

The labor-supply elasticity and borrowing constraints: Why estimates are biased

David Domeij^a, Martin Flodén^{a,b,*}

^a Department of Economics, Stockholm School of Economics, PO Box 6501, SE-113 83 Stockholm, Sweden

^b CEPR

Received 31 August 2004; revised 16 May 2005

Available online 20 December 2005

Abstract

The intertemporal labor-supply elasticity is often a central element in macroeconomic analysis. We argue that assumptions underlying previous econometric estimates of the labor-supply elasticity are inconsistent with incomplete-markets economies. In particular, if the econometrician ignores borrowing constraints, the elasticity will be biased downwards. We assess this bias using artificial data generated by a model in which we know the true elasticity and real-world data from the Panel Study of Income Dynamics. When applying standard econometric methods on the artificial data, we estimate an elasticity that is 50 percent lower than the true elasticity. We find evidence of a similar bias when using real-world data.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: E24; J22

Keywords: Frisch labor-supply elasticity; Liquidity constraint; Panel Study of Income Dynamics; Monte Carlo experiment

1. Introduction

Intertemporal substitution of labor is often a central element in macroeconomic analysis.¹ But it is not clear that individuals and households are willing or able to substitute labor supply and leisure over time in response to fluctuating wages. If intertemporal substitution were important for labor–leisure choices, individuals anticipating higher future wages would tend to work little

* Corresponding author.

E-mail addresses: david.domeij@hhs.se (D. Domeij), martin.floden@hhs.se (M. Flodén).

¹ The classical article emphasizing the role of intertemporal labor supply is Lucas and Rapping (1969).

today and more in the future. And similarly individuals anticipating lower future wages would work hard today and little in the future. Most empirical studies however find that anticipated wage fluctuations only lead to small changes in hours worked. For men, most estimates of the intertemporal labor-supply elasticity are between 0 and 0.5.²

The microeconomic evidence thus suggests that the elasticity is small. We argue, however, that previous estimates of the elasticity may have been biased downwards since liquidity constraints have been ignored.³ To understand this bias, consider an individual with little wealth that is hit by a temporary negative wage shock. If there are no liquidity constraints, this individual will reduce hours worked and borrow to smooth consumption. But if borrowing is not possible, consumption smoothing can only be achieved by an increase in labor supply. The labor-supply response of liquidity-constrained individuals is therefore smaller or of the opposite sign than what is predicted by an analysis that ignored such constraints.

Although previous estimates of the elasticity may reflect the average labor-supply response to wage fluctuations even in the presence of liquidity constraints, these estimates may be misleading in many settings and applications. For example, the labor-supply response to business cycle fluctuations or a tax reform may differ between wealth groups. Careful policy analysis therefore requires that such heterogeneity as well as liquidity constraints are modeled explicitly, and using a labor-supply elasticity that can actually be mapped to preferences.⁴

The goal of this paper is to quantitatively assess the bias generated by liquidity constraints. We do this by first applying standard econometric methods on synthetic data generated by an economic model in which we know the labor-supply elasticity. We then estimate the elasticity using data from the Panel Study of Income Dynamics (PSID), with and without controls for liquidity constraints.

In the next section we describe the model that is used to generate the synthetic data. The model economy is populated by a large number of infinitely lived households that face uninsurable idiosyncratic wage risk, supply labor elastically and trade a single asset. We impose an exogenous no-borrowing constraint implying that negative asset holdings are ruled out. The model results in that households self-insure against income fluctuations by accumulating a buffer stock of savings when income is high. Variations in labor supply both enhance households' possibilities to avoid low income and their possibilities to self-insure. Households that are borrowing constrained and have unusually low wages can avoid having low income by increasing labor supply. Self-insurance is facilitated since households with unusually high wages can increase labor supply and quickly accumulate large buffer stocks of savings.

Section 3 outlines the most important estimation procedures that have been applied in the empirical literature. We put particular focus on MaCurdy (1981) and Altonji (1986) because they used the PSID data set which occasionally contains detailed wealth data, and because their papers were the first, and still are among the few, that explicitly focus on the intertemporal elasticity rather than on some static elasticity. In this section, we also demonstrate how the presence of

² See for example MaCurdy (1981), Altonji (1986), Blundell and MaCurdy (1999), Ziliak and Kniesner (1999), and French (2004). Heckman and MaCurdy (1980) find higher elasticities for women.

³ A large fraction of US households hold virtually no wealth and many households do not even have a bank or checking account (see Deaton, 1991 and Diaz-Gimenez et al., 1997). It seems unlikely that these households can use credit to smooth consumption. Jappelli (1990) reports more direct evidence of liquidity constraints.

⁴ Browning et al. (1999) similarly note that microeconomic estimates often are incompatible with macroeconomic models.

borrowing constraints affects the equations to be estimated, and why estimates of the elasticity may be downward biased if borrowing constraints are ignored.

In Section 4, we report the results of using these econometric methods to estimate the labor-supply elasticity from artificial data generated by the model. We find that the downward bias is 50 percent. So when the labor supply elasticity is 0.5 in the model, estimates of the elasticity are around 0.25.

In Section 5, we apply the same econometric methods on PSID data. We first follow previous studies and do not control for liquidity constraints. As in the previous studies, this results in low estimates of the labor-supply elasticity. We then exclude households that are likely to be liquidity constrained and show that the estimated labor-supply elasticity rises as theory suggests.

The estimates that we obtain are based on log-linearization of the household's Euler equation. Recently, Carroll (2001) and Ludvigson and Paxson (2001) have demonstrated that similar methods can lead to a downward bias in estimates of the intertemporal elasticity of substitution (IES). They use artificial data generated by standard models to argue that the omitted high-order terms in the log-linear approximation are correlated with the variable used to identify the IES (the interest rate or the expected volatility of consumption growth). Attanasio and Low (2004) use a similar method but find that the bias disappears if long enough samples are used and if the interest rate is used as the identifying variable.

In the present paper, we also use artificial data to examine estimates of a parameter in the utility function based on the Euler equation. Our analysis deviates from the papers above in two ways. First, we consider the intertemporal substitution of labor rather than consumption, and we try to identify another parameter in the utility function.⁵ Second, we use wage changes rather than the interest rate or the variability of consumption growth to identify this parameter. We argue that wage fluctuations are dominated by idiosyncratic events whereas the return on capital to a larger extent is dominated by aggregate shocks. Therefore, even though we use a panel with only a few observations in the time dimension, we do not think that Attanasio and Low's findings directly apply to our study. Our estimates based on artificial data, where wage shocks are entirely idiosyncratic, confirm that having few observations in the time dimension is not a problem in principle.

We find, however, that the log-linearization of the Euler equation induces a downward bias in the estimated labor-supply elasticity, even if liquidity-constrained individuals are excluded from the artificial data sample. This bias is in nature similar to that identified by Ludvigson and Paxson and we are able to separate this bias from the bias induced by ignored liquidity constraints. We find that this approximation bias is quantitatively important but smaller than the bias from ignored liquidity constraints. The two sources of bias are closely related since liquidity constraints make decision rules for low-wealth households very non-linear. When we exclude households with little wealth from our sample, the approximation bias disappears together with the bias induced by liquidity constraints.

The idea that estimates based on the Euler equation are problematic in the presence of liquidity constraints is not new; Zeldes (1989) split his sample by liquid wealth and found that the Euler equation was rejected for the sample with little wealth but not for the sample with high wealth; Guvenen (2002) and references therein discuss the potential downward bias of the IES if some households are liquidity constrained; Ogawa (1991) showed analytically that estimates

⁵ In work parallel to ours, Chang and Kim (2003) also use artificial data to examine estimates of the labor-supply elasticity.

of the intertemporal labor-supply elasticity will be downward biased in the presence of liquidity constraints; and French (2002) also discusses this possibility.

