Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply

Luigi Pistaferri, *Stanford University, Center for Economic Policy Research, and Stanford Institute for Economic Policy Research*

In this article, I estimate how labor supply responds to anticipated wage growth, unanticipated wage growth, and wage risk using the 1989–93 panel section of the Bank of Italy Survey of Households' Income and Wealth (SHIW), which collects individual expectations of future wages. The use of subjective expectations has several advantages. First, they provide information on the evolution and riskiness of future wages that the econometrician may never hope to observe. Finally, this avoids the need for specifying instruments for the growth rate of wages. Finally, forecast errors can be directly controlled for, thus avoiding inconsistency in short panels.

I. Introduction

A long-standing question in labor economics is whether and to what extent individual labor supply responds to anticipated wage changes (also known as evolutionary wage changes). This effect is measured by the intertemporal elasticity of substitution, a popular parameter among macroe-

I am grateful to Tullio Jappelli, Alan Krueger, Tom MaCurdy, Charles Manski, John Pencavel, Steve Pischke, and seminar participants at Stanford University and the 2000 CEPR Summer Symposium in Labor Economics for comments. Financial support from the Taube Faculty Research Fund at the Stanford Institute for Economic Policy Research (SIEPR) is gratefully acknowledged. The views expressed in this article are mine.

[*Journal of Labor Economics*, 2003, vol. 21, no. 3]

0734-306X/2003/2103-0009$10.00

conomists. The available evidence on the magnitude of this parameter is, at best, mixed. Several studies conducted in the 1980s on samples of prime-age males in the United States (MaCurdy 1981; Altonji 1986) concluded that the magnitude of the coefficient was likely negligible or playing only a second-order effect.[1] More recent studies, such as two by Mulligan (1995, 1998), have reached different conclusions and suggested the possibility of a large impact of anticipated wage changes on changes in individual labor supply. The evidence presented in these papers exploits particular episodes (such as World War II or the Exxon-Valdez oil spill accident) or focuses on particular groups (such as agricultural workers or stadium vendors), where perhaps a more convincing source of anticipated variation in wages can be identified.

A related, yet less explored, issue, is whether and to what extent individual labor supply responds to unanticipated wage changes (also known as parametric wage changes). Attempts to estimate the impact of those changes on labor supply include work by MaCurdy (1981) and Bover (1989). These studies typically construct forecasts of future wages by making explicit hypotheses concerning the evolution over time of the individual wage profile and assuming that the econometrician's forecasts are unbiased. However, as outlined by Card (1991, p. 36), "it is possible that individuals have better information with which to forecast future wages than is unavailable to an outside analyst. In this case, wage innovations in [a] statistical model . . . do not necessarily represent new information."

A question that has not received attention so far is whether and to what extent individual labor supply responds to wage changes that increase the variability of the lifetime wage profile, that is, whether labor supply exhibits any precautionary behavior. In analogy with models of precautionary saving, a riskier wage profile should delay the consumption of leisure (if leisure is a normal good), and thus it should negatively affect the growth rate of labor supply (Low 1999). In particular, individuals work longer hours when they are young so as to accumulate precautionary balances, rather than provide all precautionary savings through sacrificing consumption, as in the standard model. Moreover, accumulated wealth increases both consumption and leisure when the saver becomes old, that is, it leads to a reduction in labor supply later in life.

These three issues represent the motivation for this article. First, the resurgence of interest on the magnitude of intertemporal substitution in labor supply warrants further research; this is an important structural parameter for policy analysis. Lucas and Rapping (1969) suggested that a large intertemporal substitution in labor supply could potentially explain

[1] Similar evidence was available for the United Kingdom (Browning, Deaton, and Irish 1985).

business cycle fluctuations. Second, it is important to understand whether unanticipated changes in wages play a role in explaining fluctuations in labor supply over the life cycle and over the business cycle. Card (1991, p. 11) noticed that the literature concentrates "on one aspect of intertemporal hours variation—the response to wage growth along a known life cycle trajectory—and . . . ignore[s] another, namely, the response to unanticipated wage innovations." On the same point, Mulligan (1998, p. 47) notes that "previous econometric studies of the [intertemporal substitution] hypothesis which used panel micro data [are] unable to distinguish anticipated wage changes from those that [are] unanticipated." Finally, one would like to assess whether precautionary labor supply is important or negligible, as implicitly assumed thus far in the empirical literature. As reported by Mulligan (1998, p. 47), "there is no empirical evidence that precautionary motives for delaying leisure are important" (or otherwise).

To address these questions, one would ideally use a panel of observations on the entire life cycle of hours and wages and other demographic information. Unfortunately, this ideal data set does not exist in the United States or elsewhere. Applied researchers have to sidestep this problem by using either the available limited panel data (such as the Panel Study of Income Dynamics [PSID]) or by constructing pseudopanels, as in Ghez and Becker (1975) and Browning et al. (1985). The latter approach, however, eliminates the individual heterogeneity that is likely to be most responsible for fluctuations in hours.

In the the current article, I take another route. I make use of subjective expectations collected at the individual level. It is often argued, but rarely put into practice, that the nature of the variables one wants to study in the intertemporal consumption or labor literature (the expectation, the innovation, and the expected variability of wage rates, say) calls for the use of individual expectation data. I test for the effect of anticipated and unanticipated wage growth and expected wage risk on intertemporal labor supply using the 1989-93 panel section of the Bank of Italy Survey of Households' Income and Wealth (SHIW). This data set has the unique advantage of collecting individual-based quantitative expectations of future wage changes. The decomposition between anticipated and unanticipated wage growth is exact because individual subjective expectations are used to construct the anticipated component and the difference between the actual realization and the subjective expectation is used to construct the unanticipated component. The same data also allow me to construct a measure of anticipated variability in the lifetime profile of wages and therefore to test for precautionary labor supply. To my knowledge, this is the first attempt to estimate the impact of wage risk on intertemporal labor supply using microeconomic data (for evidence on aggregate data, see Blau and Grossberg [1989]).

The article proceeds as follows. Section II reviews the standard theory of intertemporal labor supply, while Section III discusses the main features of the data and presents descriptive evidence on the intertemporal variability in labor supply. Here I also show how subjective expectations data can be used to disentangle anticipated from unanticipated wage growth, and I demonstrate how to compute a measure of overall expected variability in the wage profile. Section IV presents the main results and considers some extensions. Section V concludes.

II. Wage Changes, Wage Risk, and Intertemporal Labor Supply

In this section, I derive the Euler equation for labor supply and show how the structural parameters of interest can be identified. The derivation follows MaCurdy (1981) and Card (1991). The microeconomic unit of interest is a married, prime-age male. This individual solves the following problem:

$$\max E_t \sum_{\tau=0}^{T-t} (1+\delta)^{-\tau} u\left(c_{it+\tau}, h_{it+\tau}, \mathbf{Z}_{\mathrm{it}+\tau}\right)$$

subject to the intertemporal budget constraint

$$a_{it+1} = (1 + r_{t+1})(a_{it} + w_{it}h_{it} - c_{it}), \tag{1}$$

where $E_t(.) = E(.|\Omega_{it})$ is the subjective expectation that conditions on the individual information set Ω_{it} at time t, δ is the intertemporal discount rate, c is consumption, h is hours of work, $\mathbf{Z}$ is a vector of preference shifts, a represents assets, r is the real interest rate, and w is the real hourly wage. I assume that the preferences and the budget constraint are intertemporally separable and that there are no restrictions on borrowing, saving, and labor supply (i.e., the solution is interior).[2] However, I do not assume that preferences are intratemporally additive. The first-order conditions for this problem are

$$\frac{\partial u\left(c_{it}, h_{it}, \mathbf{Z}_{\mathrm{it}}\right)}{\partial c_{it}} = \lambda_{it} \tag{2}$$

$$-\frac{\partial u\left(c_{it}, h_{it}, \mathbf{Z}_{\mathrm{it}}\right)}{\partial h_{it}} = \lambda_{it} w_{it} \tag{3}$$

$$\lambda_{it} = E_t \frac{1 + r_{t+1}}{1+\delta} \lambda_{it+1}, \tag{4}$$

along with the intertemporal budget constraint (1), and where λ is the

[2] In the period of observation (1989–93), there were no changes in the tax code. The measures of wage, earnings, and expected earnings considered in this article are all after taxes and social security contributions.

