Estimating Labor-Supply Elasticities with Joint Borrowing Constraints of Couples

Christian Bredemeier, *University of Cologne and IZA*
Jan Gravert, *University of Wuppertal*
Falko Juessen, *University of Wuppertal and IZA*

Abstract

Estimates of Frisch labor-supply elasticities are biased in the presence of borrowing constraints. We show that this estimation bias is less pronounced for secondary earners, because, in households with two earners and joint borrowing constraints, wage-rate fluctuations of the secondary earner are less important for the couples' willingness to borrow. We illustrate the differential importance of the estimation bias for primary and secondary earners in the framework of a simple incomplete-markets model with two-earner households and provide an empirical application using household panel data for the U.S. Our results suggest that Frisch elasticities can be estimated more consistently in samples of secondary earners.

JEL classification: E24, J16, J22

Keywords: Frisch labor supply elasticity; Incomplete markets; Two-person households

1 Introduction

The Frisch elasticity of labor supply measures the percentage reaction of labor supply to a one percent change in the net wage rate holding the marginal utility of wealth constant. This elasticity concept is decisive for understanding labor-supply reactions to, e.g., transitory tax or productivity shocks. In practice, there exists a number of econometric problems when estimating the Frisch elasticity. Estimates can be biased when gross and net wages are not distinguished (Blomquist 1985; Blomquist 1988), when non-working individuals are excluded from the sample (Alogoskoufis 1987; Heckman 1993), and when the estimated model ignores home production (Rupert, Rogerson, and Wright 2000), returns to experience (Imai and Keane 2004), or household specialization (Bredemeier 2015). Domeij and Flodén (2006) have highlighted a further estimation bias which is due to borrowing constraints. If households are borrowing constrained, they tend to raise, rather than to reduce, labor supply in reaction to negative wage rate shocks since they can no longer smooth consumption through borrowing. This causes a downward bias in estimated Frisch elasticities if borrowing constraints are ignored. Since borrowing constraints are a substantial restriction to many households in the U.S. (Diaz-Gimenez, Glover, and Rios-Rull 2011), this bias is important.

In this paper, we show that the estimation bias due to borrowing constraints is less pronounced for secondary earners than for primary earners in the household. The intuition is that negative wage rate shocks for the secondary earner can be relatively easily smoothed out by labor-supply changes of the primary earner even if the household is borrowing constrained. Consequently, negative wage rate shocks for the secondary earner play a less important role for the household's total wealth, and hence cause less pronounced surges in labor supply in borrowing constrained households. A direct empirical implication of our results is that, other things equal, better estimates of the Frisch elasticity can be obtained from samples consisting of secondary earners. Further, our results suggest that differences in labor-supply elasticities between groups who tend to have different earner roles in the household, such as women and men, are overestimated when borrowing constraints are neglected.

To make this explicit, we extend the incomplete-markets model with endogenous labor supply used by Domeij and Flodén (2006) to a two-person household framework.[1] In our model, household members face idiosyncratic wage shocks, and can insure against these shocks through precautionary savings in a non-state contingent asset and labor supply of both spouses. Decisively, household members share a joint borrowing constraint.[2]

[1]In 2012, 60% of households in the U.S. were family households (U.S. Census Bureau).

[2]This is an accurate description of the legal situation in the U.S. where most debt that is taken on within marriage is considered joint debt and common financial activities imply taking responsibility also for one's

The negative relation between wage-rate changes and labor supply in borrowing-constrained households biases the estimate of the Frisch elasticity downwards if these households are pooled with non-constrained households who show a positive relation between temporary wage rate changes and labor supply governed by the Frisch elasticity. Technically, an omitted-variable bias arises because the true labor-supply condition includes the household's willingness to borrow which is unobservable in empirical data but correlated with changes in the wage rate. The latter variable is the independent variable in standard labor-supply regressions.

In our two-person framework, the household's willingness to borrow is more strongly correlated with wage rate changes of the primary earner. Expected wage growth for the primary earner translates into a stronger change in the couple's future income and, hence, affects the couple's willingness to borrow more strongly than the same relative wage change for the secondary earner. Further, for the borrowing-constrained couple, it is easier to stabilize household income and consumption in reaction to wage shocks for the secondary earner using the primary earner's labor supply than vice versa. Consequently, the household's consumption and its willingness to borrow reacts more strongly to wage rate shocks for the primary earner. Thus, the downward bias of estimated Frisch labor-supply elasticities due to borrowing constraints is less pronounced for secondary than for primary earners.

We first show this relation analytically in a simplified version of our model. We then solve the full model using numerical techniques and use it to simulate a synthetic household panel data set on which we run standard labor-supply regressions. We corroborate that the estimation bias due to borrowing constraints is smaller in samples of secondary earners reflecting the lower correlation of their expected wage rate changes with the bindingness of the couple's borrowing constraint. Quantitatively, we find that the estimate of the Frisch elasticity for the primary earner is less than 50% of its true value which is similar to the results of Domeij and Flodén (2006). By contrast, we find that the estimate of the Frisch elasticity for the secondary earner is about 20% below its true value. Hence, our analysis shows that more consistent estimates of the labor-supply elasticity can be obtained from samples of secondary earners.

In an empirical application, we estimate Frisch elasticities using household panel data from the Panel Study of Income Dynamics (PSID). Specifically, we estimate a labor-supply regression in first differences with instrumented wage rates that would yield unbiased estimates for the Frisch elasticity in the absence of borrowing constraints. Our theoretical

spouse's pre-marriage debt. Further, spouse's credit ratings are mutually dependent.

analysis predicts a less severe downward bias in the estimate for the Frisch elasticity for secondary earners. In line with this, we find smaller point estimates for the Frisch elasticity in samples of primary earners compared to samples of secondary earners.

We rule out that these results are solely driven by gender differences in true labor-supply elasticities or gender-specific biases by documenting the same pattern in relative estimates for primary and secondary earners also within gender, i.e., in samples that exist of only men or only women, respectively. Specifically, we show that the estimated elasticities for male primary earners are smaller than the ones for male secondary earners. This finding can not be a result of gender differences in the true elasticities. Similarly, we address potential differences in the true labor-supply elasticities of primary and secondary earners which may stem from, e.g., different roles in home production. We document that the differences in estimated elasticities between male primary earners and male secondary earners are larger in wealth-poor households than in wealth-rich households (where biases from borrowing constraints are negligible). This finding is in line with our theoretical mechanism that predicts that, in the estimates for borrowing constrained households, potential differences in the true elasticities are magnified by the differential importance of the estimation bias.

Quantitatively, we find that standard techniques yield an estimated Frisch elasticity of about 0.4 for primary earners while the less biased estimate for secondary earners is 0.8 and hence twice as large. Our empirical results also corroborate the model's implication that gender differences in labor-supply elasticities are in fact smaller than suggested by standard estimation techniques. While standard techniques yield an estimate for the female elasticity which is about 85% larger than the one for men, the difference between male and female secondary earners - for which estimates are expected to be more consistent - amounts to only 18%.

The remainder of this paper is organized as follows. In Section 2, we develop a quantitative incomplete markets model with two earners to study the intertemporal behavior of a two-earner household. In Section 3, the model is evaluated quantitatively. In Section 4, we perform a Monte-Carlo study where we run standard labor-supply regressions on synthetic data from our model. Section 5 provides an empirical application using PSID data, where we estimate Frisch elasticities separately for primary and secondary earners, respectively. Section 6 concludes.