Most previous empirical studies of the intertemporal labor-supply elasticity have nevertheless ignored the effects of liquidity constraints, and have not controlled for wealth. One exception is Ziliak and Kniesner (1999). In a sensitivity analysis they find evidence that the estimated intertemporal labor-supply elasticity indeed increases with wealth. A potential problem in their study is that they control for total wealth rather than liquid wealth, and that they only use indirect information on all wealth components except housing equity.

The low estimates of the labor-supply elasticity typically found in empirical studies have generated a number of suggested sources of downward bias in the estimates. For example, Imai and Keane (2004) argue that young workers invest in their human capital by working. They then supply much labor even if the wage is low, and estimates of the elasticity will be downwards biased if one does not control for this. Rupert et al. (2000) demonstrate that ignoring home production biases the estimated elasticity down if the productivity in non-market activities varies over the life cycle. Ignoring the effects of progressive taxes can also bias the estimates down (Blomquist, 1985, 1988). If taxes are progressive, gross wage rates will tend to vary more than marginal wage rates net of taxes. Labor supply will therefore react little to changes in gross wages, and there may be a bias if gross and net wages are not distinguished. Alogoskoufis (1987) and Heckman (1992) argue that variation in hours worked on the extensive margin is important and that estimates of the elasticity will be biased when non-working individuals are excluded from the sample. Although we ignore all these sources of bias in the present study, we do not find them unimportant.

2. The model

Our economy is populated by infinitely lived households, endowed with one unit of time which is divided between labor, h , and leisure, l .⁶ Households choose labor supply and consumption c to maximize expected discounted utility

$$U_0 = E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, h_t), \quad (1)$$

where β is the subjective discount factor. Let a_t denote a household's assets in the beginning of period t . We assume that households are unable to borrow, i.e., $a_{t+1} \geq 0$.

Let r denote the real return on savings. The household's budget constraint in period t is then given by

$$c_t + a_{t+1} = (1 + r)a_t + w_t h_t, \quad (2)$$

where w denotes the household's productivity, which evolves through time according to a first-order Markov chain. The timing convention is that w_t is observed before decisions are made in period t , and we assume that there is no aggregate uncertainty.

The solution to the household's problem is characterized by decision rules for consumption, saving, and labor supply as functions of household-specific asset holdings and productivity levels. We assume that $\beta(1+r) < 1$, which implies that the economy settles down in an equilibrium where aggregates are constant. Although the distribution of households in wealth–wage space is constant in this equilibrium, there are fluctuations and uncertainty at the household level.

⁶ The model is similar to Deaton (1991) but allows for endogenous labor supply.

Table 1
Benchmark parameter values

			Separable utility	Cobb–Douglas utility
β	0.950			
γ	0.500	μ	1.50	2.28
ρ	0.900	α	30.00	0.39
σ_ε	0.210			
σ_ψ	0.340			

Note. The parameter μ has been chosen so that the degree of risk aversion for consumption fluctuations is 1.5, and α is chosen so that average labor supply is approximately 0.33.

2.1. Parameterization

The model period is one year and the discount factor β is set to 0.95. We assume that logged productivity follows an AR(1) process with fixed effects,

$$\ln w_t = \psi + z_t, \quad (3)$$

$$z_t = \rho z_{t-1} + \varepsilon_t.$$

Various authors have estimated similar stochastic processes for logged labor productivity using data from the PSID. These processes can be summarized by ρ , σ_ε , and σ_ψ – the serial correlation coefficient, the standard deviation of the innovation term ε , and the standard deviation of fixed effects ψ , respectively. Allowing for the presence of measurement error and the effects of observable characteristics such as education and age, work by Card (1991), Hubbard et al. (1995), and Flodén and Lindé (2001) indicates a ρ in the range 0.88 to 0.96, and a σ_ε in the range 0.12 to 0.25. We adopt the estimates of Flodén and Lindé and use the parameter values $\{\rho, \sigma_\varepsilon, \sigma_\psi\} = \{0.90, 0.21, 0.34\}$.⁷

The fraction of households that are or are close to being borrowing constrained will be an important statistic in our analysis. According to Diaz-Gimenez et al. (1997) the bottom 40 percent in the US wealth distribution own 1.4 percent of the total capital stock. We set the interest rate to match this value, implying that $r = 0.02$.

The parameter values are summarized in Table 1, and the utility function is specified below.

3. Estimation procedure

In this section we describe the most common estimation procedures used in the literature.⁸ These are based on first-order conditions for household optimization. Additional assumptions such as separability of the utility function or an explicit functional form for the utility are then added in order to obtain equations that can be estimated.

The first-order conditions for the household's utility maximization in this framework are

$$u_{c_t} = \lambda_t, \quad (4)$$

$$\lambda_t w_t = -u_{h_t}, \quad (5)$$

$$\lambda_t - \phi_t = (1 + r)\beta E_t \lambda_{t+1}, \quad (6)$$

⁷ When solving the model, we approximate the productivity process by a discrete Markov process using seven grid points for z and two grid points for ψ .

⁸ See Blundell and MaCurdy (1999) for a more elaborate discussion.

where λ_t is the marginal utility of wealth in period t and ϕ_t is the marginal utility of borrowing in period t . In principle we can use these conditions to estimate the household’s willingness to intertemporally substitute. In practice, this may only be done if consumption and hours worked enter the instantaneous utility function in a tractable way.

Before turning to the details of the estimation procedures, it is important to specify exactly which labor-supply elasticity we intend to estimate. The Frisch (or constant marginal utility of wealth) labor-supply elasticity is defined as

$$\eta^\lambda \equiv \left. \frac{dh}{dw} \frac{w}{h} \right|_{\lambda}.$$

From the first-order conditions (4) and (5), we get

$$\eta^\lambda = \frac{u_h}{hu_{hh} - \frac{hu_{hc}^2}{u_{cc}}}. \tag{7}$$

This elasticity shows how labor supply responds to an intertemporal reallocation of wages that leaves marginal utility of wealth unaffected. As Blundell and MaCurdy (1999, p. 1595) point out, this ‘is the correct elasticity for assessing the impact of wage changes through time on labor supply.’

3.1. Estimation with separable utility

To obtain tractable results, it is often assumed that utility is separable in consumption and leisure, and this assumption has frequently been employed in the empirical literature. As the starting point for our analysis, let us therefore assume that the instantaneous utility function is separable and takes the following form

$$u(c, h) = \frac{c^{1-\mu}}{1-\mu} - \alpha \frac{h^{1+1/\gamma}}{1+1/\gamma}. \tag{8}$$

This specification of the utility function is convenient since γ is the Frisch intertemporal elasticity of substitution, which we aim to estimate. Further, μ is the coefficient of relative risk aversion, and α is a parameter determining the level of labor supply.

The first-order conditions can be rewritten into the following log-linear equations

$$\ln c_t = -\frac{1}{\mu} \ln \lambda_t, \tag{9}$$

$$\ln h_t = \text{constant} + \gamma [\ln w_t + \ln \lambda_t], \tag{10}$$

$$\ln \lambda_t - \frac{\phi_t}{\lambda_t} = \ln \lambda_{t+1} + \ln \beta(1+r) - \xi_{t+1}, \tag{11}$$

where ξ_{t+1} is the forecast error and (11) is a first-order approximation.

Let $\Delta x_{t+1} = x_{t+1} - x_t$. Equation (10) can then be written in first differences,

$$\Delta \ln h_{t+1} = \text{constant} + \gamma \Delta \ln w_{t+1} - \frac{\gamma \phi_t}{\lambda_t} + \gamma \xi_{t+1}. \tag{12}$$

MaCurdy (1981) and Altonji (1986) estimate the Frisch intertemporal labor-supply elasticity by regressing the changes in hours on changes in wages,

$$\Delta \ln h_{t+1} = \text{constant} + \gamma \Delta \ln w_{t+1} + \gamma \xi_{t+1}, \tag{13}$$

which follows from their assumption of perfect capital markets.