Lagrange multiplier. Equations (2) and (3) can be solved for consumption and hours of work to obtain Frisch (or λ-constant) demand functions; in this context, consumption and hours of work will depend on the wage rate, the marginal utility of wealth λ_{it}, and preference shifts.[3] The main feature of Frisch demand functions is that they keep constant the marginal utility of wealth. In particular, for hours of work, one can write the approximated log-linear specification (Altonji 1986) as follows:

$$\ln h_{it} \simeq \mathbf{Z}'_{\mathrm{it}}\alpha + \eta \ln w_{it} + (\eta + \varphi)\ln \lambda_{it}. \tag{5}$$

The term $\ln \lambda_{it}$ captures all we need to know about future wages and interest rates. However, unless restrictive assumptions are made, it very rarely has a workable closed form. I comment later on this point. The parameter η is the intertemporal elasticity of substitution for labor supply. It measures the impact of an anticipated change in wages (an evolutionary change) on labor supply while keeping constant the marginal utility of wealth. The parameter $(\eta + \varphi)$ is the elasticity of hours with respect to the marginal utility of wealth. If preferences are strictly concave and if leisure and consumption are normal goods, the following sign restrictions hold: $\eta > 0$ and $\eta + \varphi > 0$. As for φ, its sign will be negative if consumption and leisure are complements and positive if they are substitutes in utility. If consumption is independent of wages, which occurs if preferences are intratemporally additive, then $\varphi = 0$.[4]

One can combine (4) and (5) to obtain the (approximated) Euler equation for labor supply:

$$\begin{aligned} \Delta \ln h_{it} \simeq\ & \Delta \mathbf{Z}'_{\mathrm{it}}\alpha + \eta\Delta \ln w_{it} - (\eta + \varphi)(E_{t-1}r_t - \delta) \\ & + (\eta + \varphi)(\ln \lambda_{it} - E_{t-1}\ln \lambda_{it}) \\ & - (\eta + \varphi)\ln E_{t-1}(e^{\ln \lambda_{it} - E_{t-1}\ln \lambda_{it}}), \end{aligned} \tag{6}$$

where $(\ln \lambda_{it} - E_{t-1}\ln \lambda_{it})$ is the innovation in the logarithm of the marginal utility of wealth. In previous studies, the term $\ln E_{t-1}(e^{\ln \lambda_{it} - E_{t-1}\ln \lambda_{it}})$

[3] The use of both consumption and labor supply data goes beyond the scope of this article and is left to future research.

[4] Defining with u_i the derivative of $u(.)$ with respect to the argument i ($i = h, c$) and u_{ij} the derivative of u_i with respect to the argument j ($j = h, c$), it can be shown that

$$\eta = \frac{u_h u_{cc}}{u_{cc}u_{hh} - u_{ch}u_{hc}}\frac{1}{h},$$

$$\varphi = -\frac{u_{hc}u_c}{u_{cc}u_{hh} - u_{ch}u_{hc}}\frac{1}{h}.$$

If $u_{ch} = u_{hc} = 0$, then $\varphi = 0$ and $\eta = (u_h/u_{hh})(1/h)$.

is part of the intercept; I allow it instead to vary across individuals in order to capture heterogeneity in wage uncertainty.

To obtain an empirically tractable specification, I make two assumptions similar to those made in previous empirical work. The first assumption is that the wage rate follows a geometric martingale process with drift, that is, that

$$\ln w_{it} = \Delta X'_{it}\sigma + \ln w_{it-1} + \zeta_{it}, \tag{7}$$

where $\mathbf{X}$ is a vector of variables affecting the deterministic component of wage growth (such as age, education, and other information available to the individual but not to the econometrician) and ζ_{it} is the innovation in the wage rate, that is, $\zeta_{it} = (\ln w_{it} - E_{t-1}\ln w_{it})$. It is fair to say that this assumption is consistent with aggregate data, but less so with microeconomic data, where instead the evidence is more favorable to ζ_{it} following an MA(1) process.[5] As a partial justification, one can note that a transitory unanticipated component has a negligible impact on labor supply because it affects $\ln\lambda_{it}$ only marginally. I address this point in the empirical section below.

The second assumption concerns the approximated form of the log of the marginal utility of wealth, and it is the same as in MaCurdy (1981):

$$(\eta + \varphi)\ln\lambda_{it} = \gamma_a a_{it} + \sum_{\tau=0}^{T-t}\gamma_\tau E_t(\ln w_{it+\tau}) + \nu_i, \tag{8}$$

where ν_i is an approximation error.[6] Note that concavity of the utility function requires that γ_a and γ_τ be negative for all τ. Given the characteristics of the wage process, it follows that the innovation in the logarithm of the marginal utility of wealth can be written as

$$(\eta + \varphi)(\ln\lambda_{it} - E_{t-1}\ln\lambda_{it}) = \sum_{\tau=0}^{T-t}\gamma_\tau(E_t - E_{t-1})\ln w_{it+\tau} = \Gamma\zeta_{it},$$

where $\Gamma = (\sum_{\tau=0}^{T-t}\gamma_\tau)$ is the wealth effect of a parametric permanent shift in the wage profile.[7] Note, finally, that

$$-(\eta + \varphi)\ln E_{t-1}(e^{\ln\lambda_{it} - E_{t-1}\ln\lambda_{it}}) = \vartheta^\zeta \operatorname{Var}_{t-1}(\zeta_{it}), \tag{9}$$

where $\vartheta^\zeta = -\Gamma^2/[2(\eta + \varphi)] < 0$. The equality in (9) holds if the inno-

[5] The MA(1) component reflects transitory shocks and measurement error in the wage variable.

[6] The inclusion of predictable demographic variables does not affect our estimation strategy. Allowing for higher moments of the distribution of wages has little effect on our estimates (results are available from the author on request).

[7] Note that $a_{it} - E_{t-1}a_{it} = 0$ if innovations in the real interest rate are small. Removing this assumption leaves the results qualitatively unchanged. These are available on request from the author.

vation in the marginal utility of wealth is normally distributed or, as an approximation, for small risks. The labor supply equation (6) can then be rewritten (neglecting from now on the approximation sign) as

$$\Delta \ln h_{it} = \Delta \mathbf{Z}'_{\mathrm{it}}\alpha + \eta \Delta \ln w_{it} + \vartheta^{\zeta}\,\mathrm{Var}_{t-1}(\zeta_{it}) - (\eta + \varphi)(E_{t-1}r_t - \delta) + \Gamma\zeta_{it}. \tag{10}$$

To see more clearly the distinction between evolutionary and parametric wage growth effects on labor supply, take the conditional expectation of (10) to yield

$$E_{t-1}\Delta \ln h_{it} = \Delta \mathbf{Z}'_{\mathrm{it}}\alpha + \eta E_{\mathrm{t}-1}\Delta \ln w_{it} + \vartheta^{\zeta}\,\mathrm{Var}_{t-1}(\zeta_{it}) - (\eta + \varphi)(E_{t-1}r_t - \delta), \tag{11}$$

and subtract (10) from (11) to obtain the innovation in labor supply growth:

$$\begin{aligned}\Delta \ln h_{it} - E_{t-1}\Delta \ln h_{it} &= \eta(\Delta \ln w_{it} - E_{t-1}\Delta \ln w_{it}) + \Gamma\zeta_{it} \\ &= (\eta + \Gamma)\zeta_{it}.\end{aligned} \tag{12}$$

Replace (11) in (12) to obtain

$$\Delta \ln h_{it} = \Delta \mathbf{Z}'_{\mathrm{it}}\alpha + \eta E_{\mathrm{t}-1}\Delta \ln w_{it} + \vartheta^{\zeta}\,\mathrm{Var}_{t-1}(\zeta_{it}) - (\eta + \varphi)(E_{t-1}r_t - \delta) + (\eta + \Gamma)\zeta_{it}. \tag{13}$$

Equation (13) predicts that individuals vary their labor supply in response to numerous factors: first, because of shifts in preferences, represented by the vector $\Delta\mathbf{Z}'_{\mathrm{it}}$; second, because of anticipated evolutionary changes in the wage rate, $E_{t-1}\Delta \ln w_{it}$; third, because of a gap between the real interest rate and the intertemporal discount rate, $(E_{t-1}r_t - \delta)$; fourth, because of unexpected (permanent) parametric shifts in the wage rate, ζ_{it}; and, finally, because of a higher dispersion in the conditional distribution of future wage rates, the term $\mathrm{Var}(\zeta_{it})$.