2 A simple incomplete-markets model with two-earner households

2.1 Model set-up

Decision problem. The economy is populated by heterogeneous households consisting of two household members. Households differ from one another by asset holdings and household members' wage rates. Hence, there is heterogeneity within as well as between households. Members of a household are subject to joint budget and borrowing constraints and take decisions jointly. Hence, the resulting allocations are Pareto optimal. Households are potentially borrowing constrained and use precautionary savings and labor supply of both household members to insure against bad wage rate realizations, which is similar to behavior in the model of Ortigueira and Siassi (2013).

We use indices i to identify individuals through their relative earnings potential in the household, $i = 1, 2$. We refer to the person with the higher earnings potential as the primary earner. Specifically, the stochastic wage process accounts for fixed effects, and we attach the higher fixed effect to household member 1, so that member 1 is the primary earner and member 2 is the secondary earner.

For household member 1, the wage process is given by

$$\begin{aligned} \ln w_1 &= \psi_1 + z_1, \\ z_1' &= \rho_1 \cdot z_1 + \varepsilon_1', \end{aligned} \tag{1}$$

where ψ_1 is the fixed effect of primary earners. z_1 is the stochastic wage component which follows an AR(1) process with innovations ε_1 and autocorrelation ρ_1. Throughout, a prime denotes next period values. The wage process for household member 2 is given by

$$\begin{aligned} \ln w_2 &= \psi_2 + z_2, \\ z_2' &= \rho_2 \cdot z_2 + \varepsilon_2', \end{aligned} \tag{2}$$

where ψ_2 is the fixed effect of secondary earners, $\psi_1 > \psi_2$.

Individuals have preferences over consumption c and hours worked n, which are described by the standard utility function

$$u(c, n_i) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \alpha \cdot \frac{(n_i)^{1+1/\eta}}{1+1/\eta},$$

We assume that market consumption is a household public good, i.e. there is no rivalry between spouses, $c_1 = c_2 = c$, and that an individual's utility is reduced by his or her individual hours worked.[3]

By assumption, intra-household decision making leads to Pareto optimal allocations, so that the decision process can be represented by the decisions of a household planner who maximizes a household target function v which is a weighted sum of members' utilities with weights μ and $1-\mu$ for the two household members 1 and 2, respectively,

$$v\left(c, n_1, n_2\right) = \mu \cdot u\left(c, n_1\right) + (1-\mu) \cdot u\left(c, n_2\right). \tag{3}$$

We assume full commitment between the two household members. This means that, at any time, decisions reflect individuals' initial bargaining positions, i.e. spouses' Pareto weights in a household μ are constant over time. Since we focus on reactions to short-run fluctuations in wage rates, the determination of the long-run Pareto weights is not relevant for our purposes and we can treat them as parameters. Agents discount future utility with the discount factor $\beta \in (0,1)$ and so does the household planner future values of v.

Asset markets are incomplete. Agents have access to a non-state-contingent risk-free asset a. This asset can be purchased at price $1/(1+r)$ and pays out one unit of the consumption good in the next period. Thus, r denotes the constant risk-free interest rate, which is exogenous in our model.[4] As Domeij and Flodén (2006), we assume that $\beta(1+r) < 1$. There is a restriction on borrowing, i.e., assets can not fall below an exogenous minimum level denoted by $a_{\min}$:

$$a' \geq a_{\min}. \tag{4}$$

The household budget constraint is given by

$$c + \frac{a'}{1+r} \leq w_1 n_1 + w_2 n_2 + a, \tag{5}$$

where a' are next period asset holdings to be determined and purchased at price $1/(1+r)$ this period.

Recursive formulation. We consider the problem of a household taking as given the Pareto weights μ and $1-\mu$. The household's problem can be solved with techniques of dynamic programming. The state space of the recursive maximization problem is $\Omega = [a_{\min}, a_{\max}] \times S$,

[3]Together with additively separable preferences, the public good assumption implies that a constant marginal valuation of wealth by the household is equivalent to constant marginal utility of consumption for the individuals which allows a simply notion of Frisch elasticities. The same could be achieved with purely private goods or any combination of private and public goods, where private goods are distributed among household members in a way that equalizes weighted marginal utilities, if individuals have the same risk aversion with respect to consumption.

[4]We consider a partial equilibrium set-up because we neither analyze policy nor parameter changes.

where $[a_{\min}, a_{\max}]$ denotes the interval of possible asset choices, and $S = S_1 \times S_2$ denotes the state space of the wage pair $\omega = (w_1, w_2)$.[5]

The household problem in recursive formulation is given by:

$$V(a,\omega) = \max_{a',c,n_1,n_2} \mu \cdot u(c, n_1) + (1-\mu) \cdot u(c, n_2) + \beta E\left[V\left(a', \omega'\right)|\omega\right], \tag{6}$$

subject to the borrowing constraint (4) and the budget constraint (5), for exogenous wage rates ω and a given initial value of assets a_0. The solution to the maximization problem is described by the policy functions

$$x = x(a, \omega), \tag{7}$$

with $x \in X = \{c, n_1, n_2, a'\}$.

Equilibrium conditions. The first order conditions of the recursive problem are

$$\frac{\partial v(c, n_1, n_2)}{\partial c} = \frac{\partial V(a,\omega)}{\partial a} = \lambda \tag{8}$$

$$\phi = \lambda - (1+r)\beta E\left[\lambda'|\omega\right] \tag{9}$$

$$\lambda \cdot w_1 = \mu \cdot \alpha \cdot n_1^{1/\eta} \tag{10}$$

$$\lambda \cdot w_2 = (1-\mu) \cdot \alpha \cdot n_2^{1/\eta} \tag{11}$$

$$\phi \geq 0 \tag{12}$$

$$a' \geq a_{\min} \tag{13}$$

$$\phi\left(a' - a_{\min}\right) = 0, \tag{14}$$

together with the budget constraint (5), given the exogenous wage rates w_1 and w_2 and the initial asset stock a_0. λ is the Lagrange multiplier on the budget constraint (5). ϕ is the Kuhn-Tucker multiplier on the borrowing constraint (4). Condition (9) is the consumption Euler equation of the household. Conditions (10) and (11) are the labor-supply conditions of the household members which also reflect that an individual's labor supply depends negatively on his or her Pareto weight within the household. Conditions (12)-(14) are the Kuhn-Tucker conditions associated with the borrowing constraint (4).

From (10) and (11), it can be seen that the Frisch labor-supply elasticity follows from the curvature of disutility of labor and are equal to the parameter η, independent of whether the household is borrowing constrained or not. Put differently, the true Frisch labor-supply

[5] The upper bound $a_{\max}$ guarantees a bounded state space which is needed in order to describe the economy by a probability measure. It is set high enough to never be reached.

elasticity is not affected by borrowing constraints.

