Labor supply response in macroeconomic models: Assessing the empirical validity of the intertemporal labor supply response from a stochastic overlapping generations model with incomplete markets

Juan Contreras and Sven Sinclair

Congressional Budget Office

17. September 2008

Working Paper Series
2008-07

Congressional Budget Office
Washington, D.C.
September 2008

Labor Supply Response in Macroeconomic Models

Assessing the Empirical Validity of the Intertemporal Labor Supply Response from a Stochastic Overlapping Generations Model with Incomplete Markets

Juan M. Contreras and Sven H. Sinclair
Congressional Budget Office

Authors: juan.contreras@cbo.gov, sven.sinclair@cbo.gov.

Labor Supply Response in Macroeconomic Models

Assessing the Empirical Validity of the Intertemporal Labor Supply Response from a Stochastic Overlapping Generations Model with Incomplete Markets

Juan M. Contreras and Sven H. Sinclair[1]

Congressional Budget Office

September 2008

Abstract

We evaluate the labor supply response in a stochastic overlapping generations model with incomplete markets and a non separable utility function in labor and consumption. Using a simulated panel from the model, we calculate the labor supply response to anticipated changes in wages (holding the marginal utility of wealth constant-that is, the Frisch elasticity) and to unanticipated change in wages (which describes the effect of uncertainty in labor supply responses). The model's Frisch elasticity estimate is 0.33, which is slightly higher than the empirical estimates in the earlier literature but somewhat lower than more recent estimates. The paper also shows that the borrowing constraints in the model reduce substantially the estimates of the Frisch elasticity. The labor supply response to an unanticipated change in wages is small because of large wealth effects. Having all the variables required and no measurement error, we calculate the omitted variable bias of not controlling for the level and variance (risk) of the unexpected changes in wages. Omitting both variables biases the estimates of the Frisch elasticity downward by a factor of 8; omitting measures of wage risk alone biases it by a factor of 1.4.

JEL CODES: J22, D91, D58

Keywords: labor supply, intertemporal substitution, computable GE models.

[1]We thank Doug Hamilton for very useful comments and for important insights about the relevance of the wealth effect in the labor supply response, and Bob Dennis and Thomas Woodward for comments on earlier drafts.

1 Introduction

The validity of dynamic macroeconomic models for analyzing policy experiments depends on key behavioral parameters that determine the agents' responses to changes in policy variables. The intertemporal elasticity of labor supply, which determines the labor supply response, is among the most important of those parameters.

In this paper, we analyze the empirical validity of the labor supply response in a stochastic overlapping generations model with incomplete markets.[2] We estimate the labor supply response in the model to anticipated changes in wages, holding constant the marginal utility of wealth (Frisch elasticity), and also to unanticipated changes in wages.[3] We calculate the Frisch elasticity on the intensive margin (that is, we account for the hours decision and not for the participation decision given) because that approach is consistent with most of the empirical studies that estimate the wage elasticity of labor supply.

This paper makes three main contributions. First, it reconciles the Frisch elasticity estimates from the empirical labor literature with the values used in the macroeconomic literature, using a methodology that is consistent across fields. Second, it emphasizes the importance and quantifies the magnitude of the effect of wage uncertainty in the labor supply response in macroeconomic models. Third, it quantifies the omitted variable bias of not controlling for the level and variance (risk) of the unexpected changes in wages.

Many macroeconomic models focus on the Frisch elasticity values because their goal is to analyze real business cycles. Equally important but less studied is the life cycle

[2]This is the model that CBO regularly uses in its analysis of the President's budget.

[3]Economists usually consider four basic elasticities of labor supply (Blundell and MaCurdy [2008]): compensated and uncompensated elasticity of labor supply, Frisch elasticity of labor supply, and a life cycle elasticity that considers unanticipated (parametric) changes in wage changes. Because the compensated and uncompensated elasticities of labor supply usually are derived and estimated from static models, they are unsuitable for application in dynamic macroeconomic models. The Frisch elasticity of labor supply is defined as the percentage change in labor supply resulting from a one percent increase in the expected wage rate, holding the marginal utility of wealth constant. The Frisch elasticity measures the response of wages to anticipated wage changes, as opposed to the other life cycle elasticity, which measures the response of work hours to unanticipated wage changes or shifts in the wage profile.

elasticity that considers labor supply responses to unanticipated changes in wages. This elasticity is also very important for the study and the analysis of tax and benefit reforms, which almost always cause unanticipated shifts in wages today and in the future (see Blundell and MaCurdy 1999 and 2008). Of course, the relative importance of the two elasticities depends on the type of tax policy implemented and on the mix of wealth and intertemporal substitution effects it triggers.

The estimates of the Frisch elasticity of labor supply that come from the model are consistent with estimates from the empirical literature, once borrowing constraints and the level and variance of unexpected changes in wages are accounted for. The Frisch elasticity from macroeconomic models often differs from that estimated in microeconomic studies.[4] Previous work has posited that those differences could be accounted for in two ways. The first way is considering the entry to and exit from the labor market—the extensive margin at which labor supply is adjusted.[5] Although the elasticity at the intensive margin is generally small, adjustments at the extensive margin are larger.

The second way to account for the differences between the Frisch elasticity coming from macroeconomic models and from the empirical studies is the consideration of borrowing constraints.[6] Commonly used utility functions, including the constant relative risk aversion (CRRA) utility, can be shown analytically to have high Frisch elasticities when there is no uncertainty or binding constraints. Most empirical studies find much lower elasticities. However, borrowing constraints play a central role in our model; indeed, we also show in this paper that once we account for borrowing constraints, the Frisch elasticity that comes from the model using the CRRA utility is reduced by a factor of nearly 4. Borrowing constraints are important because without them house-

[4]Blundell and MaCurdy (2008) argue that the aggregation of individual agents into a representative agent to describe the aggregate labor supply elasticity requires stringent assumptions. That criticism does not apply to the class of models used in this paper, because the agents are ex-post heterogeneous in wealth, and we can simulate a longitudinal panel of agents to estimate the elasticity of labor supply with microeconomic techniques.