Previous papers in this literature consider the impact of preference shifts, interest rates, and anticipated wage growth, but they neglect the effect of wage innovations and wage risk, which are therefore subsumed in the error term. Moreover, expected wage growth is unobservable in typical data sets, so it must be inferred from instrumental variable techniques. This approach has at least four drawbacks. First, it is hard to find good predictors for wage growth (e.g., a promotion or a change in the wage policy of the firm can be observed by the individual but is unknown to the analyst). Second, the typical predictors used in these papers are age and education, which thus have to play the role of excluded instruments. However, age and education are also the most obvious preference shifts one may think of, and they should, therefore, be included in the vector

Z. In other words, these instruments are invalid if they also play the role of taste shifts. Moreover, the marginal utility of wealth is a function of all characteristics over the entire lifetime, so there are no exclusion restrictions. Third, consistent estimation requires the forecast error to be uncorrelated with the instruments, a property that holds when the time dimension of the panel is long but not necessarily in short panels (Chamberlain 1984). Finally, it has been noticed that the estimates of the elasticities of interest tend to be sensitive to the set of instruments used.

If the marginal utility is kept constant, then $\zeta_{it} = 0$ and all wage changes are anticipated. The reaction of labor supply to expected changes in the wage rate is measured by η, the intertemporal elasticity of substitution. If the marginal utility of wealth is kept constant, a higher expected wage growth suggests that one should work less in the current period and more in the future, and thus $\eta > 0$.

To understand the role of wage innovations, consider equation (13) and suppose that the consumer faces an unexpected positive change in the wage rate, $\zeta_{it} > 0$. By the characteristics of the wage process, all unanticipated parametric wage shifts are permanent. Such a shift induces a revision in $\ln \lambda_{it}$ (which decreases under the maintained assumption that $\Gamma < 0$) and therefore on labor supply. This is then a fully parametric wage change. The elasticity of current labor supply with respect to the parametric wage change, $(\partial \ln h_{it})/(\partial \zeta_{it})$, is therefore $(\eta + \Gamma)$, a parameter that combines intertemporal substitution effects ($\eta > 0$) and wealth effects ($\Gamma < 0$). The sign of $(\eta + \Gamma)$ is thus ambiguous. Blundell and MaCurdy (2000) argue that $(\eta + \Gamma)$ is the main parameter of interest for policy analysis. In fact, since changes in the tax code are usually one-shot reforms that shift the entire wage profile in the current and future periods, the full effect of tax changes on labor supply can only be appreciated by accounting for both anticipated and unanticipated wage changes.[8]

The empirical specification I consider in this article has to take into account the fact that the Bank of Italy Survey of Household Income and Wealth (SHIW) elicits subjective expectations of earnings growth rather than expectations of wage growth. However, I can use the equilibrium condition between hours and wages and the following change of variable, as suggested by MaCurdy (1985). Since $y_{it} = w_{it}h_{it}$, it follows that $\Delta \ln w_{it} = \Delta \ln y_{it} - \Delta \ln h_{it}$. A conditional version of this equality is obtained by replacing realizations with expectations (i.e., $E_{t-1}\Delta \ln w_{it} = E_{t-1}\Delta \ln y_{it} - E_{t-1}\Delta \ln h_{it}$). Moreover, the relationship between the innovation in wages and the innovation in earnings is $\zeta_{it} = (1 + \eta + \Gamma)^{-1}\psi_{it}$,

[8] Of course, the assumptions I make on the statistical form of $\ln \lambda$ and on the characteristics of the wage process are somehow "ad hoc." Nevertheless, both assumptions are very popular choices among labor economists.

with $\psi_{it} = \Delta \ln y_{it} - E_{t-1} \Delta \ln y_{it}$.[9] Replacing the last two equalities in equation (13) and rearranging the terms yields

$$\Delta \ln h_{it} = \Delta \mathbf{Z}'_{it} \frac{\alpha}{1+\eta} + \frac{\eta}{1+\eta} E_{t-1}(\Delta \ln y_{it}) + \vartheta^{\psi} \operatorname{Var}_{t-1}(\psi_{it})$$

$$- \frac{\eta + \varphi}{1+\eta}(E_{t-1} r_t - \delta) + \left(\frac{\eta + \Gamma}{1 + \eta + \Gamma}\right) \psi_{it}, \qquad (14)$$

where $\vartheta^{\psi} = \vartheta^{\zeta}/[(1+\eta+\Gamma)^2(1+\eta)]$. The specification that I test in the empirical analysis is then the following:

$$\Delta \ln h^*_{it} = \Delta \mathbf{Z}'_{it} \beta_0 + \beta_1 E_{t-1}(\Delta \ln y_{it}) + \beta_2 \operatorname{Var}_{t-1}(\Delta \ln y_{it})$$

$$+ \beta_3 (i_t - E_{t-1} \pi_t) + \beta_4 [\Delta \ln y_{it} - E_{t-1}(\Delta \ln y_{it})] + e_{it}, \qquad (15)$$

where h^*_{it} denotes measured annual hours; e_{it} includes a multiplicative measurement error in hours, measurement error in the independent variables, unobserved taste shifts, and approximation errors; $E_{t-1}(\Delta \ln y_{it})$ is the subjective expectation of real earnings growth, which is directly elicited in the SHIW; the term $[\Delta \ln y_{it} - E_{t-1}(\Delta \ln y_{it})]$ denotes the innovation, which is constructed by subtracting the subjective expectation of earnings growth from the actual earnings growth realization; $\operatorname{Var}_{t-1}(\Delta \ln y_{it})$ is the conditional variance of earnings growth; and $(i_t - E_{t-1}\pi_t)$ is the real interest rate, constructed as the difference between the nominal interest rate at time t and inflation expectations, which are also elicited in the survey.

The use of subjective expectations in this context has several advantages. First, it allows consistent estimation of all the parameters of interest, in particular of η and $(\eta + \Gamma)$, by controlling directly for $E_{t-1}(\Delta \ln y_{it})$ and ψ_{it} in estimation. In other words, I estimate an almost exact form of the intertemporal labor supply equation under the assumptions spelled out above. Since the forecast error is directly controlled for in estimation, no short panel problem should arise in this context (Chamberlain 1984). Our estimates will thus be consistent even in the cross-section. Second, controlling for anticipated and unanticipated wage (earnings) growth avoids the need for specifying instruments for wage growth, which is usually hard to predict in microeconomic data, and it avoids imposing untenable exclusion restrictions. Moreover, estimates will not be sensitive to the

[9] From the expression $\Delta \ln w_{it} = \Delta \ln y_{it} - \Delta \ln h_{it}$, the wage growth innovation is the difference between the earnings growth innovation and the hours growth innovation:

$$\zeta_{it} = \psi_{it} - (\eta + \Gamma)\zeta_{it},$$

where I have used the fact that, under my assumptions, the hours growth innovation equals $(\eta + \Gamma)\zeta_{it}$ (see [12]). The presence of unobserved taste shifts may create an endogeneity bias. I comment on this point later.

instruments used, because no IV procedure is needed. Finally, I exploit directly individual information on the evolution and riskiness of future wages (earnings), information that the econometrician may never hope to observe. This allows me to isolate truly expected wage changes from those that are not and provides important identifying information.

III. Data

A. Description

I estimate the intertemporal labor supply equation (15) using the 1989–93 panel section of the Bank of Italy SHIW. This data set contains measures of hours, earnings, consumption, income, and demographic characteristics of the household. The SHIW surveys a representative sample of the Italian resident population. From 1987 through 1995 the survey was conducted every other year and covered about 8,000 households, defined as groups of individuals related by blood, marriage, or adoption and sharing the same dwelling. The most recent survey refers to 1998. Starting in 1989, each SHIW has reinterviewed some households from the previous surveys. The panel component has increased over time: 15% of the sample was reinterviewed in 1989, 27% in 1991, 43% in 1993, and 45% in 1995. Details on sampling, response rates, processing of results, and comparison of survey data with macroeconomic data are provided by Brandolini and Cannari (1994). The 1989–91 panel section of the SHIW includes 2,187 households, and the 1991–93 panel section includes 3,470 households.

For the purpose of this article, the most important feature of the SHIW is that it collected subjective information on future income in both 1989 and 1991. Several surveys contain subjective income expectations, but the surveys vary considerably as to the way expectations are elicited. In the case of the SHIW, in 1989 and 1991 each labor income and pension recipient interviewed was asked to attribute probability weights, summing to 100, to given intervals of inflation and nominal income increases 1 year ahead. The appendix of this article details the wording of the survey questions and describes how the responses were used to construct the subjective mean and variance of future earnings.