2.2 Estimation equation and estimation bias

Empirical studies estimate the relation between changes in the wage rate and changes in labor supply to infer the Frisch elasticity. We use the first-order conditions above to determine the relation between these variables in our model. Taking logs of the labor-supply conditions (10) and (11), and first differences, we obtain (for $i = 1, 2$)

$$\Delta \ln w_i' = \frac{1}{\eta} \cdot \Delta \ln n_i - \Delta \ln \lambda'. \tag{15}$$

In log-linear form, the Euler equation (9) implies

$$\Delta \ln \lambda' = -\ln \beta - \ln (1 + r) - \frac{\phi}{\lambda} + \xi', \tag{16}$$

where $\xi' = \ln \lambda' - E \ln \lambda'$ denotes an expectation error. Inserting (16) in (15) and rearranging yields

$$\Delta \ln n_i' = \eta \cdot \Delta \ln w_i' - \eta \cdot \frac{\phi}{\lambda} - \eta \cdot \ln \beta - \eta \cdot \ln (1 + r) + \eta \cdot \xi'. \tag{17}$$

In this equation, an estimate of the parameter η corresponds to an estimate for the Frisch elasticity.[6]

Independent of borrowing constraints, the expectation error ξ' in (17) which is part of the residual is correlated with the regressor $\Delta \ln w_i'$. As has been discussed in the literature (see, e.g., Blundell and MaCurdy (1999)), this endogeneity problem can be resolved by an instrumental variable estimator (2SLS). When using synthetic data, we can use the (mathematical) expected relative wage change $E\Delta \ln w_i'$ to instrument for $\Delta \ln w_i'$:

$$E\Delta \ln w_i' = E(z_i' - z_i) = E((\rho - 1) \cdot z_i + \varepsilon_i') = (\rho - 1) \cdot z_i. \tag{18}$$

The covariance between $E\Delta \ln w_i'$ and $\Delta \ln w_i'$ is

$$\begin{aligned} cov\left((\rho - 1) \cdot z_i, (\rho - 1) \cdot z_i + \varepsilon_i'\right) &= (\rho - 1)^2 \cdot var\,(z_i) + (\rho - 1) \cdot cov\left(z_i, \varepsilon_i'\right) \\ &= (\rho - 1)^2 \cdot var\,(z_i) > 0, \end{aligned}$$

which shows that the instrument is informative. The instrument is valid, since the covariance between $E\Delta \ln w_i'$ and ξ' is zero, $cov\left(E\Delta \ln w_i', \xi'\right) = (\rho - 1) \cdot cov(z_i, \xi') = 0$.

A second estimation problem originates from the presence of borrowing constraints: the multiplier ratio ϕ/λ affects labor supply but is not observable and is therefore part of the

[6] The terms $\eta \cdot \ln \beta$, $\eta \cdot \ln(1 + r)$, and $\eta \cdot \xi'$ on the right-hand side of (17) can be captured by fixed effects and residuals.

combined residual in (17). In practice, one therefore estimates[7]

$$\Delta \ln n_{ijt+1} = const. + \eta \cdot \Delta \ln w_{ijt+1} + u_{ijt+1}, \tag{19}$$

where $\Delta \ln w_{ijt+1}$ is instrumented by $E_t \Delta \ln w_{ijt+1}$ and u_{ijt+1} is the combined error term, also picking up the unobserved multiplier ratio. The index $ijt+1$ refers to member i of household j in period $t+1$. If there was no correlation between ϕ_{jt}/λ_{jt} and $E_t \Delta \ln w_{ijt+1}$, the 2SLS estimator of η, $\widehat{\eta}$, would be a consistent estimator of the Frisch elasticity. However, as discussed in Section 2, the multiplier ratio, which enters the true relation (17) negatively, is positively correlated with $E_t \Delta \ln w_{ijt+1}$ and therefore the combined residual correlates negatively with the expected wage change. We now investigate this correlation analytically in a simplified version of our model, separately for primary and secondary earners, respectively.

2.3 Analytical results

In this section, we demonstrate analytically that the multiplier ratio ϕ/λ is more stronly correlated with the instrument $E\Delta \ln w_i'$ in the primary earner's labor-supply condition than it is in the secondary earner's condition. To do so, we consider a simplified version of the model, where we assume $\eta_1 = \eta_2 = \eta$, $\alpha_1 = \alpha_2 = \alpha$, $\rho = 0$, and $var(\varepsilon_1) = var(\varepsilon_2)$.

First, we can state that the multiplier ratio in the model's labor-supply conditions (17) is weakly decreasing in a symmetric, increasing function of spouses' wage levels:

Proposition 1 *Define $\Lambda = W_1^{1+\eta} + W_2^{1+\eta}$. The multiplier ratio ϕ/λ is weakly decreasing in this symmetric function of spouses' wage rates,*

$$\frac{\partial(\phi/\lambda)}{\partial \Lambda} \leq 0.$$

Proof. See appendix. ■

Second, the symmetric function of the spouses' wage *levels* is less correlated with the secondary earner's wage *shocks* which translates to the correlation of the multiplier ratio with the wage shocks:

Proposition 2 *The multiplier ratio ϕ/λ is, in absolute terms, more strongly correlated with the primary earner's wage-rate shock than it is with the secondary earner's wage-rate shock,*

$$corr(\phi/\lambda, z_1) < corr(\phi/\lambda, z_2) < 0.$$

[7] Here, we introduce household and time indices in order to clarify the panel dimension of the estimation.

Proof. See appendix. ■

This leads to the following statement regarding the relations between the multiplier ratio ϕ/λ on the one hand and the next-period wage change $\Delta \ln w_i'$ or the expected wage change $E_t\Delta \ln w_i'$, respectively, of the two spouses on the other hand. Note that the wage change $\Delta \ln w_i'$ is the regressor in an estimation of (17) and the expected change $E_t\Delta \ln w_i'$ is the instrument.

Proposition 3 *The multiplier ratio ϕ/λ is more strongly correlated with the primary earner's next-period wage-rate change than it is with the secondary earner's next-period wage-rate change,*

$$corr\left(\phi/\lambda, \Delta \ln w_1'\right) > corr\left(\phi/\lambda, \Delta \ln w_2'\right) > 0.$$

The same relation holds for the expected wage-rate change,

$$corr\left(\phi/\lambda, E\Delta \ln w_1'\right) > corr\left(\phi/\lambda, E\Delta \ln w_2'\right) > 0.$$

Proof. See appendix. ■

In the proposition, we also consider the actual wage change $\Delta \ln w_i'$ in order to clarify that the discussed biases and their relative importance do not result from any specific instrument chosen.[8]

Since the omitted variable ϕ/λ is negatively correlated with both the dependent variable $\Delta \ln n_i'$ as well as with the independent variable $\Delta \ln w_i'$ and its instrument, the estimated coefficient on $\Delta \ln w_i'$ underestimates the true parameter η. The impact of ϕ/λ on the dependent variable is the same for primary and secondary earners but the correlation with the independent variable and the instrument is stronger in the condition for the primary earner. Hence, the downward bias is larger in labor-supply regressions for primary earners.

3 Numerical analysis

In this section, we relax the simplifying restrictions from the previous analysis and solve the full model numerically. Section 3.1 introduces parameter choices and in Section 3.2, we present policy functions for the key variables to discuss the model's implications. Thereafter, in Section 4, we will use the model to simulate a synthetic panel data set of labor supply decisions for both spouses from which we will estimate Frisch labor-supply elasticities.

[8]In fact, in the primary earner's labor-supply condition, the omitted variable ϕ/λ is also more strongly correlated with the regressor itself than in the secondary earner's labor supply condition - and not only with the instrument.