Note that the marginal utility of borrowing, ϕ , is positive if the borrowing constraint binds, and zero otherwise. If the borrowing constraint binds, the household expects wages to increase in the future, and the more future wages are expected to increase, the higher is ϕ . All this implies that ϕ_t and $E_t \Delta \ln w_{t+1}$ are positively correlated. Consequently, the estimate of γ will be biased downwards if the term $-\gamma \phi_t / \lambda_t$ is ignored.

The error term, ξ , also includes an approximation error due to the log-linearization of the first order conditions. Ludvigson and Paxson (2001) show that these approximation errors may cause substantial downward bias in estimates of other parameters derived from the Euler equation. The log-linearization works particularly poorly when decision rules are non-linear as they are in states where households are or are close to being borrowing constrained, something which will make this bias larger. Appendix A contains a detailed description of this bias and shows how we, when using artificial data, can separate the bias due to omitting ϕ/λ from the bias induced by the log-linearization.

Even if we could ignore the liquidity constraints and the log-linearization bias, there are econometrical obstacles to estimating (13). The error term ξ is correlated with the explanatory variable $\Delta \ln w$.⁹ Equation (13) is therefore estimated with instruments for the productivity changes.¹⁰ To instrument for $\Delta \ln w_{t+1}$ in our synthetic framework we use the mathematical expectation of $\Delta \ln w_{t+1}$ at time t , $E_t \Delta \ln w_{t+1}$, which contains all information about the productivity change except what is contained in ξ_{t+1} . This approach is not feasible when using real-world data, since the true productivity process is then unknown and the exact productivity levels, fixed effects, etc. are not observed.

Altonji used lagged wages and lagged wage changes to instrument for future wage changes. Note that the wage process (3) implies that

$$\Delta \ln w_{t+1} = z_{t+1} - z_t = (\rho - 1)z_t + \varepsilon_{t+1} = (\rho - 1)(\ln w_t - \psi) + \varepsilon_{t+1}. \quad (14)$$

This equation shows that $\ln w_t$ is a useful instrument since it is correlated with the wage change $\Delta \ln w_{t+1}$, but there are two potential problems when using $\ln w_t$ as the instrument. First, $\ln w_t$ is correlated with the error term ξ_{t+1} , since $\xi_{t+1} = \varepsilon_{t+1} / (E_t \lambda_{t+1} + \varepsilon_{t+1})$ and $E_t \lambda_{t+1}$ is on average higher for households with low current income. Second, $\ln w_t$ is correlated with the fixed effect ψ . This motivates, as is sometimes done in the empirical literature, using additional household variables when estimating (14). In our synthetic sample, the results are not sensitive to the choice of instruments. When using real-world data, we follow Altonji and use lagged wages to instrument for the wage change.

An alternative procedure for estimating the elasticity is pursued in Altonji (1986). He uses (9) to rewrite Eq. (10) in log-levels,

$$\ln h_t = \text{constant} + \gamma \ln w_t - \frac{\gamma}{\mu} \ln c_t, \quad (15)$$

and uses data on hours worked, wages, and food consumption to estimate γ .¹¹ The advantage with this procedure is that liquidity constraints do not enter the equation. However, the use of

⁹ An exception is when households know their wage one period ahead, which is assumed by MaCurdy and in some of Altonji's specifications. In our model the *only* innovation between periods is the shock to household productivity.

¹⁰ Another reason to instrument for $\Delta \ln w$ is the occurrence of measurement errors. MaCurdy used year dummies and individual specific information such as age and education as instruments. Altonji used two different wage series for each household.

¹¹ Browning et al. (1999, Section 3.2.2) discuss a similar approach.

consumption data is problematic. It is difficult to find good microdata containing both total consumption and labor supply. The PSID, used by Altonji, contains food consumption and income data, but using food consumption as a proxy for total consumption requires that food consumption is sufficiently separable from other consumption goods in the utility function.

To handle participation constraints, some authors (e.g. Browning et al., 1985, and Blundell et al., 1993, 1998) have estimated semi-log equations of the form¹²

$$\Delta h_{t+1} = \text{constant} + \gamma \Delta \ln w_{t+1} + \varepsilon_{t+1}. \tag{16}$$

Again, a correlation between $\Delta \ln w_{t+1}$ and the error term is likely. When testing this specification, we therefore instrument for $\Delta \ln w_{t+1}$ using its mathematical expectation. Blundell and MaCurdy (1999) point out that this equation cannot be derived from any standard utility function. Furthermore, we have found that this formulation performs poorly on the synthetic data.

3.2. Non-separable utility

The assumption that preferences are separable in consumption and leisure is generally regarded as restrictive and possibly unrealistic. Quite naturally, therefore, attempts have been made to allow for more general preferences in the empirical literature. Altonji (1986) argues that adding cross-substitution terms to Eqs. (9) and (10) results in an approximation of the log-linearized first-order conditions. As long as measurement errors are negligible, the elasticity can still be estimated from the difference form (as in Eqs. (12) and (13)). The log-level form (15) will, however, result in a biased estimate of the labor-supply elasticity.

It is not clear what restrictions on preferences are needed for Altonji’s approximation to be valid. Blundell and MaCurdy (1999) indeed argue that estimating the Frisch elasticity is not possible unless preferences are separable.

Unless a specific functional form for preferences is assumed, it will be difficult to further assess this issue. Let us therefore also consider the Cobb–Douglas utility function that has been used frequently in the macroeconomic literature,

$$u(c, h) = \frac{[c^\alpha (1 - h)^{1-\alpha}]^{1-\mu}}{1 - \mu}.$$

For this utility function, we can derive the intertemporal elasticity of leisure as

$$\eta^l = \frac{1 - \alpha(1 - \mu)}{\mu}.$$

The Frisch elasticity, $\eta^\lambda = \frac{1-h}{h} \eta^l$, then depends on each household’s labor supply. If participation constraints and borrowing constraints are not binding, one can derive the following relationships¹³

$$\Delta \ln(1 - h_{t+1}) = \text{constant} - \eta^l \Delta \ln w_{t+1} - \frac{\xi_{t+1}}{\mu} \tag{17}$$

and

$$\ln(1 - h_t) = \text{constant} - \ln w_t + \ln c_t. \tag{18}$$

¹² The participation constraint is never binding with the utility function in Eq. (8).

¹³ See Appendix A for the expressions with constraints.

If $\Delta \ln h_{t+1}$ is used on the left-hand side of (17), there is no reason to expect that the Frisch elasticity for labor supply will be estimated. First, the change in log hours is not identical to the negative of changes in log leisure. Further, even if that were the case, the elasticity for leisure is different from the Frisch elasticity. Equation (18) shows that the log-level equation, even if estimated on leisure, is not related to the intertemporal labor-supply elasticity.

The straightforward approach for obtaining an estimate of the average labor-supply elasticity is thus to estimate η^l from (17) and multiply by average $(1-h)/h$ (see Heckman and MaCurdy, 1980; Browning et al., 1999). One problem with this approach is to measure leisure in the data.¹⁴ Moreover, there will be a downward bias in the estimate of η^l if borrowing constraints are binding but ignored. Finally, with the Cobb–Douglas utility function, participation constraints will bind for households with much wealth and low wages. If participation constraints bind, Eq. (5) does not hold with equality and the estimate of η^l from (17) will be biased.¹⁵ In empirical work households with no or low labor supply are often excluded. Since h_{t+1} is correlated with the wage change, that procedure may still induce a bias in the estimates.