[5]The elasticities of labor supply can refer to the extensive margin (labor force enter or exit), or to the intensive margin (change in the number of hours worked). Rogerson and Wallenius (2007) propose an elegant model that reconciles the differences between the relatively small micro estimates of the Frisch elasticity with the relatively large macro responses.

[6]See Domeij and Floden (2006).

holds respond to a bad wage shock by reducing labor and increasing borrowing so as to smooth consumption over time. With borrowing constraints, however, households can't completely smooth consumption, so they don't reduce labor supply as much as they otherwise would. That dampens the labor supply response in the model.

We analyze a third factor that affects the difference between the Frisch elasticity from empirical estimates and the Frisch elasticity from macroeconomic models: the omitted variable bias that is present in many empirical studies. If we do not control for the variance of the unexpected changes in future wages as a measure of wage risk or for both the variance and the level of unexpected changes in future wages, the Frisch elasticity estimates are biased downward. If we omit measures of level and variance of unexpected changes in wages, the point estimate of the Frisch elasticity is reduced by a factor of 8 (0.33 vs. 0.04). If we omit measures of the variance of unexpected changes in wages (wage risk) the point estimate is 1.4 times smaller (0.33 compared with 0.24). The intuition behind the omitted variable bias is related to the presence of the wealth and the intertemporal substitution effect. If the econometrician does not control for unexpected wage changes, the error term would be correlated with the coefficient on wages that measures the intertemporal substitution (Frisch) effect.

Our estimates from the model of the labor supply response to unanticipated wage changes are small because wealth effects in the model largely offset the intertemporal substitution effect. That result comes from important wealth effects in our OLG model. Unexpected wage changes shift the wage profile, generating a labor supply response from a wealth effect and from an intertemporal substitution (Frisch) effect. A wage increase generates a wealth effect and a consequent labor supply reduction, but the same wage increase generates an intertemporal substitution effect that tends to increase the labor supply. Depending on which effect is bigger, the labor supply response to unanticipated wage changes would be negative or positive. In our model, wealth effects have almost the same magnitude as intertemporal substitution effects.

Our analysis used a stochastic OLG model with incomplete markets that was similar to those of Nishiyama (2002), Nishiyama and Smetters (2005) and CBO (2008).

Households in the model have a CRRA utility function that is non separable between consumption and leisure. First, we use that model to show analytically how borrowing constraints affect the Frisch elasticity measures. We then generate a simulated panel from the model and use the Pistaferri (2003) econometric method to estimate the Frisch elasticity of labor supply. Next we compare the estimates with estimates from the empirical literature and quantify the omitted variable bias that is induced by the lack of measures of the level and variance of unexpected changes in future wages. We can calculate that bias because the variables from the simulated panel are known exactly and thus have no measurement error.

In the next section we review the literature that estimates the Frisch elasticity of labor supply and put the different estimates in perspective. In section 3, we discuss methodological and empirical issues regarding the measurement of the Frisch elasticity, present our estimates of various measures of intertemporal labor supply elasticities, and quantify the omitted variable bias from the empirical estimates. We present our conclusions in section 4.

2 Literature Review

Most empirical studies of the intertemporal elasticity of labor supply focus on the labor supply response to anticipated wage changes or Frisch elasticity (with the additional qualification of holding the marginal utility of wealth constant); just a few, among them MaCurdy (1981) and Pistaferri (2003), estimate the intertemporal elasticity of labor supply in response to unanticipated wage changes. MaCurdy estimates a male labor supply response to unanticipated wage changes of 0.08 in the median range; Pistaferri's estimate is 0.5.[7]

[7] Early studies showed that the female intertemporal response is much larger. Heckman and MaCurdy (1980, 1982) estimated a Frisch labor supply elasticity for women of 1.6. Blundell, Meghir, and Neves (1993) estimated a Frisch elasticity between 0.8 and 1.2 for married women with children and an elasticity of 0.6 for childless women. The high elasticities obtained for women in part reflect the importance of considering the decisions of entry and exit from the labor market not present in the studies that consider only the decision of hours worked for males. Heckman (1979, 1993) and Blundell and MaCurdy (1999) argued that most of the labor supply response stems from choices at the extensive

Empirical measures of the intertemporal Frisch elasticity of labor supply present a huge variation. Estimates range from close to zero to about 3, although most have values of less than 1. The variation arises from differences in methodology, in the samples used, and in the type of population and identification strategies.

MaCurdy (1981) and Altonji (1986) were among the first to identify the intertemporal response of hours worked to changes in wages for a sample of males in the United States. MaCurdy's estimates of the Frisch elasticity of labor supply are between 0.1 and 0.45. Altonji's estimates are between 0 and 0.35. Pencavel (1986) surveys this early literature and presents a mean value of 0.2 and a range of 0 to 0.45.

Among the difficulties associated with early work estimating the Frisch elasticity of labor supply are measurement error, lack of data about expected wages and difficulty in accounting for wealth effects through the life cycle.[8] Angrist (1991) addressed some of the problems associated with measurement error and estimated Frisch labor supply elasticities of 0.6 to 0.8. French (2004) also tried to address measurement error problems finding a Frisch elasticity close to zero with a standard error of 0.25.

Moreover, Lee (2001) addresses the finite sample bias in estimates of intertemporal labor supply. He finds that the estimates of the Frisch elasticity of many studies in the field are biased downward because of small samples and, when that bias is corrected, the Frisch labor supply elasticities for men rise from 0–0.2 to about 0.5.