3.1 Calibration

We calibrate the model parameters using household panel data for the U.S. from the Panel Study of Income Dynamics (PSID). The PSID is the most widely used source of data in the applied microeconometric literature on labor-supply elasticities, see e.g. MaCurdy (1981) and Altonji (1986), and has also been used by Domeij and Flodén (2006). Our sample is based on observations from the years 1972-1997.[9] Note that our baseline sample is relatively large in comparison to related studies. For instance, Domeij and Floden (2006) have to restrict the data to three subpanels around the years 1984, 1989 and 1994, where the PSID contained a supplement on household wealth. For our study, a relatively long sample period is important, as we need to select on households where both, primary and secondary earners are observed, next to the usual sample selection requirements.

Our sample selection is as follows. Due to our focus on double-earner households, we consider married individuals for whom both spouses' wage rates are observed. Further, we apply similar sample selection criteria as Altonji (1986) and Domeij and Flodén (2006): We consider individuals between age 25 and 60 and drop the Survey of Economic Opportunity (SEO) sample which is not representative for the U.S. We require that the reported age of the household head does not fall between periods, and does not increase by more than two years between years. We drop individuals with reported annual hours of work larger than 4860 (more than 92 hours in 52 weeks) and exclude individuals for which hours worked or the wage rate fall by more than 40 percent or increase by more than 250 percent between two consecutive years. We require that individuals are participating in the labor force, i.e., are either classified as working, temporarily laid off, or unemployed. For both partners, we calculate the wage rate as total labor income divided by total hours worked and deflate wages to 1983 prices using the CPI. To eliminate the influence of extreme observations and data errors, we drop wage observations falling in the top 0.5 percentiles of male and female wages, respectively.

A key point of our model is to distinguish between primary and secondary earners. For the quantitative evaluations of our model, we calibrate parameters which are specific to either primary or secondary earner using empirical information on men or women, respectively. Thereby, we interpret the average man in our sample as the representative primary earner and the average woman as the representative secondary earner. Of course, this interpretation is not perfect, though empirically reasonable and in line with the PSID's convention that

[9] Before 1972, there is no information on wives' education which is used an instrument in the 2SLS labor-supply regression. After 1997, the PSID switched from annual to biennial interviews.

the husband is typically regarded as the household head.[10] In the empirical investigations in Section 5, we will distinguish between men and women as well as between primary earners and secondary earners based on actual relative wage rates within couples. An individual is then defined as a primary earner if his or her average wage rate in the sample is larger than the one of his or her spouse.

Table 1 summarizes our baseline calibration. A first step in the calibration is to estimate the parameters of the wage processes for primary and secondary earners, respectively, see (1) and (2). We estimate the gender-specific autocorrelation and the gender-specific innovation variances for residual wage rates using a Generalized Method of Moments (GMM) estimator. We obtain residual wages by filtering out deterministic cross-sectional variation using a first-stage OLS regression. We then identify auto-correlations ρ_m, ρ_f and innovation variances $\sigma^2_{m,\epsilon}$, $\sigma^2_{f,\epsilon}$ from gender-specific GMM estimations where we take into account a permanent component capturing individual-specific fixed characteristics and a transitory component with persistence captured by an ARMA(1,1) process. In addition, we allow for time effects in both the permanent and the transitory wage component. Appendix B provides details on the wage process estimation. In line with the literature on earnings dynamics, we find that idiosyncratic wages display relatively high persistence with an estimated annual autocorrelation of $\rho_m = 0.84$ and $\rho_f = 0.82$. For men, the estimated standard deviation of idiosyncratic wage shocks is with $\sigma_{m,\epsilon} = 0.22$ very similar to the estimate used by Domeij and Flodén (2006). For women, we estimate that idiosyncratic labor market risk is with $\sigma_{f,\epsilon} = 0.45$ about twice that large. In the PSID data, the mean wage of women amounts to 67% of the mean wage of men. We use this statistic to calibrate the wage gap in our model by setting ψ_1 and $\psi_2 < \psi_1$ accordingly.[11]

In the PSID data, men's average yearly hours worked amount to 2260 and women work 1574 hours on average. We set men's preference parameter on the disutility of working to $\alpha_m = 43$, leading to an average working time of 1/3, and choose the corresponding parameter value for women so that the ratio of mean hours of market work of the two household members is about $1574/2260 \approx 0.7$. For the sake of illustration, we will also consider cases of our model where, besides the intra-couple wage gap, there is no further heterogeneity in the model parameters between both household members. Thus, in these specifications, we will assume that neither the disutility of work nor the parameters of the wage process depend on gender.

[10] In 76.5% of all married couples in the U.S., the wife has lower labor earnings than the husband. In married couples where both spouses work, this number is still 71% (data source: 2014 BLS databook 'Women in the Labor Force').

[11] In the model, the wage process is discretized using Tauchen's (1986) algorithm with 13 grid points per household member.

Table 1: Calibration.

Description	Parameter	Value
Discount factor	β	0.95
Relative risk aversion	σ	1.5
Frisch labor-supply elasticity	η	0.9
Labor disutility weight (prim. earner)	α_m	44
Labor disutility weight (sec. earner)	α_f	50
Pareto weight primary earner	μ	0.5
Risk-free interest rate	r	0.005
Fixed effect primary earner	ψ_1	0
Fixed effect secondary earner	ψ_2	-0.26
Autocorrelation wage process (prim. earner)	ρ_m	0.84
Autocorrelation wage process (sec. earner)	ρ_f	0.82
Std. dev. wage innov. (prim. earner)	$\sigma_{m,\epsilon}$	0.22
Std. dev. wage innov. (sec. earner)	$\sigma_{f,\epsilon}$	0.45
Borrowing limit	$a_{\min}$	0

The remaining preference parameters, i.e., the rate of time preference, β, and the coefficient of risk aversion, σ, are set to standard values in the literature. Following Domeij and Flodén (2006), we calibrate the interest rate so that the bottom 40% of the wealth distribution own 1.4% of total wealth. Most importantly, we assume that the true Frisch elasticity is identical for both household members, and given by the preference parameter $\eta = 0.9$. This modelling choice is taken in order to highlight differences in the estimation biases for primary and secondary earners. We will also consider a specification where the true Frisch elasticities differ between household members.

3.2 Policy functions

We now present policy functions for the endogenous variables to illustrate the key property of the model that wage rate fluctuations of the secondary earner affect the household's willingness to borrow less strongly than those of the primary earner as well as the mechanisms leading to this result. Figure 1 shows the own-wage labor-supply curves for both partners, i.e. labor supply of the primary earner as a function of the primary earner's wage rate, holding constant the wage rate of the secondary earner, and, labor supply of the secondary earner as a function of the secondary earner's wage rate, holding constant the wage rate of the primary earner. To highlight the role of earner status, we condition on the extreme realizations of the partner's wage rate. That is, when showing the decision rules for the primary earner, we

Figure 1: Labor supply as a function of own wage rate.