4. Estimation on synthetic data

The previous section shows that estimates of the labor-supply elasticity will be biased if liquidity constraints are ignored. To examine the quantitative importance of this bias we estimate the elasticity using synthetic data generated by the model in Section 2. As the benchmark, we have used the utility function (8) and the parameter values $\mu = 1.5$ and $\gamma = 0.5$ for the risk aversion and labor-supply elasticity, respectively. The synthetic sample consists of 500 individuals simulated for 400 periods.

Table 2 shows results from estimations of (13) on data generated by the model. Columns 1–4 contain results based on the full sample while columns 5–8 contain average results of random samples that are similar in size to those in the PSID data samples used in Section 5.¹⁶ The first column reports regression results for the full sample, thus ignoring liquidity constraints captured by ϕ . The estimated elasticity is then 0.23 (recall that the true elasticity is $\gamma = 0.5$). Column 2 shows that the estimate is much closer to the true value when liquidity-constrained households are excluded from the sample, but the estimate is still less than the true value because of the approximation error associated with log linearization.

Appendix A shows how to separate the bias stemming from liquidity constraints and the approximation error. The results of separating the bias accordingly is presented in Table 3. Most of the bias is due to liquidity constraints but the bias due to approximation errors is not negligible. For the full sample (column 1), liquidity constraints account for 70 percent of the bias. As expected, the bias due to liquidity constraints is calculated to be zero for samples that exclude households with no wealth (columns 2 and 4). The bias due to approximation error is

¹⁴ Heckman and MaCurdy (1980) found that the estimates for η^l (around 0.4 in their study) were not sensitive to the assumption of total available time. However, the ratio $(\text{total time} - h)/h$ is sensitive to that assumption. If available time is 8760 hours (24 hours per day), the ratio is 5.5, but if instead available time is 5000 hours, the ratio falls to 2.7. The implied labor supply elasticities are 2.2 and 1.1 respectively. According to Browning et al. (1999), the ratio is 4 and the implied Frisch elasticity is 1.6.

¹⁵ Our simulations show that participation constraints may be important. For example, estimates of (18) are biased downwards by 10 percent. When using the difference form in (17), the bias due to participation constraints in periods t and $t+1$ almost cancel.

¹⁶ Three random years were chosen, and estimates were calculated based on these three years, each year containing 500 individuals. This was repeated 100 times.

Table 2
Labor-supply estimates, separable utility and synthetic data

	All synthetic data				Average of small samples			
	$a_{t+1} \in$				$a_{t+1} \in$			
	$[0, \infty)$	$(0, \infty)$	$[0, 0.1\bar{a})$	$[\bar{a}, \infty)$	$[0, \infty)$	$(0, \infty)$	$[0, 0.1\bar{a})$	$[\bar{a}, \infty)$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
φ	0.00 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	0.01 (0.00)	0.00 (0.00)	0.01 (0.00)
$\Delta \ln w$	0.23 (0.01)	0.44 (0.00)	-0.09 (0.00)	0.50 (0.01)	0.24 (0.03)	0.45 (0.06)	-0.06 (0.06)	0.50 (0.10)
R^2	0.06	0.07	0.03	0.13	0.04	0.05	0.03	0.11
# obs.	199,000	157,418	68,184	59,608	1500	1183	517	448

Notes. This table shows results from estimation of $\Delta \ln h_{t+1} = \varphi + \gamma \Delta \ln w_{t+1} - \xi_{t+1}$. Standard errors in parenthesis. Columns 1–4 report results based on the full sample of artificial data with 500 households simulated 400 time periods, and using $E_t \Delta \ln w_{t+1}$ to instrument for $\Delta \ln w_{t+1}$. Columns 5–8 report the averages of 100 random three-year samples as described in the text. In columns 5–8, lagged wages are used as instruments for $\Delta \ln w_{t+1}$, and the first-stage regression results in $R^2 = 0.04$, and $F = 6.6$.

Table 3
Separation of bias due to liquidity constraints and approximation error

	$a_{t+1} \in$			
	$[0, \infty)$	$(0, \infty)$	$[0, 0.1\bar{a})$	$[\bar{a}, \infty)$
	(1)	(2)	(3)	(4)
$\hat{\gamma}$	0.23	0.44	-0.09	0.50
Bias ($\gamma - \hat{\gamma}$)	0.27	0.06	0.59	0.00
Bias due to:				
liquidity constraints	0.18	0.00	0.54	0.00
linearization	0.08	0.05	0.04	0.01

Notes. $\hat{\gamma}$ refers to the estimated elasticities in Table 2. The sources of bias are separated as in Appendix A.

also most problematic for households with little wealth, since decision rules are then particularly non-linear. This approximation bias is therefore negligible when only wealthy households are included as in column 4.

A further confirmation that the liquidity-constrained households cause the downward bias in the estimate is evident from the third column in Table 2. Only the households with no or little wealth were included in that regression, which shows a negative labor-supply response to wage increases.

Figure 1 is useful for understanding the bias. The figure shows labor supply decisions as a function of the current wage and wealth.¹⁷ In a model with no liquidity constraints, households would choose to work hard in periods with high wages and to consume more leisure when wages are low. Here, however, the top lines in the figure show that labor supply is falling in the wage rate for households with little wealth and low wages.¹⁸ These households are (or are close to being) borrowing constrained. They would consequently have to reduce consumption drastically

¹⁷ The figure only shows decision rules for households with the high fixed effect. Decision rules for households with the low fixed effect are similar.

¹⁸ Liquidity constrained households effectively face a static problem, and their labor supply curve is backward bending as in Fig. 1 if consumption and leisure are gross complements.




Fig. 1. Decision rules for labor supply and distribution of households. The different lines represent labor supply decision rules for individuals with different wealth levels. Wealth increases from top to bottom. The top line shows labor supply decisions as a function of the idiosyncratic wage for a household with no wealth. The bottom line shows decision rules for a household with asset holdings sixteen times as high as average asset holdings. The figure only displays decision rules for a few selected wealth levels. The dots indicate wage and labor combinations for a subsample of the simulated data.

if they did not increase labor supply in response to falling wages. This kind of behavior is only quantitatively important for our estimations if it is displayed by a substantial fraction of households in the economy. The figure also plots the wealth–wage distribution of households in the model economy. It is clear that many households are in regions where labor supply is flat or even falling in the wage rate.

Even though our model is simplistic, its distributions over wealth, wages, and hours worked are similar to what is observed in the data. The Gini coefficient for wealth is 0.69 in the model and 0.78 in US data, and the interest rate was chosen so that the bottom 40 percent in the wealth distribution hold 1.4 percent of all assets, as in the data. The top 20 percent hold 71 percent in the model and 79 percent in the data.¹⁹ The correlations between earnings and wealth, and between disposable income and wealth are 0.44 and 0.52 in the model and 0.23 and 0.32 in the data.

The parameter estimates based on small samples reported in columns 5–8 in Table 2 are similar to the estimates based on the full sample size. Note also that these parameter estimates

¹⁹ US data based on Diaz-Gimenez et al. (1997).

Table 4
Cobb–Douglas utility and synthetic data: estimates of η^l

	Sample: households with assets, a_{t+1} , belonging to			
	$[0, \infty)$	$(0, \infty)$	$[0, 0.1\bar{a})$	$[\bar{a}, \infty)$
φ	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
$\Delta \ln w$	0.39 (0.00)	0.59 (0.00)	0.09 (0.00)	0.64 (0.00)
R^2	0.15	0.12	0.02	0.19
# obs.	199,000	161,879	68,400	58,671

Notes. This table shows results from estimation of $\Delta(1 - h_{t+1}) = \varphi - \eta^l \Delta \ln w_{t+1} - \xi_{t+1}$ using $E_t \Delta \ln w_{t+1}$ as instruments. The true η^l is 0.66. Non-working households have been excluded in all regressions. Standard errors in parenthesis.

are relatively precise although the instruments in the first-stage equation has little explanatory power.