A problem with the SHIW data is that they are not available for consecutive years but only at 2-year intervals; moreover, subjective expectations stretch over a single calendar year. The interviews take place between March and September, although income, consumption, and wealth data refer to the previous calendar year.[10] We thus need to assume either that people do not update their information set between the end of 1989 (1991) and the date of the interview or that their updating does not affect

[10] The reason for surveying in May is that previous experience has shown that people report income more accurately when filing the income tax forms, which must be returned by May 31.

subjective expectations of income. This can be a strong assumption if people receive important news about the evolution of their future income between the end of 1989 (1991) and the date of the interview.

So far, the SHIW has been used to test a series of reduced-form propositions in the consumption literature, most notably the impact of precautionary behavior on wealth accumulation (Guiso, Jappelli, and Terlizzese 1992). The data have also been used to address more structural questions, such as estimating the excess sensitivity of consumption to predicted income growth in a standard Euler equation framework (Jappelli and Pistaferri 2000) or estimating the "saving for a rainy day" equation (Pistaferri 2001), that is, how transitory and permanent income shocks affect household saving in the standard framework of the permanent income hypothesis. However, the data have not yet been used to test labor market behavior.

In the next section, I present more structural evidence on the life cycle relationship between wage rates and hours of work. Here I limit myself to a descriptive analysis. The sample used in the empirical analysis below includes only male heads of households who were married, who were between 26 and 59 years of age in 1989, and who reported usable information on the variables of interest, in particular, subjective expectations, earnings, and weekly normal hours. I excluded the self-employed and the multiple job holders. There are 1,461 individual-year observations. Sample statistics are reported in table 1, separately for 1989, 1991, and 1993. The table shows that there is little fluctuation over time in both demographic and economic characteristics. I also report cross-sectional averages of earnings growth expectations and innovations. Both the expectation and the innovation for 1991 earnings are smaller than those for 1993; moreover, they are opposite in sign. Actual earnings growth was negative in both 1991 and 1993, but it was more strongly so (−.018) in 1993, a period in which Italy entered a recession.

Before turning to the empirical analysis, it is useful to describe the distribution of hours and wages. Descriptive evidence on the life cycle pattern of hours and wages is provided in figures 1 and 2, where I plot the age profile of weekly normal hours of work and hourly wages for 12 cohorts whose membership is defined according to the year of birth. Here I use the repeated cross-sections of the SHIW from 1989 to 1998 but retain the sample selection criteria detailed above (i.e., I consider only males who are heads of their household and who are married individuals and who are between 26 and 63 years of age, and I exclude the self-employed and the multiple job holders). The hours profile is slightly declining over the life cycle. The age profile of the hourly wage rate is hump-shaped, as in many microeconomic data sets, and the wage decline occurs relatively late in the life cycle (at around age 55). Note that, in the raw data, hours and wages are negatively correlated, a matter to which I will return later.

Table 1
Sample Characteristics

Variable	1989	1991	1993
Age	42.07	43.26	44.81
	(8.59)	(8.69)	(8.40)
Education	10.33	9.96	9.66
	(4.30)	(4.17)	(4.07)
North	.38	.49	.50
	(.49)	(.41)	(.43)
South	.50	.49	.48
	(.45)	(.40)	(.37)
Family size	3.86	3.81	3.79
	(1.07)	(1.03)	(.99)
Number of income recipients	1.64	1.68	1.75
	(.73)	(.74)	(.73)
Children aged 0–5 years	.35	.29	.24
	(.59)	(.55)	(.50)
Children aged 6–13 years	.61	.57	.54
	(.76)	(.73)	(.72)
Children aged 14–17 years	.35	.33	.33
	(.61)	(.58)	(.55)
Working wife	.36	.35	.35
	(.48)	(.48)	(.48)
Family income ($)	21,415.63	21,033.32	21,617.61
	(10,642.93)	(10,752.47)	(21,033.32)
Working months	11.85	11.83	11.66
	(.92)	(.96)	(1.42)
Weekly normal hours	40.83	40.67	39.50
	(5.78)	(5.92)	(7.18)
Earnings ($)	12,443.24	12,193.92	12,173.83
	(4,921.95)	(5,709.74)	(4,642.11)
Wage rate ($)	6.49	6.50	6.88
	(2.36)	(4.67)	(3.58)
ψ_{it}		.0028	.0132
		(.2521)	(.03438)
$E_{t-1}(\Delta \ln y_{it})$		−.0079	−.0300
		(.0508)	(.0594)

NOTE.—Standard deviations are in parentheses. The figures for income, earnings, and wages are in euros. ψ_{it} = the innovation in earnings growth; $\Delta \ln y_{it} - E_{t-1}(\Delta \ln y_{it})$.

B. Reliability

Economists are usually skeptical about the informational content of subjective expectations data. Manski (2000, p. 17) reports that "economists often assert that respondents to surveys have no incentive to answer questions about their preferences or expectations carefully or honestly; hence, there is no reason to believe that subjective responses reliably reflect respondents' thinking." Interestingly, economists "do not apply this reasoning to self-report of *objective* data" (p. 17). Based on his experience with the Survey of Economic Expectations data, Manski concludes that such conventional wisdom is unfounded: "Survey respondents do provide coherent, useful information when queried about their expectations" (p. 17).

The reliability of the expectations data from the SHIW has already been examined elsewhere (Guiso et al. 1992; Jappelli and Pistaferri 2000). In particular, Jappelli and Pistaferri confront inflation expectation reported by SHIW respondents with those elaborated by international institutions and find that subjective expectations are close on average to both actual reali-

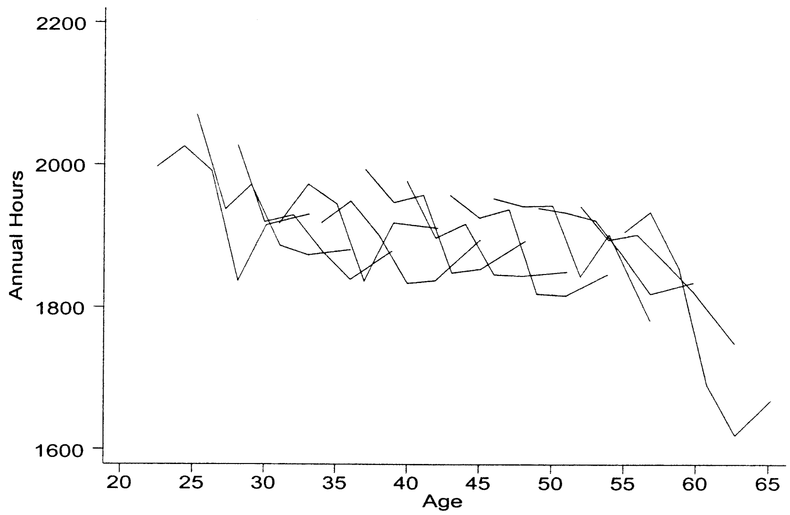

FIG. 1.—The life-cycle profile of annual hours

zations and to the forecasts elaborated by the analysts of the Organization for Economic Cooperation and Development (OECD). An interesting feature is that subjective inflation expectations do not in fact mask a great number of implausible extreme values. More than 50% of the sample bunches the entire probability distribution for inflation between 5% and 7%. Moreover, there is no clear pattern of subjective expectations by region, age, education, or income.

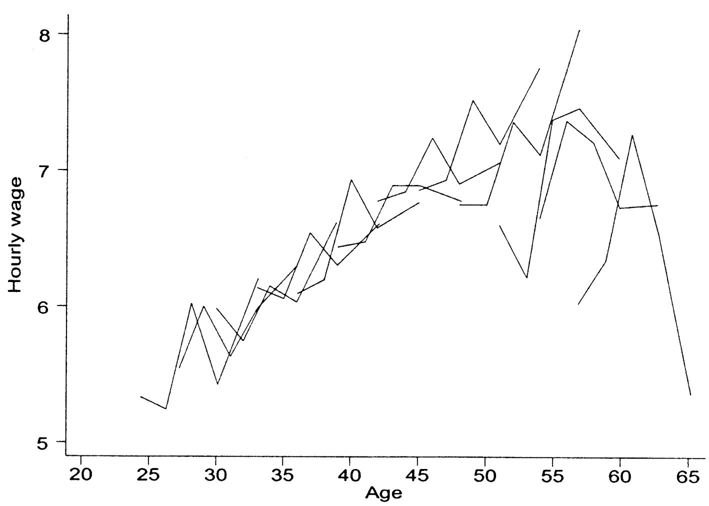

FIG. 2.—The life-cycle profile of hourly wages

Jappelli and Pistaferri (2000) also compare earnings expectations with realizations for various demographic groups. Subjective expectations can be criticized because respondents may not fully understand the survey questions: households with better education might, therefore, give more accurate income forecasts simply because they understand the survey questions better. However, Jappelli and Pistaferri find that individuals with less education do not appear to answer the survey questions less accurately than those with more education.