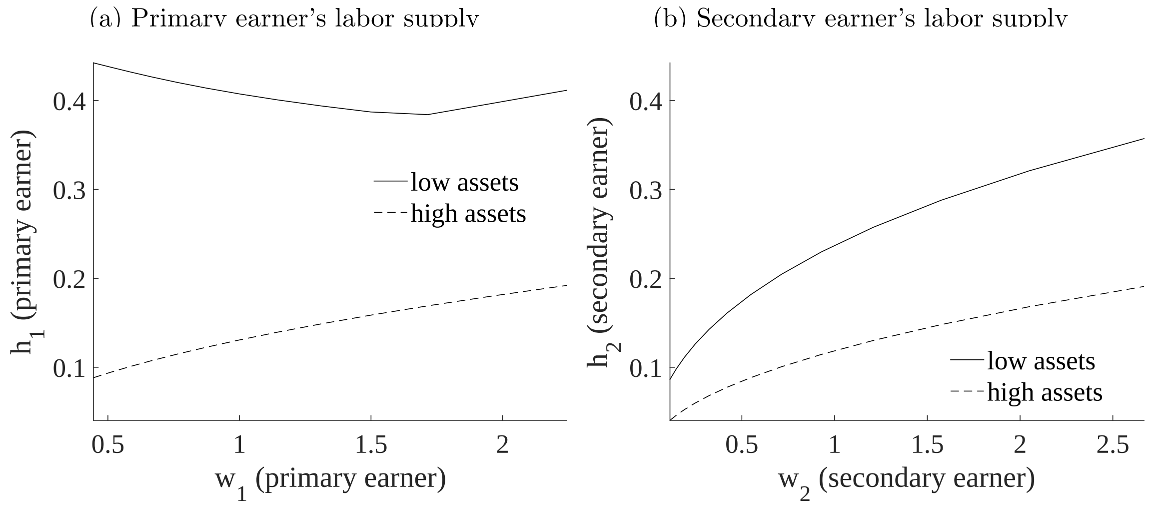

Notes: In the left panel, the wage rate of the secondary earner is held constant at its lowest possible value. In the right panel, the wage rate of the primary earner is held constant at its highest possible value.

hold the partner's wage constant at its lowest value. Vice versa, when we look at the policy functions for the secondary earner, we fix the primary earner's wage rate at its highest value. To highlight the effect of borrowing constraints, we show the policy functions for two levels of asset holdings. The dashed line refers to a situation where household wealth is high, whereas the solid line shows the decision rules when the household is borrowing-constrained (i.e., has no assets).

Figure 1 shows that labor-supply curves for low asset holdings (solid) lie above those for high assets (dashed), reflecting a standard wealth effect. The left panel shows that, for primary earners, household asset holdings are relevant for the shape of the labor-supply curves. The labor-supply curve is globally upward-sloping if the household is wealth-rich (dashed line). However, for the case where the household is wealth poor (solid line), the labor-supply curve of the primary earner has a downward-sloping range at small to medium wage rates. Thus, when wages are sufficiently low, a further wage decrease triggers an increase in labor supply (rather than a decrease), giving rise to a non-standard shape of the labor-supply function. This particular shape of the labor-supply curve for the primary earner resembles the one in the model by Domeij and Flodén (2006), where the household consists of one person only. Thus, for borrowing-constrained households, the labor-supply reaction to transitory wage rate changes is not governed by the Frisch elasticity alone. This biases the estimate of the Frisch elasticity downwards if these households are pooled with non-

Figure 2: Cross-wage labor-supply curves.

(a) Primary earner's labor supply (b) Secondary earner's labor supply

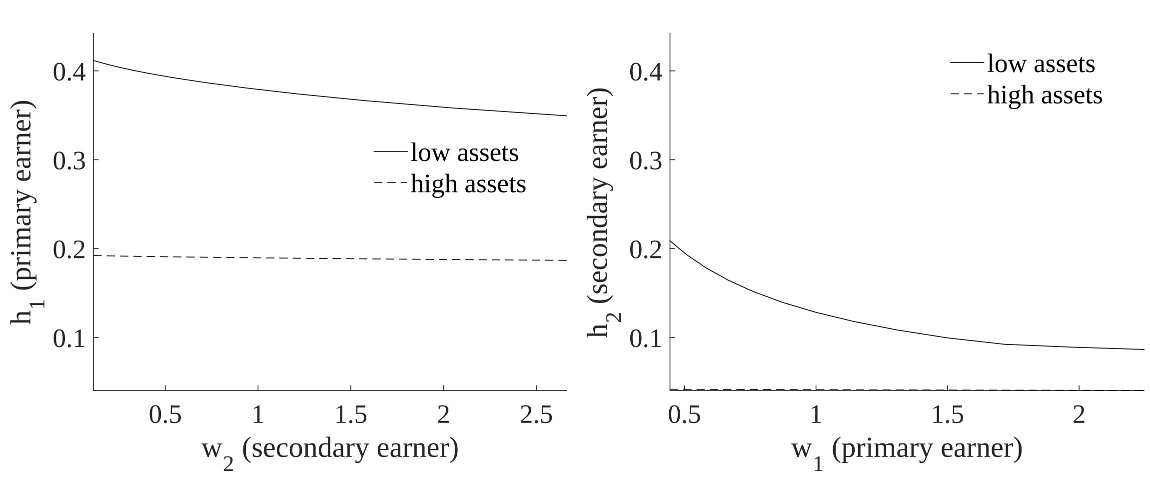

Notes: In the left panel, the wage rate of the primary earner is held constant at its highest possible value. In the right panel, the wage rate of the secondary earner is held constant at its lowest possible value.

constrained households who show a positive relation between temporary wage rate changes and labor supply governed by the Frisch elasticity.

By constrast, the right panel in Figure 1 shows that, for secondary earners with well earning spouses, labor-supply curves are globally upward-sloping, independent of whether the household is wealth-rich or wealth-poor. Thus, for secondary earners, the labor-supply reaction to transitory wage rate changes is (mostly) governed by the Frisch elasticity, so that an estimate for the Frisch elasticity based on data for secondary earners can be expected to be less biased than an estimate based on data for primary earners.

The differences in the shape of the labor supply curves for primary and secondary earners is due to the differential effect of borrowing constraints on labor supply decisions in a partnership. Specifically, consider a situation where a household is borrowing constrained and faces a negative wage rate shock. For such household, the only opportunity to stabilize household income and with it consumption is to increase labor supply. For efficiency reasons, the household finds it optimal to increase labor supply of the household member with the higher wage rate by more than the hours of the other household member. The former is typically the primary earner, the latter is usually the secondary earner. Thus, in particular for primary earners, a negative wage rate shock can actually lead to an increase (rather than to a decrease) of this spouses' labor supply, as depicted in the left panel of Figure 1.

To further illustrate this relation, we show that it is generally the primary earner's wage

Figure 3: Consumption functions.