The results using the Cobb–Douglas utility function are reported in Table 4, and reconfirm the previous findings.²⁰ With our specification, the true elasticity of leisure is 0.66, while our estimate on the full sample is 0.39. If households with little wealth and households with low labor supply are excluded, we obtain estimates close to the true elasticity.²¹

To investigate how robust the results are to our assumptions we have re-estimated the labor-supply elasticity on data generated under various model specifications. First, we solved the model with a higher labor-supply elasticities ($\gamma = 1$). The bias due to ignoring liquidity constraints is then still around 50 percent.

Second, we considered a higher coefficient for risk aversion ($\mu = 3$). In principle this has two opposing effects. On the one hand, higher risk aversion increases households preference for consumption smoothing, thus making households with little wealth less willing to reduce labor supply as wages fall. This effect tends to increase the bias. On the other hand, for a given interest rate households want to hold more precautionary wealth when risk aversion increases. This would result in fewer households being constrained and a lower bias. In our experiment, however the interest rate is recalibrated to match the fraction of households with little wealth. The first effect therefore dominates so that the bias increases when risk aversion increases.

5. Estimation on PSID data

The analysis based on synthetic data showed that ignoring liquidity constraints can cause a substantial bias in estimates of the labor-supply elasticity. The natural next step is to turn to real-world data and re-estimate the elasticity with controls for liquidity constraints. Following much of the previous empirical literature, we use the Michigan Panel Study of Income Dynamics (PSID) and focus on male household heads.²²

²⁰ We now set $\mu = 2.28$ and $\alpha = 0.39$ so that the level of labor supply and the degree of risk aversion are the same as with the separable utility function.

²¹ We have also ran regressions like (13) and (16) on the data generated by this model. The first regression yields estimates that are clearly biased downwards. The latter regression yields estimates far from the true elasticity. As noted in the previous section, these regressions have no foundation in the model.

²² Other studies using the PSID to estimate the labor-supply elasticity include MaCurdy (1981), Altonji (1986), and Ziliak and Kniesner (1999). MaCurdy used data for 1967–1976, Altonji for 1968–1981, and Ziliak and Kniesner for 1978–1987.

5.1. Using the 1984, 1989, and 1994 wealth data

The model suggests that liquidity-constrained households should be excluded from the sample before we estimate an equation like (13). We use two different approaches to identify households that are likely to be liquidity constrained. The first approach is to exclude households with little wealth from the sample. This approach was used by Zeldes (1989) and is fully consistent with our model—households in the model choose to hold no assets when the liquidity constraint binds. Our second approach is to exclude households that are likely to be liquidity constrained according to Jappelli's (1990) estimated logit model. Jappelli identified variables that are good predictors of liquidity constraints by using direct evidence on households that had been rejected credit according to the Survey of Consumer Finances.

Common for both approaches is that we need data on wealth in addition to the data on hours worked and wages. The 1984, 1989, and 1994 PSID waves contain a supplement on household wealth. Our main sample is therefore based on observations from these years.²³ An alternative approach would have been to impute wealth from the information on housing equity and capital income that is available in all PSID waves (see Zeldes, 1989; and Ziliak and Kniesner, 1999). There are many potential problems when imputing wealth from capital income. For example the return may vary between assets and over time. And some of the most liquid assets, such as cash and checking accounts, have not paid any return at all.

In the 1984, 1989, and 1994 waves, households report asset holdings at the time of the interview (typically in February or March) and income and hours worked during the preceding year. For example, we obtain h_{83} , w_{83}^* , w_{83}^{**} , a_{84} , and some household characteristics from the 1984 wave.²⁴ The first wage measure, w^* , is calculated as the household head's total labor income divided by total hours worked. The second wage measure, w^{**} , is the reported hourly wage rate, which is only available for hourly rated workers. To calculate the relevant variables for running regression (19) below, we also need to know h_{84} , w_{84}^* , and w_{84}^{**} . We obtain this information from the 1983 and 1985 waves.

We calculate two wealth measures. Liquid wealth is measured as the sum of checking and savings account balances, bonds, and stocks minus 'other debts'.²⁵ Total wealth is the net value of all assets. Note that, while the model focuses on total wealth, liquid assets may be a more relevant indicator of a household's ability smooth consumption. Houses account for a large fraction of total wealth, but a house is an illiquid asset with large transaction costs that limit its use in smoothing consumption. Angeletos et al. (2001) also show that within an overlapping generations economy, households use liquid wealth to buffer income shocks, suggesting that we should use liquid wealth rather than total wealth to control for constraints.

To obtain an alternative indicator of liquidity constraints, we use Jappelli's (1990) estimated logit equation and calculate a probability p that the household is constrained. The probability is then a polynomial function of total family income, household wealth, the head's age, education, employment status, marital status, race, family size, homeownership, and a savings dummy.²⁶

²³ See Appendix B table for details on the sample selection criteria.

²⁴ Recall that we define a_t as assets brought into period t . The household was not borrowing constrained in period $t - 1$ if $a_t > 0$.

²⁵ 'Other debts' mostly consists of credit card charges, student loans, medical or legal bills, and loans from relatives.

²⁶ We use the variables (except the region and area dummies) and parameter estimates reported in Jappelli's Table III. Note that Jappelli's estimates are based on another data set with other sample selection criteria, and his data are expressed

5.2. Results based on 1984–1994 sample

To estimate the Frisch elasticity γ , we follow Altonji (1986) and consider the difference specification,

$$\Delta \ln h_{t+1} = \text{constant} + \gamma \Delta \ln w_{t+1}^* - \xi_{t+1}, \quad (19)$$

and we use $\Delta \ln w_t^{**}$ and $\ln w_t^{**}$ to instrument for $\Delta \ln w_{t+1}^*$. The reason for using lagged w^{**} rather than lagged w^* as instruments is to avoid a bias implied by measurement errors in the wage data (see Altonji for a further discussion).

The first-stage equation is estimated as

$$\begin{aligned} \Delta \ln w_{t+1}^* = & 0.14 + 0.27 \Delta \ln w_t^{**} - 0.06 \ln w_t^{**}, \\ & (0.03) \quad (0.04) \quad (0.02) \end{aligned} \quad (20)$$

$$R^2 = 0.04, \quad F = 23.4, \quad n = 1280.$$

The low R^2 value and high F value show that lagged wages have little but significant explanatory power for future wages.²⁷ This is consistent with the implications of the model where wage fluctuations are dominated by an unpredictable component. Note that the first-stage estimation in the synthetic sample resulted in similar R^2 and F values (see Table 2), but still delivered precise parameter estimates in the second-stage.

Table 5 reports results of the second-stage regression. The first column reports regression results for the full sample, thus ignoring liquidity constraints. The estimated elasticity is then 0.16.

Columns 2–10 in Table 5 report regression results for subsamples where we have used different criteria to exclude individuals that are likely to be liquidity constrained.²⁸ In the first two subsamples we exclude individuals with little liquid assets. In the next two subsamples we exclude individuals with little liquid assets relative to historical earnings. The next four subsamples similarly exclude households with little total wealth. In the final two subsamples we exclude individuals that have a high probability of being liquidity constrained according to Jappelli's criteria. The estimated elasticities in these subsamples range from 0.18 to 0.55. All the estimates in the subsamples where we control for liquidity constraints are thus higher than in the full sample where we do not control for constraints, but the standard errors for the parameter estimates are large and the differences between the estimates are not statistically significant.

These results are supported when we restrict attention to married men (see the bottom panel in Table 5), although the liquid-assets subsamples now result in somewhat lower parameter estimates. To further investigate the robustness of our results we have also added combinations of year dummies, age, and education to the list of instruments. The first three columns in Table 6 show that the parameter estimates are unaffected.