I checked the reliability of nominal earnings expectations by comparing subjective expectations with the actual realizations. In the raw data, the correlation coefficient of realized earnings and expected earnings is 0.46, which is prima facie evidence that subjective expectations do have informational content.[11] The correlation coefficient between realized earnings growth and expected earnings growth is instead 0.05. Thus, the relation between expectations and realizations has a cross-sectional component that is stronger than the time-series component. Cross-sectional variability is, however, the one that I exploit in the estimation of my model, to which we now turn.

IV. Estimation Results

A. Basic Specification

My purpose in this section is to report and comment on the empirical estimates of the Euler equation for labor supply (15), which is reproduced here with the account for the biennial nature of the SHIW data:[12]

$$\begin{aligned}\ln h_{it}^{*} - \ln h_{it-2}^{*} = {} & (\mathbf{Z}_{\mathrm{it}}' - \mathbf{Z}_{\mathrm{it-2}}')\beta_0 + \beta_1 E_{t-2}(\ln y_{it} - \ln y_{it-2}) \\ & + \beta_2 \operatorname{Var}_{t-2}(\ln y_{it} - \ln y_{it-2}) + \beta_3 (i_t - E_{t-2}\pi_t) \\ & + \beta_4 [(\ln y_{it} - \ln y_{it-2}) - E_{t-2}(\ln y_{it} - \ln y_{it-2})] \qquad (16) \\ & + e_{it}.\end{aligned}$$

Due to nonlinearities in the parameters, a nonlinear least squares procedure is used to identify the parameters of interest (see below) and to obtain the appropriate standard errors. The real interest rate $(i_t - E_{t-2}\pi_t)$ is constructed as the difference between the after-tax nominal interest rate on

[11] Bernheim and Levin (1989) use the Health and Retirement Survey (HRS) to compare expected social security benefits with actual realizations, and they find a raw correlation coefficient of 0.41.

[12] The dependent variable is the growth in annual hours of work. The SHIW elicits weekly overall hours h^n (including overtime) and months of employment m. In the empirical analysis below, annual hours are defined as the product of weekly hours and weeks of employment: $h = 4h^n m$. The wage variable, whenever used, is defined as the ratio between annual earnings and annual hours.

saving accounts and the subjective expectation of the inflation rate. By virtue of heterogeneity in inflation expectations, this measure of the real interest rate displays cross-sectional variability.[13] I include the following taste shifters: age, education, regional dummies, number of children in three age bands (0–5, 6–13, 14–17), the number of additional income recipients, a dummy for working wife, and time dummies.

The following structural parameters are estimated:

η is the elasticity of intertemporal substitution. This measures the reaction of labor supply to an evolutionary change in wages (i.e., keeping constant the marginal utility of wealth). The theory requires that $\eta > 0$.

Γ is the wealth effect of a parametric permanent shift in the marginal utility of wealth. The theory requires that $\Gamma < 0$. In turn, $\Gamma = \sum_{\tau=0}^{T-t} \gamma_\tau$.

φ is a parameter that measures the extent of intratemporal nonadditivity between consumption and leisure. If $\varphi > 0$ (< 0), consumption and leisure are substitutes (complements) in utility.

ϑ^ξ is a parameter that measures the reaction of labor supply to wage risk. The theory requires $\vartheta^\xi < 0$.

$\eta + \Gamma$ captures the reaction of labor supply to a parametric permanent shift in the wage profile. The theory has no sign predictions concerning this parameter.

$\eta + \varphi$ measures the response of labor supply to the gap between the expected real interest rate and the intertemporal discount rate (and also the elasticity of the marginal utility of wealth). The theory requires that $\eta + \varphi > 0$.

In the basic specification (reported in table 2), the elasticity of intertemporal substitution, η, displays a point estimate of .7, with a small standard error of .09. This implies that, along a known wage profile, an expected evolutionary 10% increase in the wage rate produces a decline in current labor supply (and a concurrent increase in next period labor supply) of about 7%. This estimate is slightly larger than the one obtained by MaCurdy (1981) and Altonji (1986) for the United States and by Browning et al. (1985) for the United Kingdom, where η is in the 0–0.6 range. On the other hand, it is still quite far from the range of estimates (between 1 and 4) reported by Mulligan (1995) or in studies based on aggregate

[13] I use the nominal interest rate on saving accounts because these are very widespread in our sample, as opposed to other financial assets. In particular, 92% of our sample has a savings account, as opposed to 23% holding short-term government bonds and 25% holding other assets (excluding cash). Moreover, conditioning on ownership, roughly 62% of the household assets are held in savings accounts. The before-tax nominal interest rates in 1989, 1991, and 1993 were 7.3%, 7.79%, and 7.8%, respectively. A flat tax rate of 30% was levied on the interests accrued.

Table 2
The Estimate of the Intertemporal Labor Supply Equation

Variable or Structural Parameter	Estimate	SE
Constant	−.0032	.0739
1993	−.0683	.0228
Age	−.0001	.0014
Education	−.0014	.0026
North	.0213	.0301
South	−.0306	.0299
Δ number of children aged 0-5	−.0252	.0319
Δ number of children aged 6-13	−.0003	.0296
Δ number of children aged 14-17	−.0138	.0249
Δ number of income recipients	.0161	.0216
Δ working wife	.0597	.0391
η	.7038	.0928
$\eta + \Gamma$	.5053	.0506
$\eta + \varphi$	.5928	.2892
φ	−.1111	.2748
Γ	−.1985	.0908
ϑ^{ζ}	−.1060	.0265

Note.—Definition of variables: η is the elasticity of intertemporal substitution, $\eta + \Gamma$ is the elasticity with respect to parametric wage growth, $\eta + \varphi$ is the elasticity of the marginal utility of wealth, φ captures within-period nonadditivity between consumption and leisure, Γ is the wealth effect of parametric wage growth, and ϑ^{ζ} is the coefficient attached to the conditional variance of wage growth.

data. One reason why my results are different is that I am able to control in a more effective way for uncertainty.

The effect of a wage innovation on the rate of growth of labor supply (an estimate of $\eta + \Gamma$) has a point estimate of .51, with a small standard error of .05. It is worth recalling that the sign of the impact of a permanent parametric wage shock on labor supply is generally ambiguous. However, in this particular case, the intertemporal substitution effect (η) dominates the wealth effect (Γ, which has a point estimate of $-.2$ with a standard error of .09), and, hence, the overall effect on labor supply is positive. Note that the results are in agreement with the theoretical prediction that $\Gamma < 0$. The implication of the estimate of $(\eta + \Gamma)$ is that a 10% permanent upward shift in the wage profile increases labor supply in the current period and at all future ages by approximately 5%–7% (depending on the point of the life cycle when this is measured, with young individuals being more responsive on average because they have a longer horizon over which to allocate the wealth effect of a permanent change in the wage).[14] Such an effect is larger than the one estimated with a different

[14] From the model above, we have $\Gamma = \sum_{\tau=0}^{T-t}\gamma_{\tau}$. If one assumes γ_{τ} constant over the life cycle, then $\Gamma = (T - t + 1)\gamma$. For young individuals ($T - t + 1 = 40$, say), $\Gamma = -0.2$, but for those only 10 years away from retirement ($T - t + 1 = 10$), $\Gamma = -0.05$. Thus, the wealth effect declines in absolute value, and the elasticity with respect to parametric permanent wage shocks $(\eta + \Gamma)$ increases over the life cycle (from 0.5 at age 25 to 0.7 at age 65).

identification strategy by MaCurdy (1981), who finds that $(\eta + \Gamma) =$ .08 for a sample of prime-age U.S. males. On the other hand, it agrees with what is outlined by Card (1991, p. 36). He stresses that "one source of systematic hours variation that is described by the labor supply model but ignored in most studies is idiosyncratic variation in the marginal utility of wealth." He adds that "a typical person-specific wage innovation results in a significant revision to lifetime wealth" (p. 36).[15]

The effect of the interest rate on labor supply $(\eta + \varphi)$ agrees with theoretical predictions. In particular, it implies a point estimate of the elasticity of labor supply with respect to the marginal utility of wealth of .59 (with a standard error of .29). The estimate of φ is $-.11$, but it is imprecisely measured (a standard error of .27). One reason for this low precision is that we are using an unsophisticated approximation to individuals' preferences for consumption and leisure. At face value, the evidence is consistent with consumption and leisure being complements in utility, a finding similar to that of Browning et al. (1985). In light of the large standard errors, however, I do not wish to put too much emphasis on this result.