Domeij and Floden (2006) argue that when the econometrician ignores borrowing constraints, the estimated intertemporal Frisch elasticities have a downward bias. They also identify an additional downward bias in many studies from the approximation errors of log-linearizing the Euler equation.

Most recent studies tend to find higher values than was the case earlier. Estimates generated by Ziliak and Kneisner (2005) are around 0.5. Pistaferri (2003) estimates a

(entry-exit) margin. In his Nobel lecture, Prescott (2006) emphasized this point; he argued for a Frisch elasticity of 3.

[8]See Card (1994) for a critical review of the basic intertemporal model. He points out the difficulty of identifying the parameters that govern the intertemporal substitution of labor supply and the ambiguity in the literature that is attributable to the presence of wealth effects.

Frisch elasticity of 0.75 in a paper that improves on previous studies by avoiding the use of weak instruments to estimate expected wage rates. The author instead relies on a unique data set from the Bank of Italy that contains data both on an individual's hours and on their expectations of future earnings (but not wages).[9]

Researchers who take a more structural approach and estimate directly the intertemporal elasticity of labor supply from the life cycle model also report much higher values. Gourinchas and Parker (2002) present estimates that range from 0.7 to almost 2. Imai and Keane (2004) report that accounting for human capital accumulation can lead to huge Frisch elasticities of labor supply (more than 3). Laitner and Silverman (2005) present estimates that are around 0.9.

Our paper is related to Chang and Kim (2006), who use an infinite horizon model with incomplete markets and assume that consumption and leisure are separable in the utility function. In contrast, we use an OLG model and a CRRA utility function non-separable between leisure and consumption. Chang and Kim's utility function implies a direct mapping between the intertemporal substitution elasticity of leisure and the Frisch elasticity; in our model it does not. Another difference is that we can compare our results directly with previous results from the empirical literature because we apply standard econometric techniques used in previous microeconomic studies for comparison. Finally, Chang and Kim do not estimate the sources of bias present in the empirical studies nor the labor supply response to unexpected changes in wages

This paper also is similar to the work of Domeij and Floden (2006) in that it considers the effects of uncertainty and borrowing constraints on the estimates of the Frisch elasticity of labor supply. It differs from Domeij and Floden (2006) in three major ways: First, instead of using a model with infinitely lived agents, we use a life cycle model to compare estimates from a simulated panel with the estimates that do not consider borrowing constraints and uncertainty. Second, we focus on a different type of bias present

[9]Fehr and Götte (2007) estimate an intertemporal elasticity of substitution of about 1.1 by using a unique data set that measures individual workers' responses to a preannounced change in wage for a specific piece rate (service provided by bicycle messenger) service. Because they use high-frequency data and estimate an elasticity that does not hold the marginal utility of wealth constant, their results are not relevant for the model used in this paper or for the calculation of the Frisch elasticity.

in the Frisch elasticity estimates: the omitted variable bias. Finally, we use a different method to measure the Frisch elasticity of labor supply. We borrow Pistaferri's (2003) methodology, which accounts for the labor supply response to unexpected changes in wages.

3 Quantifying the Intertemporal Labor Supply Response Using a Stochastic OLG Model with Incomplete Markets

In this section we lay down a theoretical framework to analyze two types of intertemporal elasticity of labor supply. The first measures the response of hours of work to anticipated wage changes, holding the marginal utility of wealth constant (that is, the Frisch elasticity). The second measures the response of hours of work to unanticipated wage changes. Both elasticities are relevant to the analysis of tax and benefit reforms, and their relative importance depends on whether the reforms favor wealth or intertemporal substitution effects differently. We discuss the empirical challenges faced in estimating the elasticities and the issues that arise when we want to map the results from the model to the empirical estimates. We quantify the Frisch elasticity directly from the model, not accounting for borrowing constraints, and we quantify the same elasticity accounting for the borrowing constraints using data from a simulated panel. We estimate from the model the labor supply response to unanticipated wage changes. Finally, we quantify the omitted variable biases of previous estimates of the Frisch elasticity of labor supply.

3.1 OLG Model Used to Generate Synthetic Data

This paper uses an OLG growth model with uninsurable idiosyncratic working ability shocks and uncertain life span.[10] The unit of analysis is married households, but we

[10] The base model is similar to those used Aiyagari (1994), Huggett (1996), and many others, although Aiyagari assumed infinitely-lived agents. The model is an extension of that used by Nishiyama (2002)

treat them as a single decision unit.[11] Accordingly, we calibrate the number of hours worked in the model as the sum of two individuals' total hours. The consumer side is the one relevant to the determination of the intertemporal elasticity of labor supply. The production side, which features a perfectly competitive representative firm, and the government sector with exogenously determined policies, only determine the interest rate and the wage rate that feed into the consumer's problem. For this reason, we just briefly analyze here the consumer part of the model. For more detailed information on the model, see Nishiyama (2002) and Nishiyama and Smetters (2005). Table 1 lists the parameters used in the calibration.

Table 1. Parameters for the model

Time preference parameter	β	0.986
Share parameter for consumption	α	0.449
Coefficient of relative risk aversion	γ	2.0
Capital share of output	θ	0.300
Depreciation rate of capital stock	δ	0.047
Long-term real growth rate	μ	0.018
Long-term population growth rate	ν	0.010
Total factor productivity	A	0.949

Subject to a budget constraint, households choose consumption and leisure. They face uninsurable idiosyncratic working ability shocks and mortality shocks and can hold only one type of assets. Formally, the household problem is expressed as follows:

$$\max_{c,h} E_t \sum_{t=0}^{T} \left(\frac{1}{1+\rho}\right)^t u(c_t, h_t, Z_t) \tag{1}$$

$$s.t. \quad a_{t+1} = (1+r_t)a_t + w_t h_t - c_t \quad \text{and} \quad a_t \geq \bar{a} \tag{2}$$

where ρ is the discount rate, c is the consumption level, h is the number of hours worked, Z is a vector of preference shifts, a is the asset level, r is the interest rate, $\bar{a}$ is the lower limit of assets a household can have, T is the total number of periods a household lives

and is very similar to the model presented by Nishiyama and Smetters (2005).