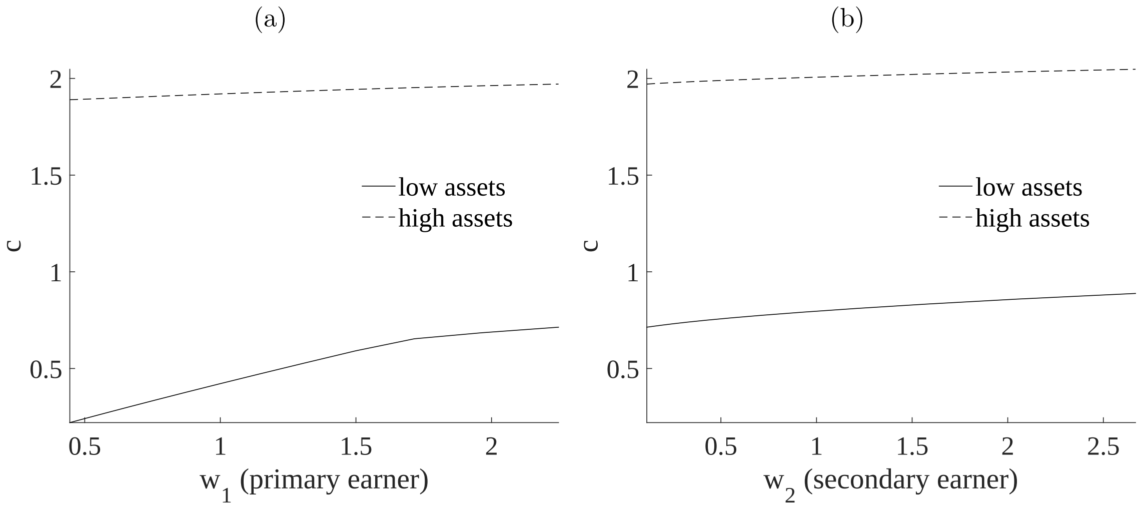

Notes: In the left panel, the wage rate of the secondary earner is held constant at its lowest possible value. In the right panel, the wage rate of the primary earner is held constant at its highest possible value.

rate that impacts strongly on labor supply in wealth-poor households. Figure 2 shows cross-wage labor-supply relations, i.e., labor supply as a function of the partner's wage rate. If the household is wealth-poor (solid lines), working time of the secondary earner reacts strongly to low wage rates of the primary earner (right panel), but the cross-wage relation between the primary earner's labor supply and the secondary earner's wage rate is substantially flatter (left panel). Since the primary earner's hourly earnings are relatively high, wage rate shocks for the secondary earner can be smoothed out by relatively small adjustments of the primary earner's labor supply.

Finally, we investigate the policy functions for consumption which is most closely related to the household's willingness to borrow. Figure 3 shows consumption as a function of the primary earner's wage rate (left panel) and the secondary earner's wage rate (right panel), respectively. The dashed lines in both panels show that wage rate changes affect consumption relatively little when the household is wealth-rich, as those households can smooth consumption through dissaving. By contrast, if a wealth-poor household faces low wage rate realizations for the primary earner (solid line in the left panel), consumption falls more strongly - the consumption policy has a downward kink . Although the household raises labor supply to cushion the drop in consumption (see the solid lines in the left panels of Figures 1 and 2), it does not smooth consumption perfectly but rather balances the marginal disutilities of more labor and less consumption. Hence, in this scenario, the household is substantially constrained by the impossibility to borrow. The household has a strong wish to borrow and to smooth consumption but, since it cannot borrow, this drives up the shadow price on the borrowing constraint, ϕ.

Importantly, these effects are less pronounced in response to negative wage rate realizations of the secondary earner. The right panel of Figure 3 shows that fluctuations in the secondary earner's wage rate affect consumption only moderately, even if the household is borrowing-constrained. Also wealth-poor households are able to smooth out the consumption effects of wage rate changes for the secondary earner, by increasing labor supply. Since, here, substantial consumption smoothing can be achieved through relatively small increases in the primary earner's labor supply, the household's willingness to borrow is limited.

Thus, consumption is most strongly correlated with the primary earner's wage rate realizations in wealth-poor households. Since it is the marginal utility of consumption that determines the household's willingness to borrow, see equation (9), the strongest correlation between the shadow price on the borrowing constraint and wage rates occurs for primary earners, as we have also shown analytically for a simplified model version in Section 2.3. This has important consequences for standard labor-supply regressions: the (downward) bias in the estimated Frisch elasticity arising from borrowing constraints is more substantial for primary earners than for secondary earners.

4 Estimating Frisch labor-supply elasticities from synthetic data

In this section, we use the calibrated model to simulate synthetic household panel data. From this data, we estimate standard labor-supply regressions to quantify the differential estimation biases for primary and secondary earners, respectively. We do so by estimating labor-supply regressions for different subsamples of the simulated data.

4.1 Sample selections

For the Monte-Carlo estimations, we use the incomplete markets model outlined in Section 2.1 to simulate a synthetic panel data set with similar size as the PSID. Specifically, we simulate $N = 2500$ households for a long period of time and keep the final $T = 6$ periods of this simulation, corresponding to the average number of years for which we observe a household in the PSID.

We quantify the differential estimation biases for primary and secondary earners by estimating the regression (19) on different subsamples of the simulated data. We consider alternative ways of identifying an individual's earner status in the household. Under a first approach, we discriminate between primary and secondary earners by conditioning on the fixed effect in the wage process. Given the baseline calibration based on gender-specific empirical information, this selection is tantamount to estimating the Frisch elasticities separately for men (household members 1) and women (household members 2), respectively. Yet, in the regressions based on real-world data presented in Section 2.1, we cannot oberve perfectly the long-run wage potentials of individuals. Therefore, we apply a second approach where we use both spouses' realized mean wage rates in the sample. Specifically, for both household members, we calculate the mean wage rate in the stochastic simulation in the final $T = 6$ periods, respectively, and classify the spouse with the higher mean wage rate as the primary earner and the spouse with the lower mean wage as the secondary earner. Thus, in this specification, the assigned earner status does not necessarily coincide with long-run wage potentials and, hence, gender. Put differently, the sample of primary earners in this specification also comprises some individuals who are secondary earners in the long run. Yet, when estimating the Frisch elasticities from real-world instead of simulated data, an estimation for women may be subject to additional complications next to borrowing constraints, such as how to handle non-participation and that also true Frisch elasticities may differ, e.g., due to differences in home production across spouses. To address this point, we consider a third approach of identifying an individual's earners status by estimating separate labor-supply equations for primary and secondary earners (identified through mean realized wages) *within gender*. For comparability with the real-world estimations, we proceed analogously with the simulated data.

We first consider the case where there are no differences in the true Frisch elasticity between household members. This assumption helps to illustrate the mechanisms discussed before, as any differences in the estimated labor-supply elasticities that we find in different subsamples of the population must be due to estimation biases. We will also consider a specification where the true Frisch elasticity is larger for women than for men. Note that, in all cases, the true Frisch elasticities for both household members are given by the preference parameters η_m and η_f, see equations (10) and (11). Next to the point estimates, we also report the percentage bias of the estimated Frisch elasticity compared to its true value.

4.2 Estimation results from synthetic data

Table 2 summarizes a first set of estimation results, where we assumed that the true value for the Frisch elasticity is identical for men and women, $\eta_m = \eta_f = 0.9$. The first column shows the results for a specification where we pooled all individuals. In the pooled sample, the

Table 2: Estimation results from synthetic household panel data.