The effect of the conditional variance of wage growth (an estimate of ϑ^{ς}) is consistent with precautionary behavior (a negative coefficient of $-.11$), and it is statistically significant (a standard error of .03). In agreement with what is argued by Mulligan (1995), however, the effect of wage risk on labor supply appears negligible: a 10% increase in the conditional variance of future wages increases current labor supply by about 0.5% (evaluated at the median). Even doubling wage risk (a very extreme experiment) would thus hardly affect labor supply. The small precautionary labor supply effect is consistent with the low wealth effect noticed above. If there are small wealth effects of a particular wage realization, there will be only second-order effects on utility and hence only small precautionary effects. Note that, from the model in Section II, $\vartheta^{\varsigma} = -\Gamma^2/[2(\eta + \varphi)]$. While I have not imposed such a restriction in our estimation, we can test whether it holds in the data. The test statistic has a borderline p-value of 4.3%, which supports the model.

As for the demographics, I find that, once I control for expected wages, wage surprises, and wage risk, they have a generally small and statistically insignificant impact on intertemporal labor supply. A time dummy for 1993 is statistically significant. The change in the working wife indicator captures family labor supply effects. It is theoretically justified within a conditional leisure demand framework. The estimate is positive, but it is

[15] Another interpretation of $(\eta + \Gamma)$ is that it measures cross-sectional difference in hours for individuals with different wage profiles. In our case, individuals with a 10% higher wage profile work approximately 5% more hours than individuals with lower permanent wages.

not well measured (.06 and a standard error of .04). It implies complementarity between household head's hours and the wife's labor market participation, which is consistent with the notion that family members enjoy leisure together.

The evidence for intertemporal substitution seems to contradict the raw data evidence (figs. 1 and 2), where hours and wages appear to be negatively correlated in the early part of the life cycle. A positive correlation between the age profile of hours and that of wages is typical of a simple model in which only intertemporal substitution effects exist. As equation (13) makes clear, however, the effect of wages on labor supply is only one of the possible factors at work. The factor $(\eta + \varphi)(E_{t-1}r_t - \delta)$, for instance, may induce a decline in hours over time even when wages are rising, as long as $E_{t-1}r_t - \delta > 0$. This could be due, for example, to lower intertemporal discount rates (lower impatience) at the beginning of the life cycle. In the same direction works the factor $\vartheta_{it}^{\zeta}\mathrm{Var}_{t-1}(\zeta_{it})$. Wage risk is more likely to be higher at the beginning than at the end of the life cycle (because of uncertainty resolution issues, initial job mismatch, etc.). Coupled with the interest rate effect, this may induce a decline in labor supply at the beginning of the life cycle even when wages are rising. Finally, the presence of unobserved taste shifts (which we discuss below) may come in the form of time-increasing preferences for leisure.

In sum, the comparison between estimation results and raw data evidence suggests that, at least in the early part of the life cycle, intertemporal substitution and wealth effects of wage changes (which imply positive growth rates in hours) are weaker than the combination of precautionary labor supply, interest rate, and taste heterogeneity effects (which imply negative growth rates). It confirms that the amount of variability in hours explained by the intertemporal substitution motive can be quite small, a fact that agrees with evidence available for other countries.[16]

B. Sensitivity Analysis

I considered some extensions to the basic specification, performed several tests to check the robustness of the results, and explored various sources of bias.

[16] A simple numerical example can be instructive. Suppose that the average expected wage growth and the average innovation are 1.4% and 2.1%, respectively (these figures are consistent with the data). If $\eta = 0.7$ and $\eta + \Gamma = 0.5$, the growth rate of hours due to wage growth is about 2%. Evaluated at the median, the wage risk part contributes a negative growth in hours of 0.5% if $\vartheta^{\zeta} = -0.11$. If the gap between the interest rate and the intertemporal discount rate is 2% and $\eta + \varphi = 0.6$, this contributes a negative growth in hours of 1.2%. The growth in hours predicted by the demographics is about −2%. Thus, hours will exhibit a negative growth of more than 1% even when wages grow on average by about 3.5%, as is consistent with the raw data and the empirical estimates.

Table 3
Sensitivity Analysis

Structural Parameter	Standard Labor Supply Equation (1)	Controlling for Unemployment (2)	Controlling for Overtime Constraints (3)
η	.3180	.2569	.6931
	(.3188)	(.0528)	(.0918)
$\eta + \Gamma$		.3031	.4973
		(.0348)	(.0503)
$\eta + \varphi$		.2225	.5883
		(.1844)	(.2878)
φ		−.0344	−.1061
		(.1757)	(.2872)
Γ		.0462	−.1958
		(.0571)	(.0897)
ϑ^{ξ}		−.0505	−.1048
		(.105)	(.2729)

NOTE.—Definition of variables: η is the elasticity of intertemporal substitution, $\eta + \Gamma$ is the elasticity with respect to parametric wage growth, $\eta + \varphi$ is the elasticity of the marginal utility of wealth, φ captures within-period nonadditivity between consumption and leisure, Γ is the wealth effect of parametric wage growth, and ϑ^{ξ} is the coefficient attached to the conditional variance of wage growth. Standard errors are in parentheses.

1. *Extensions and Robustness Checks*

An important question is whether there is any substantial gain produced by the use of subjective expectations. In particular, one may wonder whether the parameters that we are estimating will be the same if subjective expectations are not available, as in standard studies of intertemporal labor supply.

The benchmark I consider to evaluate this question is the standard intertemporal labor supply regression estimated by MaCurdy (1981), where there is no control for wage innovations and wage risk and where typical instruments for $\Delta \ln w_{it}$ are a cubic in age, education, and the interaction between age and education.[17] These are then excluded from the vector $\mathbf{Z}$. In our sample, this regression produces a point estimate of the intertemporal elasticity of substitution η of .32 with a standard error of the same magnitude (see col. 1 of table 3). With the obvious caveat represented by the large standard error of this estimate, it is worth noting that this number is not far, in terms of magnitude and statistical significance, from the one estimated in previous empirical work with panel data. However, it provides misleading information concerning the impact of an expected wage change on labor supply because it neglects the fact that the information available to an individual differs from that available to

[17] MaCurdy uses also family background variables, but these are unavailable in the SHIW.

the econometrician.[18] Furthermore, with the limited data available, it is impossible to obtain an estimate of the impact of parametric wage changes on labor supply or other structural parameters.

An important issue has to do with the impact of unemployment and other demand side variables on labor supply. The major criticism addressed to studies of intertemporal labor supply is in fact that employed individuals may face hours constraints. Sizable year-to-year changes in annual hours will then be associated with movements between jobs, multiple job holding, voluntary or involuntary unemployment, a switch from full-time to part-time employment (or vice versa), or measurement error.

Ham (1986) suggests testing whether unemployment is part of the optimal plan by introducing a dummy for unemployment in the λ-constant labor supply function. His argument is that, if unemployment is a voluntary choice, then the dummy should be statistically insignificant. In contrast, if unemployment reflects hours constraints, then the coefficient should be negative and statistically significant.

I experiment by adding two variables that reflect demand conditions. The first is ΔU_{it}, the first difference of a dummy that equals one if the individual worked less than 12 months and zero otherwise. This affects 6% of our sample (1% experience a fall in months of employment while 5% experience an increase). The second variable, ΔN_t, is the change in the local (province) unemployment rate. Given that I am directly controlling for the forecast error in the marginal utility of wealth through subjective expectations, least squares estimates should still be consistent. This allows me to treat the unemployment variables as exogenous.

The estimates of η and $(\eta + \Gamma)$, reported in column 2 of table 3, are now very close to each other (an equality test has a p-value of .42). The estimate of Γ is small and statistically insignificant, and so are the estimates of $(\eta + \varphi)$ and φ. The coefficient on the variance term halves to .05, but it is still statistically significant. The dummy for months of unemployment has a negative coefficient of −.65, with a standard error of .03. The local unemployment rate has a coefficient of −.004, with a standard error of .003.[19] The evidence is favorable to the hypothesis of hours constraints as put forward by Ham. Ham's argument, however, has been criticized by MaCurdy (1990), who argues that the unemployment dummy is compatible with voluntary unemployment if the offered wage falls below the reservation wage. The interpretation of the unemployment dummy remains an issue of considerable debate. For example, the significance of

[18] The predictive power of the first-stage regression is, as usual in these cases, very low. The adjusted R^2 in the regression for the reduced form is in fact only .0025; the F-statistic is 1.73. Wage growth is in fact very hard to predict in panel data.