MaCurdy (1981) and, in some specifications, Altonji use household characteristics such as age and education as the only instruments for future wage changes. Column 4 in Table 6 summarizes our results with these instruments. Since we now do not use w^{**} as an instrument we can

in 1982 dollars. Possibly, therefore, the parameter estimates that he reports are not valid for our sample. However, adjusting our data to 1982 dollars does not change the results.

²⁷ Altonji's corresponding equation (column 2 in his table A.1, and column 5 in his Table 1) resulted in even lower R^2 and F values.

²⁸ We have used the same instrumented $\Delta \ln w_{t+1}^*$ in all subsamples. There are, however, only small changes in the results when we re-estimate the first-stage regression for each subsample.

Table 5
Labor-supply estimates and PSID data

	All	Liquid assets >				Total wealth >				Jappelli's $p <$	
		0	\$1000	$\frac{1}{2}i^m$	i^m	0	\$12,000	$\frac{1}{2}i^y$	i^y	0.25	\bar{p}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
<i>Full sample</i>											
φ	0.01 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.01 (0.01)	0.00 (0.01)	0.00 (0.01)	-0.01 (0.01)	0.00 (0.01)	0.00 (0.01)
$\Delta \ln w^*$	0.16 (0.13)	0.33 (0.24)	0.36 (0.35)	0.47 (0.34)	0.49 (0.38)	0.18 (0.18)	0.23 (0.23)	0.41 (0.22)	0.49 (0.31)	0.21 (0.27)	0.55 (0.55)
# obs.	1277	831	656	694	613	1203	863	945	774	842	638
	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)
<i>Married subsample</i>											
φ	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	-0.01 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	-0.01 (0.01)	0.00 (0.01)	0.00 (0.01)
$\Delta \ln w^*$	0.19 (0.14)	0.21 (0.20)	0.21 (0.25)	0.28 (0.24)	0.24 (0.26)	0.20 (0.16)	0.24 (0.21)	0.41 (0.23)	0.46 (0.28)	0.24 (0.26)	0.54 (0.54)
# obs.	1032	696	555	582	515	984	757	815	676	746	577

Notes. This table shows results from estimation of $\Delta \ln h_{t+1} = \varphi + \gamma \Delta \ln w_{t+1}^* - \xi_{t+1}$ with $\Delta \ln w_{t+1}^{**}$ and $\ln w_{t+1}^{**}$ as instruments for $\Delta \ln w_{t+1}^*$; $i^m = 160w^*$ represents one monthly income and $i^y = 1920w^*$ represents one yearly income; $\bar{p} = 0.193$ which is the median of p in the benchmark sample. Standard errors in parenthesis. The first-stage equation results in $R^2 = 0.04$ and $F = 23.4$ in the top panel, and $R^2 = 0.04$ and $F = 20.6$ in the bottom panel.

Table 6
Labor-supply estimates and PSID data: other instruments

	(1)	(2)	(3)	(4)	(5)
w^{**}	Yes	Yes	Yes	No	No
Year dummies	Yes	No	Yes	No	Yes
Household variables	No	Yes	Yes	Yes	Yes
<i>First stage</i>					
R^2	0.04	0.04	0.04	0.00	0.01
F	12.1	7.2	5.8	2.9	3.3
<i>Elasticity</i>					
Full sample	0.17 (0.13)	0.10 (0.12)	0.11 (0.12)	0.42 (0.25)	0.51 (0.21)
$a_{liquid} > i^m$	0.38 (0.35)	0.35 (0.37)	0.29 (0.34)	1.28 (1.15)	0.80 (0.50)
$a_{total} > i^y$	0.44 (0.28)	0.35 (0.28)	0.33 (0.27)	0.62 (0.36)	0.50 (0.23)
$p < \bar{p}$	0.41 (0.47)	0.36 (0.49)	0.31 (0.46)	0.34 (0.38)	0.42 (0.29)

Notes. w^{**} indicates that $\Delta \ln w_{t+1}^{**}$ and $\ln w_{t+1}^{**}$ are used as instruments for $\Delta \ln w_{t+1}^*$. Household variables are AGE , AGE^2 , $EDUC$, $EDUC^2$, and $AGE \times EDUC$, where $EDUC$ is years of schooling. The samples in columns 4–5 include households that do not report an hourly wage w^{**} . The sample size is then 4205. $\bar{p} = 0.193$ which is the median of p in the benchmark sample.

include households that do not report hourly wages.²⁹ When we do not control for the presence of liquidity constraints, the estimated elasticity is 0.42. As in Altonji, this estimate is higher than

²⁹ When we restrict attention to the hourly rated workers, the first-stage regressions are not significant so these results are not reported.

when using lagged wages as instruments. When using liquid assets to control for liquidity constraints our estimated elasticity increases substantially to 1.28. Note, however, that the first-stage equation is barely significant and has almost no explanatory power when lagged wages are not included in the set of instruments. It should therefore be emphasized (as also Altonji remarks), that the results in the second stage are unreliable when household characteristics are used as the only instruments for wage fluctuations.

In our discussion of non-separable utility functions we concluded that, at least if utility takes the Cobb–Douglas form, the appropriate method for obtaining an estimate of the labor supply elasticity is to first estimate the leisure elasticity η^l from Eq. (17). Using the 1984–1994 sample, we find estimates of η^l that have the same pattern as the estimated labor-supply elasticities reported in Table 5. For example, when we assume that total time is 5000 hours in a year, the estimated η^l is 0.14 in the full sample, and 0.44 when excluding households with liquid assets less than one monthly income. To convert these values into labor-supply elasticities, we multiply by 1.3, which is the average of $(5000 - H)/H$. Interestingly, and contrary to Heckman and MaCurdy (1980), our implied estimates of the labor-supply elasticity are not sensitive to the assumption about the total available time.

5.3. Using savings information, 1970–1980

In 1970–1972, 1975, and 1979–1980, the PSID asked whether the respondent had “any savings such as checking or savings accounts or government bonds?” By using the answer to this question as an indicator of liquidity constraints, we are able to consider a different time period. The savings question asked in, for example, 1971 refers to the household’s gross savings early in that year. If the household responds affirmatively to the 1971 savings question, we have an indication that $a_{71} > 0$. This indication may be crude, though, since debts are not considered. In our benchmark 1984–1994 sample, we only find 831 of 1277 individuals with positive liquid assets but 1101 individuals who would answer affirmatively to this question.

We cannot use the savings information from the 1970 wave since the previous PSID waves do not contain all the necessary wage variables. The five remaining waves with savings information allow us to construct a larger sample than when we used the 1984–1994 data. The main sample selection criteria for the 1971–1980 sample are the same as for the 1984–1994 sample (see Appendix B for details).

Results based on the 1971–1980 sample are reported in Table 7. The first-stage estimates are similar to those reported in (21), but the estimated elasticities in the second stage are lower than in the 1984–1994 sample. We still find, however, that the estimated elasticity increases substantially when households with little wealth are excluded from the sample.

5.4. Indicators of liquidity constraints

In the empirical analysis presented in the previous subsections, we used a number of different indicators of liquidity constraints. It turned out that these indicators to a large extent selected the same households. For example, of the 613 households with liquid assets exceeding one monthly income in the benchmark sample, 85 percent also have total wealth exceeding one yearly income and 74 percent fulfill the criterion $p < \bar{p}$. And almost all (95 percent) of those who are selected with Jappelli’s indicator are also selected with the total-wealth indicator. It is therefore not surprising that the main message delivered by these different indicators is the same. Still, and in

Table 7
Labor-supply estimates and PSID data: 1971–1980 sample

First stage:	$R^2 = 0.02, F = 27.2$		$R^2 = 0.02, F = 24.6$	
	Full sample		Married	
	all	$a > 0$	all	$a > 0$
φ	0.00 (0.00)	-0.01 (0.01)	0.00 (0.00)	-0.01 (0.01)
$\Delta \ln w^*$	-0.32 (0.12)	0.05 (0.16)	-0.36 (0.12)	-0.01 (0.17)
# obs.	2275	1878	2140	1780

Note. This table shows results from estimation of $\Delta \ln h_{t+1} = \varphi + \gamma \Delta \ln w_{t+1}^* - \xi_{t+1}$.

particular when we do not include lagged wages in the set of instruments, there are differences between the parameter estimates obtained with different indicators.