[11] Collective decision models within the households have been studied theoretically by Chiappori (1988, 1992) and Blundell et al (2007). We abstract from this issue in our analysis and treat the households as if it is a single decisionmaker.

and w is the stochastic wage rate. The optimality conditions for this problem are

$$\frac{\partial u(c_t, h_t, Z_t)}{\partial c_t} = \lambda_t \quad (3)$$

$$-\frac{\partial u(c_t, h_t, Z_t)}{\partial h_t} = \lambda_t w_t \quad (4)$$

$$E_t \left(\frac{1 + r_{t+1}}{1 + \rho} \right) \lambda_{t+1} + \mu_t = \lambda_t \quad (5)$$

λ_t represents both the marginal utility of wealth in period t and the Lagrange multiplier associated with the budget constraint in equations (3), (4), and (5). μ_t represents the Lagrange multiplier associated with the borrowing constraint. This multiplier would be equal to zero when the household does not face liquidity constraints and different from zero otherwise.

From equations (3) and (4), we can derive the Frisch elasticity of labor supply, or how the number of hours responds intertemporally to changes in the wage rate, holding the marginal utility of wealth constant, if the borrowing constraint in equation (2) does not bind. Dropping the time subscript, the expression for the Frisch elasticity without considering borrowing constraints is as follows:

$$\left.\frac{dh}{dw}\frac{w}{h}\right|_{\lambda}^{\text{no constraint}} = \frac{u_h u_{cc}}{u_{hh}u_{cc} - u_{hc}^2}\frac{1}{h} \quad (6)$$

If the borrowing constraint binds, the Lagrange multiplier μ will be different from zero. In that case, we can use equation (5) in addition to equations (3) and (4) to express the Frisch elasticity of labor supply, as follows:

$$\left.\frac{dh}{dw}\frac{w}{h}\right|_{\lambda}^{\text{borrowing constraint}} = \frac{u_h u_{cc}}{u_{hh}u_{cc} - u_{hc}^2}\frac{1}{h} * \frac{1}{E_t\left(\frac{1+r'}{1+\rho}\right)} * \left[1 + \frac{w\mu}{u_h}\right] \quad (7)$$

Note that the additional term in equation (7) is less than 1 because the marginal utility of labor is negative.[12] This implies that if borrowing constraints are binding, the

[12]The parameters' choice in this economy make $\left(\frac{1+r'}{1+\rho}\right) \leq 1$.

Frisch elasticity of labor supply is lower than it is when borrowing constraints are not binding.

We can use equation (6) to calculate directly from the model's parameters the Frisch elasticity of labor supply . However, that calculation is more complicated if the borrowing constraint binds (equation [7]) because we don't know the value of the Lagrange multiplier.

Equation (6) implies also that if there is no change in consumption when the number of hours changes ($u_{ch} = 0$), the Frisch elasticity is given by $(u_h/u_{hh})(1/h)$. That happens, for example, when the utility function is separable in labor and consumption, as was the case in many previous studies.

The assumption that labor and consumption are additively separable was very important to many previous empirical estimates, because it implies that the Frisch elasticity can be calculated directly using the parameters of the utility function. In particular, Altonji (1986), Heckman and MaCurdy (1982), MaCurdy (1981), and others use the following functional form:

$$u(c,h) = \frac{B_c}{1+B_c} V_{ck} C_k^{1+\frac{1}{B_c}} - V_{nk} N_k^{1+\frac{1}{B_n}} \tag{8}$$

where C is consumption, V_{ck} and V_{nk} are preference parameters, N is labor supply, B_c is the elasticity of intertemporal substitution of consumption (EIS consumption). In equation (8), the Frisch elasticity of labor supply is given directly by B_n, which represents at the same time the EIS of labor.[13] That condition does not hold in utility functions for which labor and consumption are nonseparable, as will be clear in the next section.

[13]This is a particular case: The difference between the EIS and the Frisch elasticity is that the EIS does not hold the marginal utility of wealth constant.

3.2 Quantifying the Frisch Elasticity, Not Accounting for Borrowing Constraints and Uncertainty

To quantify the Frisch elasticity without considering borrowing constraints or uncertainty, we use equation (6) and assume in the model the following utility function:

$$u(c_t, h_t) = \frac{\left\{\left[\left(1+\frac{n_t}{2}\right)^{-\varsigma} c_t\right]^{\alpha} \left[h_t^{max} - h_t\right]^{1-\alpha}\right\}^{1-\gamma}}{1-\gamma} \tag{9}$$

here γ is the coefficient of relative risk aversion, α is the share parameter for consumption, n_t is the number of dependent children, ς is the consumption adjustment parameter for children, and h_t^{max} is the maximum number of working hours.

Using the model's utility function in equation (9) and applying the general theoretical formula for the Frisch elasticity described in equation (6) we arrive at an expression for the Frisch elasticity of labor supply without considering borrowing constraints or uncertainty:

$$\left.\frac{dh}{dw}\frac{w}{h}\right|_{\lambda} = \frac{h_t^{max} - h_t}{h_t}\left\{\frac{1-\alpha(1-\gamma)}{\gamma}\right\} \tag{10}$$

To calculate the Frisch elasticity in (equation 10), we need measures of the maximum number of hours and parameters for the share of consumption in the utility function and for the coefficient of relative risk aversion. Those parameters are chosen such that the average hours per household match the data for the United States, as well as the capital-to-labor ratio. Following Nishiyama and Smetters (2005), we use $h^{average} = 3414$ and $h^{max} = 8760$, representing 2 persons working 12 hours each, $\gamma = 2.0$, and $\alpha = 0.465$. We calibrate the value of α (the share of consumption in the utility function) to get consistency among the value of the risk aversion γ and the value of maximum and average hours. The intertemporal Frisch elasticity of labor supply obtained in this fashion has a value of 1.14.