[19] Inclusion of job and industry dummies has little effect.

the unemployment dummy in the labor supply equation could be the reflection of a mechanical negative correlation: workers who face an increase in unemployment from one period to the next will also report a negative growth in annual hours over the same time period.[20]

A somewhat different and perhaps more plausible way to examine the issue of hours constraint is to use direct information on the presence of these constraints. In both 1991 and 1993, the SHIW respondents are asked whether they had the option of working overtime in their job. Unless people self-select in jobs where overtime is allowed, this variable can be assumed to be exogenous. To check whether hours constraints are important, I augment my hours equation by a dummy variable for the option of working overtime. The results are reported in column 3 of table 3, and they are very similar to those reported in the basic specification of table 2. The estimate of the coefficient on the dummy variable for the option of working overtime is .04, with a standard error of .02, and this agrees with intuition: those who are given the option of working overtime may experience faster growth in annual hours (which include overtime hours). A possible interpretation of this result is that some workers are underemployed. When hours constraints are removed, faster growth in hours occurs. This, however, seems to have a negligible impact on the sensitivity of hours with respect to wage changes.

2. *Sources of Bias*

It is worth commenting on possible sources of bias in our estimates. First, we examine whether the estimated impact of unanticipated wage growth on labor supply is biased because we are neglecting transitory wage innovations. To address this point, assume that wages are the sum of a permanent martingale component (as in [7]) and a serially uncorrelated transitory shock (ε_{it}), a very popular characterization that agrees with panel data evidence. The wage innovation is, therefore, $(\zeta_{it} + \varepsilon_{it})$. The labor supply equation (13) can be rewritten as

$$\begin{aligned}\Delta \ln h_{it}^{*} = {} & \Delta Z_{\mathrm{it}}'\alpha + \eta E_{\mathrm{t}-1}\Delta \ln w_{it} - (\eta + \varphi)(E_{t-1}r_t - \delta) \\ & + (\eta + \Gamma)(\zeta_{it} + \varepsilon_{it}) + \vartheta^{\zeta}\mathrm{Var}_{t-1}(\zeta_{it}) + e_{it}.\end{aligned}$$

The error term in the Euler equation for labor supply, e_{it}, now includes the term $(\gamma_0 - \Gamma)\,\varepsilon_{it}$ (if one assumes for simplicity that the variance of the transitory shock is absorbed by the intercept). This term is obviously

[20] Given that the focus of this article is on prime-age males, I do not model participation. Implicitly, the assumption I make is that wage uncertainty affects hours choice but is not incorporated in the participation decision. In a more realistic model, participation could, of course, depend on wage uncertainty.

correlated with the wage innovation $(\zeta_{it} + \varepsilon_{it})$, for which I control in estimation.

It can be shown that, in the simplest case in which the innovation is orthogonal to the other regressors in the intertemporal labor supply equation, the bias in $(\eta + \Gamma)$ is the product of three factors: $(\gamma_0 - \Gamma)$; $\omega/(1 - \omega)$, where ω is the relative importance of the transitory innovation on the overall innovation in earnings; and $(1 + \eta + \Gamma)$. All these factors are positive (as long as $(\eta + \Gamma) > -1$). This suggests that the estimated long-run elasticity is likely an upper bound for the true elasticity. The bias is small if most of the variability in wages is of a permanent nature.

Next, consider the possibility of unobserved taste shifts. These will create a bias due to the fact that I use earnings rather than wages on the right-hand side of the labor supply equation. Equation (13) can be rewritten as

$$\Delta \ln h_{it}^{*} = \Delta Z_{it}'\alpha + \eta E_{t-1}\Delta \ln w_{it} - (\eta + \varphi)(E_{t-1}r_t - \delta)$$
$$+ (\eta + \Gamma)\zeta_{it} + \vartheta^{\zeta}\mathrm{Var}_{t-1}(\zeta_{it}) + e_{it}.$$

The error term, e_{it}, includes error in hours and π_{it}, a taste shock, with $E_{t-1}\pi_{it} = 0$. Note that earnings, wages, and taste shocks are linked through the relationship $\psi_{it} = (1 + \eta + \Gamma)\zeta_{it} + \pi_{it}$.

The source of bias is that, while ζ_{it} and π_{it} are (or can be assumed to be) orthogonal, ψ_{it} and π_{it} are not. In particular, $\mathrm{Cov}(\psi_{it}, \pi_{it}) = \sigma_\pi^2$. If ψ_{it} is orthogonal to the other regressors of the labor supply equation, one can prove that the bias in $(\eta + \Gamma)$ is given by $\sigma_\pi^2/[(1 + \eta + \Gamma)\sigma_\zeta^2]$. This suggests an upward bias if $\eta + \Gamma > -1$. The bias tends to disappear when the proportion of unexplained earnings variability due to taste shocks is small.

Thus, the possibility of an upward bias in $(\eta + \Gamma)$ due to omitted transitory shocks or unobserved taste shifts cannot be dismissed. In the absence of additional information, information on the magnitude of these two sources of bias is hard to come by. The inference one can draw is that the wealth effect (Γ) is probably higher (in absolute value) than the one I have estimated above, and the effect of permanent shocks on labor supply $(\eta + \Gamma)$ is probably an upper bound of the true value.

On the other hand, consider that the earnings innovation itself could be subject to measurement error, which would create a downward bias in the estimate of $(\eta + \Gamma)$. For instance, if the innovation in earnings is measured with error and taste shocks are present, the bias in $(\eta + \Gamma)$ is $[\sigma_\pi^2 - (\eta + \Gamma)\sigma_\omega^2]/[(1 + \eta + \Gamma)\sigma_\zeta^2 + \sigma_\omega^2]$, where σ_ω^2 is the variance of the measurement error in ψ_{it}. In general, this bias is smaller than the one calculated in the case where only taste shocks are present. It is not clear (and perhaps impossible to determine) whether the various sources of bias balance each other out.

V. Conclusions

In this article, I have used unique information about expectations of future wages collected in the 1989–93 panel section of the Bank of Italy SHIW to isolate anticipated and unanticipated wage changes and to construct a measure of wage risk. These are, in turn, related to shifts in life cycle labor supply. This is, to my knowledge, a novel approach not attempted in previous empirical work.

My empirical strategy allows us to disentangle the effect of evolutionary, expected wage changes from parametric, unanticipated wage changes. Similar to other studies in the literature, I assume that preferences are intertemporally separable. This implies that current wages do not depend on past work effort; a more complex model would relax this assumption, but it would face serious data limitations.

I estimate an intertemporal substitution elasticity of around 0.7. This number is slightly larger than the one obtained by MaCurdy (1981), Altonji (1986), and others for the United States, but it is smaller than the one used by macroeconomists in calibration analyses. I estimate that the impact of a permanent wage shock on labor supply, that is, letting the marginal utility of wealth vary in response to parametric wage shifts, is positive and about .5 on average. Also, this number is relatively large when compared with previous evidence but relatively small when compared with macroeconomic guess-estimates. A possible explanation for the discrepancy between my results and previous evidence is that I am able to isolate in a more convincing way the effect of anticipated and unanticipated wage growth on labor supply; alternatively, the presence of taste shocks is creating an upward bias in our estimates.

Another finding of this article relates to the role of wage risk in explaining intertemporal movements in labor supply. While the effect of wage uncertainty on labor supply is in agreement with the theory, implying that there is some adjustment, perhaps for some categories of workers or in particular stages of one's career, it is, for practical purposes, negligible. Finally, I find some weak evidence that consumption and leisure are complements in utility.

Appendix

Survey Questions on Income Prospects

This appendix draws from Guiso et al. (1992). In both 1989 and 1991, each labor income and pension recipient interviewed in the SHIW was asked the following two questions:

Inflation expectations: "On this table [a table is shown to the respondent] we have indicated some classes of inflation. We are interested in knowing your opinion about inflation 12 months from now. Suppose that you have 100 points to be distributed between these intervals. Are there intervals you definitely exclude? Assign zero points to these intervals.