	(1) pooled	(2) $\psi_i > \psi_{-1}$ (prim.)	(3) $\psi_i < \psi_{-i}$ (sec.)	(4) $\overline{w}_i > \overline{w}_{-i}$ (prim.)	(5) $\overline{w}_i < \overline{w}_{-i}$ (sec.)
constant	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
log wage change	0.66	0.37	0.72	0.53	0.72
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)
Observations	30000	15000	15000	15000	15000
True Frisch elasticity	0.90	0.90	0.90	0.90	0.90
Bias (%)	-0.27	-0.58	-0.20	-0.41	-0.20
mean$(\psi_i - \psi_{-i})$	0	0.26	-0.26	0.15	-0.15
percent $\psi_i > \psi_{-i}$	50	100	0	78	22
percent $\overline{w}_i > \overline{w}_{-i}$	50	78	22	100	0

Notes: Second-stage results from a 2SLS regression of equation (19). Standard errors in parentheses. $\overline{w}_i$ is an individual's mean realized wage rate over a sample of $T = 6$ periods. $\overline{w}_{-i}$ is the mean realized wage rate of the spouse over the same period.

estimate for the Frisch elasticity is subject to a downward bias amounting to 27%. This bias reflects the correlation of the expected wage change with the multiplier ratio ϕ/λ as argued before. Column (2) repeats the estimation for primary earners, which, due to our calibration, can be interpreted as a sample of men.[12] In line with our previous discussion, we find that the downward estimation bias particularly strong in this sample and amounts to 58%. Our previous analysis has shown that this estimation bias originating from borrowing constraints should be less pronounced in a sample of secondary earners. This is confirmed in the third column of Table 2, where we repeat the estimation for secondary earners, who, due to the calibration of the model, can be understood as women. In this sample, the downward bias is 20% and, hence, substantially smaller compared to the sample of primary earners displayed in column (2). The differential importance of the estimation bias between household members reflects the differential importance of wage-rate fluctuations for a household's willingness to borrow. As discussed above, the shadow price on the borrowing constraint is less correlated with the secondary earner's expected wage rate change, resulting in smaller estimation biases for secondary earners.

We obtain similar results in the alternative specification where we identify primary and secondary earners by comparing the average realized wages of both spouses in the simulation. Note that this approach of identifying primary and secondary earners is feasible in real-world data while the former approach based on individuals' fixed effects is not, since the fixed effects are not readily observable. For consistency, we therefore consider this approach in the synthetic data as well. In line with our previous results, the fourth and fifth columns in Table 2 show that the estimation bias is larger for primary earners. Here, it is about twice as large in the sample of primary earners than in the sample of secondary earners. Note that the estimation bias among primary earners is larger in column (2) than in column (4). This difference reflects that it is the *potential* wage that is decisive for the size of the bias, as the instrument in the IV regression is the expected wage change. In the sample considered in column (2), a given high wage level is associated with smaller expected future wage decreases on average, as this is the sample of individuals having a high long-run wage. In turn, this makes the estimation bias particulary strong in this sample.

In summary, the regressions from synthetic data illustrate that better estimates for the Frisch elasticity can be obtained from samples of secondary earners. Yet, from a practical

[12]Note that this accords to the usual sample restriction in the applied literature, where Frisch elasticities are typically estimated from data on married men only. The sample selection on men is also comparable to the one-person household analysis performed by Domeij and Flodén (2006), where each individual is a primary earner by construction.

point of view, a challenge is that secondary earners tend to be women, and real-world labor-supply regressions for women might be subject to several complications next to how to deal with borrowing constraints, in particular how to deal with non-participation and differences in home production. In the next section where we estimate labor-supply relations with real-world data, we will address these challenges by also distinguishing between primary and secondary earners *within genders*. The results will show that the estimated labor-supply elasticities differ also between primary and secondary earners of the same gender. These differences can, hence, not only be driven by gender differences in the true elasticities or by gender-specific biases.

Table 3: Estimation results from synthetic household panel data (heterogeneous Frisch elasticities).

	(1) pooled	(2) men	(3) women
constant	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
log wage change	1.04	0.35	1.19
	(0.01)	(0.02)	(0.02)
Observations	30000	15000	15000
True Frisch elasticity		0.75	1.50
Bias (%)		-0.53	-0.21

Notes: Second-stage results from a 2SLS regression of equation (19). Standard errors in parentheses.

As a final evaluation within Monte-Carlo framework, we consider a specification where the true Frisch elasticity differs between men and women. The purpose of this exercise is to compare differences in estimated Frisch elasticities in a setting where there are two potential sources of heterogeneity in the estimates: First, because there is true preference heterogeneity across household members, and, second, because there are spurious differences in the estimates due to the differential importance of estimation biases for primary and secondary earners. In this experiment, we assume that the Frisch elasticity for women is twice as large as men's, specifically we use the parameter values $\eta_m = 0.75$ and $\eta_f = 1.5$. To facilitate comparisons, we re-calibrate the model to match the calibration targets discussed in Section 3.1. In particular, we re-calibrate the interest rate so that the asset distribution has the same properties between both model versions (the same holds for the mean hours ratio between men and women).

Table 3 summarizes the estimation results for the pooled sample (column 1), and the samples of men (column 2) and women (column 3). In both subsamples, the respective percentage bias is similar as in our previous specification without preference heterogeneity. Strikingly, while the true Frisch elasticity for women is only twice as large than the one for men, the difference in the estimated Frisch elasticities is $(1.19 - 0.35)/0.35 = 2.4$, hence, substantially larger. Thus, these results suggest that estimated differences between the Frisch elasticities of population groups based on real-world data, in particular relating to the differences between men and women, tend to be over-estimated. In other words, parts of the previously estimated difference between Frisch elasticities for men and women result from the differential importance of estimation biases for primary and secondary earners.

Table 4: Estimation results from synthetic household panel data (men, conditional on assets).

	(1)	(2)	(3)	(4)	(5)	(6)
	all		assets $> 0.01 \cdot \overline{a}$		assets $> 0.1 \cdot \overline{a}$	
	prim.	sec.	prim.	sec.	prim.	sec.
constant	0.00	-0.01	0.01	-0.01	0.01	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
log wage change	0.33	0.63	0.45	0.66	0.61	0.73
	(0.02)	(0.04)	(0.02)	(0.05)	(0.02)	(0.05)
Observations	11742	3258	9403	3002	7115	2724
Ratio $\hat{\eta}_2/\hat{\eta}_1$	1.91		1.46		1.20	

Notes: Second-stage results from a 2SLS regression of equation (19). Standard errors in parentheses. $\overline{a}$ denote mean asset holdings in the economy.

5 Estimating labor-supply elasticities from PSID data

In this section, we perform labor-supply regressions using household panel data from the Panel Study of Income Dynamics (PSID). As in the previous section, we estimate a labor-supply condition in first differences using instrumented wage rates which would deliver unbiased estimates for the Frisch elasticity in the absence of borrowing constraints. If Frisch elasticities in the real world were identical between primary and secondary earners, we would expect a smaller point estimate for the Frisch elasticity in samples of primary earners compared to samples of secondary earners, due to the more pronounced downward bias for the former group of individuals. However, when using real-world data, the assumption of homogeneous

Frisch elasticities across population groups may be violated. Therefore, a sharper test of the mechanism outlined in this paper is to investigate the *differences* in estimated Frisch elasticities in different subsamples of the data. Our theory predicts that, everything else equal, the percentage difference in estimated labor-supply elasticities between primary and secondary earners should increase, the more important are borrowing constraints in the data, as documented in Table 4 based on synthetic data. In the following, we will therefore investigate both, absolute as well as relative differences in estimates for the Frisch elasticity between primary and secondary earners in different subsamples.