The elasticity in equation (10) does not consider borrowing constraints or uncertainty as they appear in the household problem (equations [1] and [2]). But uninsurable wage risk and borrowing constraints are crucial features of the stochastic OLG model used in this paper. The borrowing constraints reduce the response of hours of work to a change in wages, and the presence of uncertainty introduces a different type of intertemporal labor supply response, which we measure later as the response of hours of work to unanticipated changes in wages.[14]

The uncertainty is implemented as idiosyncratic shocks to earnings. Bad wage shocks can slam households onto borrowing constraints that effectively break the intertemporal links across periods in the model. When the constraints bind, households in the model effectively become static, and the lifetime period is broken up into a series of shorter periods. Therefore, the estimate using equation (10) would not correspond to observable quantities in conditions of uncertainty and borrowing constraints, so we need to use a different estimating procedure. We do so in the next section by generating a synthetic panel and using established empirical procedures to estimate the two types of intertemporal elasticity (response to anticipated and to unanticipated changes in wages).

3.3 Empirical Estimates of the Labor Supply Responses to Anticipated (Frisch) and Unanticipated Changes in Wages from a Simulated Panel

There is no reason for the Frisch elasticity estimated in a simulated dynamic economy to be equal to the values analytically derived from the utility function. In the model we use in this paper, agents live through many periods and face borrowing constraints, uncertain time of death, and idiosyncratic shocks to working ability. Those factors are expected to attenuate the labor supply response to policy changes, but quantifying that attenuation is not easy. The implied Frisch elasticity cannot be calculated analytically from the model's parameters and aggregate results.

To overcome those difficulties, we use a general method that can be summarized in two steps. First, we generate a simulated panel with the OLG model (The method is

[14]This type of intertemporal labor supply elasticity exists only if uncertainty is present, because idiosyncratic shocks cannot be anticipated by the households, and they trigger this specific response.

described in the appendix.) Second, once the panel is generated, we estimate two types of intertemporal elasticities using an empirical methodology that has already been used with real data. The elasticities we estimate are the labor supply response to anticipated changes in wages, holding the marginal utility of wealth constant (Frisch elasticity), and the labor supply response to unanticipated changes in wages. We use Pistaferri (2003) approach because it accounts directly for variances and shocks, and it controls for changes in the marginal utility of wealth. Given that we can generate exact data and calculate expectations from the model, our estimate is free of measurement errors that arise when people misjudge their expected income or consumption or misreport current values.

As is clear in equation (7), it is impossible to calculate directly the Frisch elasticity accounting for borrowing constraints and uncertainty, so we need to obtain an equation that is feasible to estimate. To obtain such equation, we follow Pistaferri (2003) and obtain from equation (6) a log-linear specification:

$$\ln h_{it} = Z_{it}'\alpha + \eta \ln w_{it} + (\eta + \varphi) \ln \lambda_{it} \tag{11}$$

where η represents the Frisch elasticity of labor supply, φ is the elasticity between consumption and labor, and $(\eta+\varphi)$ is the elasticity of hours with respect to the marginal utility of wealth. Empirical estimates of the Frisch elasticity obtained using equation (11) will account for borrowing constraints and uncertainty, independent of the econometric problems present in the methodology; calculations of the Frisch elasticity of labor supply using equation (6) will not.

Equation (11) can be used to analyze the effect of borrowing constraints in the Frisch elasticity of labor supply, in the same spirit as Domeij and Floden (2006). We can express that equation in a differenced version and use equation (5) to arrive at an approximated expression for the change in the marginal utility of wealth: $\Delta \ln \lambda_{it} = (\rho - E_{t-1} r_t) + \ln[1 - (\mu/\lambda)]$. This way we arrive at the following new version of equation

(11):

$$\Delta \ln h_{it} = \Delta Z_{it}\alpha + \eta E_{t-1}\Delta \ln w_{it} + (\eta + \varphi)[(\rho - E_{t-1}r_t) + \ln(1 - (\mu/\lambda_{it}))] + \varepsilon_t \quad (12)$$

Through the term $\ln(1-(\mu/\lambda_{it}))$, equation (12) shows the negative effect of borrowing constraints μ in the Frisch elasticity of labor supply η. It can be argued that the expected change in wages is negatively correlated with the marginal utility of wealth because the marginal utility of wealth decreases by the effect of the concavity of the utility function. Then, if the marginal utility of wealth is not controlled directly, we would have a negative correlation between the terms $E_{t-1}\Delta \ln w_{it}$ and $\ln(1 - (\mu/\lambda_{it}))$, inducing a negative bias in the parameter η, the Frisch elasticity of labor supply.

Researchers face three important problems when they try to estimate equations like equation (12): First, neither the marginal utility of wealth nor the Lagrange multiplier associated with the borrowing constraints is observed by the econometrician. Second, the change in expected log of wages is not observed. Third, the change in expected log of wages is likely to be correlated with the marginal utility of wealth or with the change in tastes for leisure, Z_{it}. A widely applied option is to use instrumental variables techniques for that term. The problem, as Keane (2006) points out, is that it is difficult to develop a good instrument to predict the expected change in log wages, and the commonly used variables, such as age and education, could be associated with preference shifts or with elements present in the error term, again introducing bias in the estimates.