The explanation to these differences may be that some indicators more precisely select away the liquidity-constrained households. Another possible explanation is that even small changes in the sample composition randomly affect the results when parameters are estimated with large standard errors.

6. Concluding remarks

We have argued that labor-supply estimates will be biased downward if liquidity constraints are ignored. By using standard econometric methods on artificial data, we have also demonstrated that this bias may be substantial. Finally, using PSID data on male labor supply we estimate higher elasticities when workers that are likely to be liquidity constrained are excluded from the sample.

There is a vast literature trying to estimate the labor-supply elasticity with different methods, using different data sets, and obtaining a variety of estimates. We have chosen to contrast our results to MaCurdy's (1981) and Altonji's (1986) because they used the PSID data set which occasionally contains detailed wealth data, and because their papers were the first, and still are among the few, that explicitly focus on the intertemporal elasticity rather than on some static elasticity. Our study suggests that allowing for liquidity constraints in other empirical frameworks could further enhance our understanding of the labor-supply elasticity.

Acknowledgments

We thank Kjetil Storesletten, Eric French, and seminar participants at the SED Annual Meeting 2001, Sveriges Riksbank, Toulouse Seminar in Macroeconomics 2001, New York University, Stockholm School of Economics, Institute for International Economic Studies, NBER Summer Institute 2002, EEA Annual Meeting 2002, Midwest Macro Meeting 2003, and ESSIM 2003, for valuable comments and suggestions. Financial support from the Jan Wallander and Tom Hedelius Foundation at Svenska Handelsbanken, and from the Bank of Sweden Tercentenary Foundation is gratefully acknowledged.

Appendix A. Log-linearization and approximation bias

This appendix demonstrates how the log-linearization of the Euler equation (6) may create an additional bias.

Assuming rational expectations, the realized marginal utility is equal to the expected marginal utility plus a mean-zero forecast error ε ,

$$\lambda_{t+1} = E_t \lambda_{t+1} + \varepsilon_{t+1}.$$

The Euler equation is then

$$\lambda_t - \phi_t = \beta(1+r)(\lambda_{t+1} - \varepsilon_{t+1}). \tag{9'}$$

Take logs and use a first order linear approximation on the left-hand side and a second order approximation on the right-hand side to obtain³⁰

$$\ln \lambda_t - \frac{\phi_t}{\lambda_t} = \ln \beta(1+r) + \ln \lambda_{t+1} - \frac{\varepsilon_{t+1}}{\lambda_{t+1}} - \frac{1}{2} \left(\frac{\varepsilon_{t+1}}{\lambda_{t+1}} \right)^2. \tag{14'}$$

We obtained Eq. (11) by defining $\xi_{t+1} \equiv \varepsilon_{t+1}/\lambda_{t+1}$, and by assuming that $E_t \xi_{t+1} = 0$ and that ξ^2 is small. To allow for the second order term and for a correlation between $E_t \xi_{t+1}$ and productivity at date t , define $\tilde{\xi}_{t+1}$ by $\xi_{t+1} + \frac{1}{2} \xi_{t+1}^2 = E_t (\xi_{t+1} + \frac{1}{2} \xi_{t+1}^2) + \tilde{\xi}_{t+1}/\gamma$, implying that $\tilde{\xi}_{t+1}$ is uncorrelated with information at date t . Equation (12) is then

$$\Delta \ln h_{t+1} = \text{constant} + \gamma \Delta \ln w_{t+1} - \frac{\gamma \phi_t}{\lambda_t} + \gamma E_t \left(\xi_{t+1} + \frac{1}{2} \xi_{t+1}^2 \right) + \tilde{\xi}_{t+1}. \tag{15'}$$

A bias in estimates of γ based on Eq. (13) can arise either because of borrowing constraints (ϕ ignored) or because of the approximation error ($E_t \xi_{t+1} \neq 0$ and ξ^2 ignored). Using our synthetic data, we can account for the sources of the bias.

It is straightforward to calculate simulated values for the marginal utility of consumption, λ_t , using simulated consumption data. We calculate the marginal utility of borrowing as a function of the household's state variables, $\Phi(w, a)$ as $\lambda - \beta(1+r)E(\lambda'|w, a)$ using the household's decision rules. We use simulated data on w and a to generate ϕ (we use linear approximations between grid points for assets in Φ). The forecast errors ε are then calculated as $\lambda' - (\lambda - \phi)/[\beta(1+r)]$.

We calculate $E_t(\xi_{t+1} + \frac{1}{2} \xi_{t+1}^2)$ by regressing simulated $\xi_{t+1} + \frac{1}{2} \xi_{t+1}^2$ on information available at date t .

Let $\hat{\gamma}$ denote the estimate of γ based on Eq. (13). We then see that

$$E\hat{\gamma} = \gamma - \frac{\gamma \text{cov}(E\Delta \ln w', \frac{\phi}{\lambda})}{\text{var}(E\Delta \ln w')} + \frac{\gamma \text{cov}(E\Delta \ln w', E[\xi + \frac{1}{2}\xi'^2])}{\text{var}(E\Delta \ln w')}.$$

By calculating these moments on the simulated data, we account for the sources of the bias in $\hat{\gamma}$.

With the Cobb–Douglas utility function one also needs to allow for participation constraints in (17). Similar calculations result in

$$\Delta \ln l_{t+1} = \text{constant} - \eta^l \Delta \ln w_{t+1} + \frac{\phi_t}{\mu \lambda_t} - \frac{\xi_{t+1}}{\mu} - \frac{\xi_{t+1}^2}{2\mu} + \frac{\eta^l \zeta_t}{\lambda_t w_t} - \frac{\eta^l \zeta_{t+1}}{\lambda_{t+1} w_{t+1}},$$

where ζ is the Lagrange multiplier on the participation constraint. The bias can be separated by

³⁰ We have checked numerically that higher order terms of ϕ and ε are of little importance.

$$E\hat{\eta}^l = \eta^l - \frac{\text{cov}(E\Delta \ln w', \frac{\phi}{\lambda})}{\mu \text{var}(E\Delta \ln w')} + \frac{\text{cov}(E\Delta \ln w', E[\xi + \frac{1}{2}\xi'^2])}{\mu \text{var}(E\Delta \ln w')} \\ - \frac{\eta^l \text{cov}(E\Delta \ln w', \frac{\xi}{\lambda w})}{\text{var}(E\Delta \ln w')} + \frac{\eta^l \text{cov}(E\Delta \ln w', E\frac{\xi'}{\lambda' w'})}{\text{var}(E\Delta \ln w')}.$$

In our simulations, the last two terms in the above equation tend to cancel.

Appendix B. Sample selection and data description

We use similar sample selection criteria as Altonji (1986) and consider men that are household heads and between the age 25 and 60. Altonji only considered married men. We have few observations in our sample and we see no reason to exclude non-married men. Furthermore, the results are not very sensitive with respect to this sample criteria. MaCurdy (1981) used similar selection criteria but limited his study to ages 25 to 57. Below, we provide some details on the criteria used. For further details, see the SAS and Matlab code available at <http://www.hhs.se/personal/floden/>. The hours and wage data for the selected samples are summarized in Table B.1.

B.1. The 1984–1994 sample

We describe the selection of observations from the 1983–1985 PSID samples. Observations from the 1988–1990 and the 1993–1995 samples were selected analogously.