How many points do you assign to each of the remaining intervals?" For this and the following question the intervals on the table shown to the person interviewed are: <0; 0–3; 3–5; 5–6; 6–7; 7–8; 8–10; 10–13; 13–15; 15–20; 20–25; > 25. If the response is "<0," the person is asked: "How much less than zero? How many points would you assign to this class?"

Income expectations: "We are also interested in knowing your opinion about your labor earnings or pensions 12 months from now. Suppose that you have 100 points to be distributed between these intervals [a table is shown again]. Are there intervals you definitely exclude? Assign zero points to these intervals. How many points do you assign to each of the remaining intervals?" (Answer categories are same as for the inflation expectations question.)

To construct subjective expectations and variances of the variable of interest (either the rate of growth of nominal earnings or the rate of inflation), I set the upper bound of the distribution—the open interval-at 35%. Let $x_{i\tau}$ be the variable of interest. The subjective expectation of $x_{i\tau}$ at time $\tau - 1$ is then given by

$$E(x_{i\tau}|\Omega_{i\tau-1}) = \sum_{k=1}^{K} [\Pr(x_{k-1} \leq x_{i\tau} \leq x_k)] 2^{-1}(x_k + x_{k+1})$$

and the subjective variance by

$$\mathrm{Var}(x_{i\tau}|\Omega_{i\tau-1}) = \sum_{k=1}^{K} [\Pr(x_{k-1} \leq x_{i\tau} \leq x_k)][2^{-1}(x_k + x_{k+1}) - E(x_{i\tau}|\Omega_{i\tau-1})]^2,$$

where x_k and x_0 are, respectively, the upper and the lower bounds of the distribution and Ω is the information set of the individual. Note that the intervals are not of the same size.[21]

Expected real earnings growth is $E(\Delta \ln y_{i\tau}|\Omega_{i\tau-1}) = E(\Delta \ln y^n_{i\tau}|\Omega_{i\tau-1}) - E(\pi_\tau|\Omega_{i\tau-1})$, with $\Delta \ln y^n_{i\tau}$ denoting nominal income growth and π the inflation rate. As far as the variance of real income growth is concerned, this is defined as

$$\begin{aligned}\mathrm{Var}(\Delta \ln y_{i\tau}|\Omega_{i\tau-1}) = {} & \mathrm{Var}(\Delta \ln y^n_{i\tau}|\Omega_{i\tau-1}) + \mathrm{Var}(\pi_\tau|\Omega_{i\tau-1}) \\ & + 2\mathrm{corr}(\Delta \ln y^n_{i\tau}, \pi_\tau|\Omega_{i\tau-1}) \\ & \times \sqrt{\mathrm{Var}(\Delta \ln y^n_{i\tau}|\Omega_{i\tau-1})}\sqrt{\mathrm{Var}(\pi_\tau|\Omega_{i\tau-1})}.\end{aligned}$$

Since the correlation term is unknown, I assume that it equals one (different values do not change the results much). This is the value that maximizes the variance of real earnings growth. Moreover, to avoid imputing a zero variance to those reporting point expectations, I calculate

[21] More precisely, $x_0 = 0$ for those assigning zero probability to a negative earnings growth event; otherwise, it is a value chosen by the respondent; $x_1 = 0.03$; $x_2 = 0.05$; $x_3 = 0.06$; $x_4 = 0.07$; $x_5 = 0.08$; $x_6 = 0.1$; $x_7 = 0.13$; $x_8 = 0.15$; $x_9 = 0.2$; $x_{10} = 0.25$; and $x_K = x_{11} = 0.35$.

the variance within each interval assuming a multistep uniform distribution. The formula is a generalization of the one given above.

Finally, I need to adjust the available subjective expectations data to the biennial nature of the SHIW. In fact, in the SHIW, individuals report $E_{t-2}(\Delta \ln y_{it-1})$, with $t = 1991, 1993$, rather than $E_{t-2}(\ln y_{it} - \ln y_{it-2})$, which is the variable that appears in (13). However, the variable of interest $E_{t-2}(\ln y_{it} - \ln y_{it-2})$ can be obtained as $2 \times E_{t-2}(\Delta \ln y_{it-1})$ under the assumption that $E_{t-2}(\Delta \ln y_{it-1}) = E_{t-2}(\Delta \ln y_{it})$, that is, that expected growth rates of earnings are smooth over a period of 2 years. Using the same assumption, I approximate $\text{Var}_{t-2}(\ln y_{it} - \ln y_{it-2})$, which is the variable that appears in (13), with $4 \times \text{Var}_{t-2}(\Delta \ln y_{it-1})$, where $t = 1991, 1993$.

References

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

Bernheim, Douglas B., and Levin, Lawrence. "Social Security and Personal Saving: An Analysis of Expectations." *American Economic Review* 79 (May 1989): 97–102.

Blau, Francine D., and Grossberg, Adam J. "Wage and Employment Uncertainty and the Labor Force Participation of Married Women." Working Paper no. 3081. Cambridge, MA: National Bureau of Economic Research, 1989.

Blundell, Richard, and MaCurdy, Thomas E. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics,* vol. 3a, edited by Orley Ashenfelter and David Card, pp. 1560–1695. Amsterdam: North-Holland, 2000.

Bover, Olympia. "Estimating Intertemporal Labor Supply Elasticities Using Structural Models." *Economic Journal* 99 (December 1989): 1026–39.

Brandolini, Andrea, and Cannari, Luigi. "Methodological Appendix." In *Saving and the Accumulation of Wealth*, edited by Albert Ando, Luigi Guiso, and Ignazio Visco, pp. 369–86. Cambridge: Cambridge University Press, 1994.

Browning, Martin; Deaton, Angus; and Irish, Margareth. "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle." *Econometrica* 53 (May 1985): 503–43.

Card, David. "Intertemporal Labor Supply: An Assessment." In *Advances in Econometrics*, edited by Christopher Sims, pp. 49–78. Cambridge: Cambridge University Press, 1991.

Chamberlain, Gary. "Panel Data." In *Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intriligator, pp. 1247–1318. Amsterdam: North-Holland, 1984.

Ghez, Gilbert R., and Becker, Gary S. *The Allocation of Time and Goods over the Life Cycle.* New York: Columbia University Press, 1975.

Guiso, Luigi; Jappelli, Tullio; and Terlizzese, Daniele. "Earnings Uncer-

tainty and Precautionary Saving." *Journal of Monetary Economics* 30 (November 1992): 307–37.
Ham, John C. "Testing Whether Unemployment Represents Intertemporal Labor Supply Behaviour." *Review of Economic Studies* 53 (August 1986): 559–78.
Jappelli, Tullio, and Pistaferri, Luigi. "Using Subjective Income Expectations to Test for the Excess Sensitivity of Consumption to Predicted Income Growth." *European Economic Review* 44 (February 2000): 337–58.
Low, Hamish. "Self-Insurance and Unemployment Benefit in a Life Cycle Model of Labour Supply and Savings." Working Paper no. 1999/24. London: Institute for Fiscal Studies, 1999.
Lucas, Robert E., and Rapping, Leonard E. "Real Wages, Employment, and Inflation." *Journal of Political Economy* 77 (September 1969): 721–54.
MaCurdy, Thomas E. "An Empirical Model of Labor Supply in a Life Cycle Setting." *Journal of Political Economy* 89 (December 1981): 1059–85.
———. "Interpreting Empirical Models of Labor Supply in an Intertemporal Framework with Uncertainty." In *Longitudinal Analysis of Labor Market Data*, edited by James Heckman and Burt Singer. Econometric Society Monographs, no. 10. Cambridge: Cambridge University Press, 1985.
———. "Appraising Tests of the Intertemporal Substitution Hypothesis." In *Panel Data and Labor Market Studies*, edited by Gerard Van den Berg, pp. 215–30. Amsterdam: North Holland, 1990.
Manski, Charles F. "Economic Analysis of Social Interactions." Working Paper no. 7580. Cambridge, MA: National Bureau of Economic Research, 2000.
Mulligan, Casey B. "The Intertemporal Substitution of Work: What Does the Evidence Say?" Discussion Paper no. 1995-11. Chicago: University of Chicago, Population Research Center, 1995.
———. "Pecuniary and Nonpecuniary Incentives to Work in the United States during World War II." *Journal of Political Economy* 106 (October 1998): 1033–77.
Pistaferri, Luigi. "Superior Information, Income Shocks, and the Permanent Income Hypothesis." *Review of Economics and Statistics* 83 (August 2001): 465–76.