From here, we need to make some assumptions to arrive at closed-form solutions and calculate the intertemporal elasticities of labor supply. The Pistaferri (2003) method uses the usual consumer problem, as described in equation (3), with the following assumptions:

1. The wage process is parameterized as

$$\ln w_{it} = \Delta X'_{it}\sigma + \ln w_{it-1} + \zeta_{it} \quad (13)$$

where X is a vector of deterministic variables affecting wages (age, education) and ζ is the innovation in the wage rate $\zeta = \ln w_{it} - E_{t-1} \ln w_{it-1}$.

2. The log of the marginal utility of wealth (λ) is parameterized as

$$(\eta + \varphi) \ln \lambda_{it} = \gamma_a + \sum_{\tau=0}^{T} \gamma_\tau E_t(\ln w_{it+\tau}) + \nu_i \tag{14}$$

where η is the Frisch elasticity, φ is the response of consumption to changes in labor, and $(\eta + \varphi)$ is the elasticity of hours with respect to the marginal utility of wealth. From here, we can observe that $\Gamma = \sum_{\tau=0}^{T} \gamma_\tau$ is the wealth effect of a parametric permanent shift in the wage profile.

Using assumptions (13) and (14) and the general theoretical formula for the Frisch elasticity described in equation (6) we can develop the following equation:

$$\Delta \ln h_{it} = \Delta Z_{it}\alpha + \eta E_{t-1}\Delta \ln w_{it} - (\eta + \varphi)(E_{t-1}r_t - \rho) + \upsilon Var_{t-1}(\zeta_{it}) + (\eta + \Gamma)\zeta_{it} \tag{15}$$

η represents the value of the Frisch elasticity, which measures the change in labor supply in reaction to expected changes in wages, holding the marginal utility of wealth constant. The other parameters of interest are υ, which measures the change in labor supply in response to wage risk, and $(\eta + \Gamma)$, which measures the change in labor supply to unexpected changes in wages.[15] The variables generated using the model are hours (h), expected change in the log of wages ($E_{t-1}\Delta \ln w_{it}$), and innovation in the wage rate $\zeta = \ln w_{it} - E_{t-1} \ln w_{it-1}$. In the simulated panel, the interest rate and the discount rate are the same for all households, and the vector of preference shifts does not change. All those terms are reflected in the constant term. Because our data were generated by simulation from a life-cycle model, we have the advantage over most empirical studies of having the precise variables we need to estimate the parameters of interest and no

[15] The literature refers to anticipated changes in wages as "evolutionary," and to unanticipated changes as "parametric."

measurement error problems to deal with. The equation to be estimated is the following:

$$\Delta \ln h_{it} = \alpha' + \eta E_{t-1} \Delta \ln w_{it} + \upsilon Var_{t-1}(\zeta_{it}) + (\eta + \Gamma)\zeta_{it} + \varepsilon_t \tag{16}$$

Note that in equation (16) the presence of uncertainty about future wages is independent from the coefficient of the Frisch elasticity, but failing to account for this uncertainty in the estimation introduces bias to the results (omitted variable bias). Accounting for this uncertainty allows one to identify the estimate of the labor supply response to unanticipated changes in wages.

Table 2 shows the results of the estimation after applying equation (16) to the simulated data (and using a fixed-effect specification). We arrive at an estimate of the Frisch elasticity of labor supply of 0.327 (standard deviation = 0.003), about one-fourth of the value when not considering borrowing constraints and wage uncertainty. Our Frisch elasticity estimate is within the range of 0.1-0.45 given by MaCurdy (1981), close to the central value of 0.2 in Pencavel's 1986 survey, within the range of 0.24-0.79 given by Angrist (2005), but smaller than many recent estimates.

Our Frisch elasticity estimate is just over half of Pistaferri's (2003) estimate. The difference could have any one of several causes: First, because we generate data on wages and expected wages directly, we isolate the estimates from unobserved taste shifts and we can estimate the Frisch elasticity using wages instead of earnings. Second, our data do not have measurement error. And third, we can control the macroeconomic conditions, where unemployment and other demand-side variables can alter the labor supply responses.

The labor supply response in the model to unanticipated wage changes is very small (0.0007). That value is below previous empirical estimates. For example, MaCurdy (1981) has 0.08 and Pistaferri (2003) has 0.51. At the same time, our estimate implies a wealth effect of -0.3264, which is similar to the Frisch elasticity estimate and significantly ameliorates the intertemporal substitution effect in response to wage changes. The wealth effect estimate in our model is somewhat higher than Pistaferri's estimate of -0.2

Table 2. Frisch elasticity estimates using the simulated panel

Parameter	Estimate (standard error)	95% Confidence Interval
η (Frisch)	0.3271 (0.0029)	[0.3213, 0.3328]
υ (Wage risk)	0.0017 (0.0000)	[0.0017, 0.0018]
$\eta + \Gamma$ (Unexpected change in wages)	0.0007 (0.0002)	[0.0003, 0.0010]
α'	-0.2225 (0.0052)	[-0.2326, -0.2124]

and significantly higher than MaCurdy's estimate of -0.07. However, our estimate is within 1.5 standard errors of Pistaferri estimate.

3.4 Quantifying the Omitted Variable Bias in the Frisch Elasticity Estimation

If the econometrician does not include measures of wage risk or unexpected changes in wages-as is the case for many empirical studies-the Frisch elasticity estimates exhibit a significant downward bias. To quantify the omitted variable bias in the Frisch elasticity of labor supply, we can run the specification in equation (16) without controlling for wage risk or for unexpected changes in wages.

If measures of wage risk are not included in the regression, the estimate of the Frisch elasticity drops from 0.33 to 0.24 (see Table 3). That estimate suggests an omitted variable bias of around 30 percent. By contrast, the Frisch elasticity estimate does not change much when unexpected change in wages alone is not included (Table 4).