From the 1983 to 1985 samples we obtain (among other variables) a_{1984} , h_t , y_t , and w_t^{**} where t is 1982, 1983, and 1984 and where y is total labor income. We calculate $w_t^* = y_t^*/h_t$. Note that w_t^{**} denotes the hourly wage rate reported early in year $t + 1$ while w_t^* denotes the ‘average’ wage rate in year t .

We only include the representative SRC subsample, and all nominal variables are deflated to 1983 prices using the CPI (the 1988–1995 variables are also deflated to 1983 prices). We require that the individual is between 25 and 60 years old in 1984, that the reported age does not fall between periods, and that the reported age does not increase by more than two years between periods. We require that $0 < h_t \leq 4860$, $y_t > 0$, and $w_t^{**} > \$0.50$ for all three years, that $y_t > 0$. We exclude individuals if h , w^* , or w^{**} fall by more than 60 percent or increase by more than 250 percent. We also require that individuals continuously are in the labor market (working, temporarily laid off, or unemployed). Individuals with missing values for any of the wealth components are excluded.

At the time of conducting this study, we only have access to a preliminary release of the 1994 and 1995 PSID data (Public Release I files). This has caused some minor problems. First, the documentation for these years is scarce. Second, we have not been able to control for changes in the household composition in these years. Hours worked in 1993–1994 are collected from the ‘Hours of Work and Wage Files’ for 1994–1995, and total labor income is collected from the ‘Family Income Plus Files’ for 1994–1995.

B.2. The 1971–1980 sample

All nominal variables are deflated to 1972 prices using the CPI. The sample selection criteria are the same as for the 1984–1994 sample except that we require that if there is a ‘wife’ in the household, she is the same in all three years, we require that the $w_t^{**} > \$0.40$, and we require that there is a reported answer to the savings question.

Table B.1
Summary statistics

Year	h		w^*		w^{**}		# obs.
	mean	sd	mean	sd	mean	sd	
1969	2212	466	10.2	3.5	9.5	3.0	479
1970	2163	495	10.5	3.7	9.9	3.2	479
1971	2156	475	10.8	3.8	10.1	3.4	479
1972	2196	463	11.0	3.9	10.4	3.2	460
1973	2199	451	10.6	3.6	9.9	2.7	394
1974	2120	397	10.5	3.4	9.8	2.6	394
1975	2079	450	10.7	3.5	10.1	2.7	394
1977	2157	507	10.9	4.4	10.7	4.8	485
1978	2176	477	11.1	4.1	10.9	5.1	485
1979	2166	443	11.0	4.1	10.6	4.9	485
1980	2119	465	10.6	4.4	10.3	4.8	457
1982	1993	478	10.3	4.1	9.7	3.7	337
1983	2073	415	10.5	4.3	9.7	3.5	337
1984	2117	386	10.7	4.2	9.8	3.5	337
1987	2149	462	9.9	4.9	9.1	4.3	505
1988	2190	422	9.8	4.2	9.1	4.1	505
1989	2168	430	10.0	4.4	9.1	4.2	505
1992	2141	475	9.8	5.0	9.1	6.2	435
1993	2193	445	10.1	5.2	9.1	5.7	435
1994	2255	488	10.0	4.9	9.1	5.7	435

Note. Wages in 1983 dollars.

References

- Alogoskoufis, G.S., 1987. On intertemporal substitution and aggregate labor supply. *Journal of Political Economy* 95 (5), 938–960.
- Altonji, J.G., 1986. Intertemporal substitution in labor supply: Evidence from micro data. *Journal of Political Economy* 94 (2), S176–S215.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., Weinberg, S., 2001. The hyperbolic consumption model: Calibration, simulation, and empirical evaluation. *Journal of Economic Perspectives* 15 (3), 47–68.
- Attanasio, O.P., Low, H., 2004. Estimating Euler equations. *Review of Economic Dynamics* 7, 405–435.
- Blomquist, S., 1985. Labor supply in a two-period model: The effect of a nonlinear progressive income tax. *Review of Economic Studies* 52, 515–524.
- Blomquist, S., 1988. Nonlinear taxes and labor supply. *European Economic Review* 32, 1213–1226.
- Blundell, R., MaCurdy, T.E., 1999. Labor supply: A review of alternative approaches. In: Aschenfleter, Q., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 3A. Elsevier, Amsterdam.
- Blundell, R., Meghir, C., Neves, P., 1993. Labour supply and intertemporal substitution. *Journal of Econometrics* 59, 137–160.
- Blundell, R., Duncan, A., Meghir, C., 1998. Estimating labor supply responses using tax reforms. *Econometrica* 66 (4), 827–861.
- Browning, M., Deaton, A., Irish, M., 1985. A profitable approach to labor supply and commodity demands over the life-cycle. *Econometrica* 53 (3), 503–544.
- Browning, M., Hansen, L.P., Heckman, J.J., 1999. Micro data and general equilibrium models. In: Taylor, J., Woodford, M. (Eds.), *Handbook of Macroeconomics*, vol. 1A. Elsevier, Amsterdam.
- Card, D., 1991. Intertemporal labor supply: An assessment. Working paper No. 3602. NBER.
- Carroll, C.D., 2001. Death to the log-linearized consumption Euler equation! *Advances in Macroeconomics* 1.
- Chang, Y., Kim, S.-B., 2003. From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy. *International Economic Review*. In press.
- Deaton, A., 1991. Saving and liquidity constraints. *Econometrica* 59 (5), 1221–1248.

- Diaz-Gimenez, J., Quadrini, V., Rios-Rull, J.-V., 1997. Dimensions of inequality: Facts on the US distributions of earnings, income, and wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* 21 (2), 3–21.
- Flodén, M., Lindé, J., 2001. Idiosyncratic risk in the US and Sweden: Is there a role for government insurance? *Review of Economic Dynamics* 4 (2), 406–437.
- French, E., 2002. The effects of health, wealth, and wages on labor supply and retirement behavior. *Review of Economic Studies*. In press.
- French, E., 2004. The labor supply response to (mismeasured but) predictable wage changes. *Review of Economics and Statistics* 86 (2), 602–613.
- Guvenen, M.F., 2002. Reconciling evidence on the elasticity of intertemporal substitution: A macroeconomics perspective. Working paper No. 491. Rochester Center for Economic Research.
- Heckman, J.J., 1992. What has been learned about labor supply in the past twenty years? *American Economic Review* 83 (2), 116–121.
- Heckman, J.J., MaCurdy, T.E., 1980. A life-cycle model of female labour supply. *Review of Economic Studies* 47 (1), 47–74.
- Hubbard, R.G., Skinner, J., Zeldes, S.P., 1995. Precautionary saving and social insurance. *Journal of Political Economy* 103 (2), 360–399.
- Imai, S., Keane, M.P., 2004. Intertemporal labor supply and human capital accumulation. *International Economic Review* 45 (2), 602–631.
- Jappelli, T., 1990. Who is credit constrained in the US economy? *Quarterly Journal of Economics* 105 (1), 219–234.
- Lucas Jr., R.E., Rapping, L.A., 1969. Real wages, employment and inflation. *Journal of Political Economy* 77 (5), 721–754.
- Ludvigson, S., Paxson, C.H., 2001. Approximation bias in linearized Euler equations. *Review of Economics and Statistics* 83 (2), 242–256.
- MaCurdy, T.E., 1981. An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy* 89 (6), 1059–1085.
- Ogawa, K., 1991. Intertemporal substitution of labor supply in the imperfect capital market. *Kobe University Review* 37, 45–59.
- Rupert, P., Rogerson, R., Wright, R., 2000. Homework in labor economics: Household production and intertemporal substitution. *Journal of Monetary Economics* 46, 557–579.
- Zeldes, S.P., 1989. Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy* 97, 305–346.
- Ziliak, J.P., Kniesner, T.J., 1999. Estimating life cycle labor supply tax effects. *Journal of Political Economy* 107, 326–359.