in Prescott, Rogerson, and Wallenius (2009), hours worked are a continuous choice variable, despite the presence of the nonconvexity.

[4]This nonlinearity implies that workweeks of different lengths are not perfect substitutes in generating labor services, which is in line with Hornstein and Prescott (1993).

[5]Rogerson (2007) stresses that what the government does with the tax revenue affects the distortive effects of labor taxes on labor supply.

References

Altonji, Joseph G. 1986. Intertemporal substitution in labor supply: Evidence from micro data. *Journal of Political Economy* 94 (3): S176–S215.

Browning, Martin, Lars Peter Hansen, and James J. Heckman. 1998. Micro data and general equilibrium models. In *Handbook of macroeconomics*, vol. 1, pt. A, ed. John B. Taylor and Michael Woodford, 543–633. Amsterdam: North-Holland.

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.

Cho, Jang-Ok, and Thomas F. Cooley. 1994. Employment and hours over the business cycle. *Journal of Economic Dynamics and Control* 18 (2): 411–32.

Domeij, David, and Martin Flodén. 2006. The labor-supply elasticity and borrowing constraints: Why estimates are biased. *Review of Economic Dynamics* 9 (2): 242–62.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2010. Taxation and household labor supply. Working paper. Department of Economics, University of Iowa.

Hansen, Gary D. 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16 (3): 309–27.

Heckman, James J., and Thomas E. MaCurdy. 1980. A life cycle model of female labour supply. *Review of Economic Studies* 47 (1): 47–74.

Hornstein, Andreas, and Edward C. Prescott. 1993. The firm and the plant in general equilibrium theory. In *General equilibrium, growth, and trade II: The legacy of Lionel McKenzie*, ed. Robert Becker, Michele Boldrin, Ronald Jones, and William Thomson, 393–410. San Diego: Academic Press.

Imai, Susumu, and Michael P. Keane. 2004. Intertemporal labor supply and human capital accumulation. *International Economic Review* 45 (2): 601–41.

İmrohoroğlu, Selahattin, and Sagiri Kitao. 2009. Labor supply elasticity and Social Security reform. Working Paper 2009-5. Center for Retirement Research at Boston College.

Kaldor, Nicholas. 1957. A model of economic growth. *Economic Journal* 67 (268): 591–624.

Kimball, Miles S., and Matthew D. Shapiro. 2003. Labor supply: Are the income and substitution effects both large or both small? Working paper. University of Michigan.

Kydland, Finn E., and Edward C. Prescott. 1982. Time to build and aggregate fluctuations. *Econometrica* 50 (6): 1345–70.

Ljungqvist, Lars, and Thomas J. Sargent. 2007. Do taxes explain European employment? Indivisible labor, human capital, lotteries, and savings. In *NBER macroeconomics annual 2006*, vol. 21, ed. Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 181–246. Cambridge, MA: MIT Press.

Lucas, Robert E., Jr., and Leonard A. Rapping. 1969. Real wages, employment, and inflation. *Journal of Political Economy* 77 (5): 721–54.

MaCurdy, Thomas E. 1981. An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy* 89 (6): 1059–85.

Mulligan, Casey. 1995. The intertemporal substitution of work—What does the evidence say? University of Chicago–Population Research Center Discussion Paper Series 95-11.

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): 1353–62.

Pencavel, John. 1986. Labor supply of men: A survey. In *Handbook of labor economics*, vol. 1, ed. Orley Ashenfelter and Richard Layard, 3–102. Amsterdam: North-Holland.

Pistaferri, Luigi. 2003. Anticipated and unanticipated wage changes, wage risk and intertemporal labor supply. *Journal of Labor Economics* 21 (3): 729–54.

Prescott, Edward C. 2004. Why do Americans work so much more than Europeans? *Federal Reserve Bank of Minneapolis Quarterly Review* 28 (1): 2–13.

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.

Rogerson, Richard. 1984. Topics in the theory of labor markets. PhD dissertation. University of Minnesota.

———. 1988. Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics* 21 (1): 3–16.

———. 2007. Taxation and market work: Is Scandinavia an outlier? *Economic Theory* 32 (1): 59–85.

Rogerson, Richard, and Johanna Wallenius. 2009. Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory* 144 (6): 2277–92.

———. 2010. Fixed costs, retirement, and the elasticity of labor supply. Manuscript. Arizona State University.

Wallenius, Johanna. 2007. Human capital accumulation and the intertemporal elasticity of substitution of labor. Manuscript. Arizona State University.

———. 2009. Social Security and cross-country differences in hours worked: A general equilibrium analysis. Manuscript. Arizona State University.