However, if both the wage risk measure and the measure of unexpected change in wages is excluded, the estimate of the Frisch elasticity drops to 0.04 (Table 5). That result suggest that the omitted variable bias in regressions that exclude those control variables could be substantial and may explain why several early studies found small estimates of the Frisch elasticity.

Table 3. Frisch elasticity estimates using the simulated panel
not controlling for wage risk measures

Parameter	Estimate (standard error)	95% Confidence Interval
η (Frisch)	0.2376 (0.0025)	[0.2328, 0.2425]
$\eta + \Gamma$ (Unexpected change in wages)	0.0020 (0.0002)	[0.0016, 0.0023]
α'	-0.0297 (0.0021)	[-0.0255, -0.0339]

Table 4. Frisch elasticity estimates using the simulated panel
not controlling for unexpected changes in wages

Parameter	Estimate (standard error)	95% Confidence Interval
η (Frisch)	0.3269 (0.0029)	[0.3211, 0.3326]
υ (Wage risk)	0.0018 (0.0000)	[0.0017, 0.0018]
α'	-0.2241 (0.0051)	[-0.2341, -0.2140]

Indeed, most empirical studies do not include measures of wage risk or unexpected changes in wages, because expectations are hard to measure. The exception is the study by Pistaferri (2003), who extracted individual expectation of earnings growth from the Bank of Italy Survey of Households' Income and Wealth. He modifies further the specification in equation (16) to adjust for the fact that expectations consider earnings rather than wages.

Several important strengths of our approach enable us to measure the omitted variable bias. First, we generate simulated data on wages and expected wages, thus isolating the estimates from unobserved taste shifts. Second, our simulated data have no measurement error. Third, we can control the macroeconomic conditions under which unemployment and other demand-side variables can alter the labor supply responses. In addition, we can run several experiments that introduce different levels of transitory wage innovations

Table 5. Frisch elasticity estimates using the simulated panel not controlling for wage risk measures or unexpected changes in wages

Parameter	Estimate (standard error)	95% Confidence Interval
η (Frisch)	0.0411 (0.0009)	[0.0393, 0.0429]
α'	-0.0096 (0.0022)	[-0.0052, -0.0140]

to arrive at the same point estimate for the Frisch elasticity.

4 Conclusions

In this paper, we use a stochastic OLG model with incomplete markets and Pistaferri's (2003) empirical method to estimate the labor supply response to anticipated changes in wages, holding the marginal utility of wealth constant (Frisch elasticity), and to unanticipated change in wages. We conduct an experiment by using a simulated panel from the model, in which we can exactly observe expected wages and control for demand-side variables and uncertainty, but which does not present issues with measurement error.

Once borrowing constraints and uncertainty about future wages are considered, our Frisch elasticity estimate (0.33) is within the range of what appears in the microeconomic empirical literature. If we exclude those factors, the estimate is almost 4 times larger (1.17). The rationale is that, without borrowing constraints, households respond to a bad wage shock by reducing labor and increasing borrowing so as to smooth consumption over time.

The labor supply response to unanticipated wage changes, which also is relevant for the analysis of tax and benefit reforms, is very small (0.0007) because of large wealth effects in the OLG model. In principle, both intertemporal substitution and wealth effects are relevant for determining the labor supply response to tax or benefit changes. A wage increase generates an intertemporal substitution effect which tends to increase the labor supply, but the same wage increase generates a wealth effect and a labor supply reduction. Because reforms often differ at the extent to which they alter these two effects, it is important to explicitly consider both of them in the analysis.

Finally, we quantify the bias-attributed to omitted variables-that occurs in many empirical estimates because of lack of data about wage risk and unexpected changes in wages. That bias is 88% if both variables are absent and 30% if wage risk measures are excluded.

Appendix. Generation of the Simulated Panel

The simulated panel is generated in two steps. First, we solve for the optimal solution to the consumer problem and obtain the policy functions for each point in the discretized state space. The policy functions define a rule that maps the space into recursive decisions, which are located in the continuous space defined by the boundaries of the state variables (not limited to the grid points). For every combination of age, wealth, and work ability, the policy function gives the optimal consumption and work hours. The policy function will, in general, map a state S at time t, which corresponds to a grid point in the discretized space, to a state S' at time $t+1$, in the continuous-state space. When aggregating the results, a measure of households is assigned to each grid point by reverse interpolation: Each of the grid points surrounding state S' is assigned such a portion of the surviving households from state S that S' could be interpolated as a combination of solutions limited to grid points.

The second step in panel generation is to simulate a time series of the household's optimal decisions. In every period, those decisions are linked by the law of motion from the policy function, so that a household makes decisions at time $t+1$ depending on its state at time t.

To generate the time series, we simulate the stochastic variables of the model (the ability shocks) as the same stochastic process defined in the solution of the model. In each period t, the household is at a grid point S of the state space, determined by the random numbers generated so far. The policy function is applied to that state, and it maps the household's wealth to its value in the next period, resulting in a state S', which generally does not correspond to a grid point. At that stage, a process equivalent to the assigning of measure to the surrounding grid points (described above) is applied. The procedure assigns to the household one of the surrounding grid points, with the respective probability of each grid point proportional to the measure assigned to the grid point during the solution of the model. This step requires the generation of an additional random number or, more generally, a vector of as many random numbers as there are continuous dimensions in the state space. The household thus ends up in a state S'' given by a grid point in the discretized space, and the procedure can be repeated for the next period.[16]

The method described before generates the time series of length T (T is the number of stochastic realizations for the income process—generally equal to the maximum length of a career) for one individual. The same procedure is repeated N times to generate N households, and thus there is a panel of dimension $N \times T$. We generate 10,000 households that may live up to 109 years.

[16]There is another option for generating simulated households, not pursued here. That option includes generating a policy rule for each stochastic income process point and a recursive realization of state variables; that is, to apply the solution procedure instead of applying the obtained policy rules. In other words, the state S', resulting from the last period's policy, would not be further shocked to fall onto a grid point S'', but the period $t+1$ policy would be computed directly for the state S'.

References

Altonji, Joseph G. (1986) “Intertemporal substitution in labor supply: Evidence from micro data,” *Journal of Political Economy*, vol. 94, June, S176-S215.

Angrist, Joshua. (1991) “Grouped data estimation and testing in simple labor supply models,” *Journal of Econometrics*, vol. 47, March.

Aiyagari, S. Rao. (1994) “Uninsured Idiosyncratic Risk and Aggregate Saving,” The Quarterly Journal of Economics, vol. 109(3), August, 659-84

Blundell, Richard, Pierre-Andre Chiappori, Thierry Magnac and Costas Meghir. (2007) “Collective labour supply: Heterogeneity and non-Participation,” *Review of Economic Studies*, vol. 74(2), 417-445.

Blundell, Richard, and Thomas MaCurdy. (1999) “Labor supply: A review of alternative approaches.” In *Handbook of Labor Economics 3,* Amsterdam: Elsevier North Holland.

Blundell, Richard, and Thomas MaCurdy.(2008) “Labor supply.” In *The New Palgrave Dictionary of Economics*, Palgrave MacMillian.

Blundell, Richard, Costas Meghir, and Pedro Neves. (1993) “Labor supply: An intertemporal substitution,” *Journal of Econometrics*, vol. 59, 137-160.

Card, David. (1994) “Intertemporal labor supply: An assessment.” In Christopher Sims, editor, *Advances in Econometrics,* Sixth World Congress. New York: Cambridge University Press.

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*, vol. 47, 1-27.

Chiappori, Pierre-Andre. (1988) “Rational household labor supply,” *Econometrica*, vol. 56(1), January, 63-90.

Chiappori, Pierre-Andre. (1992) "Collective labor supply and welfare," *Journal of Political Economy*, vol. 100(3), June, 437-67.

Congressional Budget Office. (2008) "Analysis of the president's budgetary proposals for fiscal year 2009."

Domeij, David, and Martin Floden. (2006) "The labor supply elasticity and borrowing constraints: Why estimates are biased," *Review of Economic Dynamics*, vol. 9, April, 242-262.

Fehr, Ernst, and Lorenz Götte. (2007) "Do workers work more if wages are high? Evidence from a randomized field experiment," *American Economic Review,* vol. 97, March, 298-317.

French, Eric. (2004) "The labor supply response to (mismeasured but) predictable wage changes", *Review of Economics and Statistics*, vol. 86, May, 602-613.

Gourinchas, Pierre-Olivier, and Jonathan Parker. (2002) "Consumption over the life cycle," *Econometrica*, vol. 70, January, 47-89.

Heckman, James, and Thomas MaCurdy. (1982) "A life cycle model of female labor supply," *Review of Economic Studies*, vol. 49, January, 659-660.

Heckman, James, and Thomas MaCurdy. (1980) "Corrigendum on a life cycle model of female labor supply," *Review of Economic Studies*, vol. 47, January, 47-74.

Heckman, James. (1979) "Sample selection bias as an specification error," *Econometrica*, vol. 47, 153-161.

Heckman, James. (1993) "What has been learned about labor supply in the past twenty years," *American Economic Review (Papers and Proceedings)*, vol. 83, 116-121.

Huggett, Mark. (1996) "Wealth distribution in life-cycle economies," *Journal of Monetary Economics*, vol. 38(3), December, 469-494.

Imai, Susumu and Michael Keane. (2004) "Intertemporal Labor Supply and Human Capital Accumulation," *International Economic Review*, vol. 45, 601-641.

Keane, Michael. (2006) “Structural vs. Atheoretical Approaches To Econometrics” forthcoming, *Journal of Econometrics*

Laitner, John, and Dan Silverman. (2005). “Estimating life cycle parameters from consumption behavior at retirement,” Working Papers 099, University of Michigan, Michigan Retirement Research Center.

Lee, Chul-In. (2001) “Finite sample bias in IV estimation of intertemporal labor supply models: Is the intertemporal substitution elasticity really small?” *The Review of Economics and Statistics*, vol. 83, November, 638-646.

MaCurdy, Thomas E. (1981) “An empirical model of labor supply in a life cycle setting,” *Journal of Political Economy*, vol. 77, December, 1059-1085.

Nishiyama, Shinichi. (2002) “Bequests, Inter Vivos Transfers, and Wealth Distribution,” *Review of Economic Dynamics*, vol. 5(4), 892-931.

Nishiyama, Shinichi, and Kent Smetters, (2005) “Consumption Taxes and Economic Efficiency with Idiosyncratic Wage Shocks,” *Journal of Political Economy*, vol. 113(5), October, 1088-1115.

Pencavel, John. (1986) “Labor supply of men,” In Orley Ashenfelter and R. Layard, editors, *Handbook of Labor Economics*, Amsterdam: Elsevier North Holland.

Pistaferri, Luigi. (2003) “Anticipated and unanticipated wage changes, wage risk, and intertemporal labor supply,” *Journal of Labor Economics*, vol. 21, July, 729-728.

Prescott, Edward. (2006) “Nobel lecture: The transformation of macroeconomic policy and research,” *Journal of Political Economy*, vol. 114, April, 203-235

Rogerson, Richard, and Johanna Wallenius (2007). “Micro and macro elasticities in a life cycle model with taxes,” NBER Working Paper 13017.

Ziliak, James, and Thomas Kniesner. (2005) “The effect of income taxation on consumption and labor supply,” *Journal of Labor Economics*, vol. 23, October, 769-796.