When estimating the labor-supply regression (19) from empirical data, we use household characteristics to instrument for future wage changes, Δw_{ijt+1} (the mathematical expectation of future wages as an instrument is only available in synthetic data). Specifically, we follow e.g., MaCurdy (1981) and Domeij and Flodén (2006) and use as instruments age, age squared, years of schooling, years of schooling squared, and an interaction term between age and years of schooling.[13] As discussed by Altonji (1986) and Domeij and Flodén (2006), these instruments are potentially weak, as reflected by small F statistics in the first-stage regression of the 2SLS-estimator. Stock, Wright, and Yogo (2002) suggest that, when there is one endogenous regressor, inference based on the 2SLS estimator is reliable when the F statistic exceeds 10. In our baseline estimations with a relatively large number of observations, this condition is satisfied. To address potential concerns of weak instruments, we also present results where we estimate the labor-supply regression using limited information maximum likelihood (LIML) instead of 2SLS.

In real-world data, individuals's long-run wage potentials are not perfectly observable. Put differently, there is no perfect counterpart to the fixed effects ψ observable in the real-world data. We apply two sample selections to identify samples of primary or secondary earners in the data. These selections are similar to those applied in Table 2 for the synthetic data. One approach relies on observed wage rates over the sample period, the other one on long-run characteristics of the individuals. For the former approach, we apply the same restriction on realized wage rates as in Table 2, i.e., the person with the higher average wage over the sample period is classified as the primary earner in the household. For the latter approach, we apply sample splits by gender, a variable unrelated to the realizations

[13]In their preferred specifications, Domeij and Flodén (2006) use an alternative wage variable (a reported hourly wage rate) as instrument, which avoids a potential bias implied by measurement errors in the wage data. In our application, we can not use this instrument, as, in addition to the sample selection criteria imposed by Domeij and Flodén (2006), our analysis requires wage rates to be observed for both partners (and not only for one partner). An additional selection on workers who are paid reported hourly rates would result in a sample that is too small to allow reliable inference.

of the wage rate *shocks* but related to the long-run means. Empirically, the sample of men contains mostly primary earners but also some men who earn less than their respective wives. In turn, also the sample of individuals with higher average wage rate realizations comprises mostly primary earners but also some secondary earners who happened to realize particularly positive wage shocks relative to their spouses such as to appear as primary earners in the data. In any case, both samples comprise mostly primary earners and are rather similar with an overlap of 79%.

Table 5 summarizes the results of the second-stage regressions and also reports the F statistics of the first stage as well as the Sargan test of overidentifying restrictions.[14] The F statistics indicate that the household variables have significant explanatory power for future wage changes. The Sargan test rejects the null hypothesis of instrument validity only for the specification where we pool all individuals which is a specification that we only show for completeness.

Table 5: Empirical labor-supply regressions, PSID data.

	(1) pooled	(2) men	(3) women	(4) $\overline{w}_i > \overline{w}_{-i}$ (prim.)	(5) $\overline{w}_i < \overline{w}_{-i}$ (sec.)
constant	-0.01	-0.01	-0.00	-0.01	-0.01
	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)
log wage change	0.66	0.41	0.76	0.41	0.83
	(0.13)	(0.13)	(0.23)	(0.13)	(0.24)
Observations	28,700	14,350	14,350	14,350	14,350
F statistic	30.01	11.11	18.41	19.58	19.58
Sargan test (pval)	0.0234	0.219	0.268	0.359	0.317
mean($\ln \overline{w}_i - \ln \overline{w}_{-i}$)	0	0.40	-0.41		
percent men	50	100	0	79	21
percent $\overline{w}_i > \ln \overline{w}_{-i}$	50	79	21	100	0

Notes: Second-stage results from a 2SLS regression of equation (19). Standard errors in parentheses. $\overline{w}_i$ is an individual's mean realized wage rate in the sample. $\overline{w}_{-i}$ is the mean realized wage rate of the spouse.

The point estimate for the Frisch elasticity in the pooled sample is 0.66. In line with

[14] The null hypothesis of the Sargan test is that the instruments are uncorrelated with the error term and that the excluded instruments are correctly excluded from the estimated equation. Thus, a rejection would cast doubt on the validity of the instruments.

our theoretical analysis, we find that the estimated Frisch elasticity is larger in the sample of women (who are mainly secondary earners) compared to the sample of men (who are mainly primary earners), see columns (2) and (3). Yet, when using empirical data, we have to provide additional evidence in order to show that that these differences are indeed related to differences in earners status between men and women and do not primarily pick up gender differences in the true Frisch elasticities. We will do so by presenting within-gender estimates, as well as by conditioning on household asset holdings. Before, columns (4) and (5) repeat the estimation for a sample selection where the earners status is determined by comparing the average wage rates of husbands and wives in a marriage. In line with our estimations based on synthetic data (compare Table 2), we find that in samples where individuals have higher average wage rates than his/her spouse, the estimate for the Frisch elasticity is substantially smaller (about 50%) than in a sample of mostly secondary earners, compare columns (4) and (5).

Table 6: Empirical labor-supply regressions, PSID data (within gender).

	(1)	(2)	(3)	(4)
	prim.	sec.	prim.	sec.
	(wihtin	men)	(within	women)
constant	-0.01	-0.02	0.01	-0.01
	(0.00)	(0.02)	(0.02)	(0.01)
log wage change	0.35	0.74	0.26	0.87
	(0.14)	(0.46)	(0.41)	(0.26)
Observations	11,640	2,710	2,710	11,640
F statistic	11.77	16.48	2.198	2.167
Sargan test (pval)	0.256	0.851	0.724	0.390

Acknowledging potential gender differences in true Frisch elasticities and that an empirical estimation for women may be subject to additional biases next to the ones resulting from borrowing constraints, we now turn to the empirical counterpart to Table **??**, where we discriminated between primary and secondary earners within genders. In these samples, differences in estimated coefficients between primary and secondary earners can not result from gender differences in true elasticities or any gender-specific (rather than earner-role-specific) biases. Columns (1) and (2) of Table 6 show that, also in the PSID data, the estimated Frisch elasticity is larger in a sample of male secondary earners compared to a sample of male primary earners. In fact, the point estimate for male secondary earners is

0.74, hence, about twice as large as the one for primary earners. When we repeat the within-gender evaluation for women, we also find that the point estimate is substantially larger for secondary than for primary earners.

As a robustness check, we re-estimated the specifications shown in Tables 5 and 6 using limited information maximum likelihood (LIML). Stock, Wright, and Yogo (2002) and Poi (2006) have shown that this estimator may outperform the 2SLS estimator when instruments are weak. Appendix C provides detailed estimation results. Overall, we find that the differences in the estimates for primary and secondary earners are even larger when we use LIML instead of 2SLS.

As a final step, we address the possibility that there are differences in the true Frisch elasticities between primary and secondary earners even within the same gender. For example, a male secondary earner may contribute more to home production than a male primary earner, leading to a higher Frisch elasticity (Alesina et al., 2011; Bredemeier, 2015). To address this point, we extend our empirical analysis to take into account information on household wealth. Borrowing constraints are less important in wealth-rich households. As a consequence, our theory implies that, even if there are differences in the true Frisch elasticities between primary and secondary earners, the *differences* in the estimated elasticities should be larger in wealth-poor households where the differences in the true elasticities are magnified by differently strong biases.

As discussed by Domeij and Flodén (2006), the 1984, 1989, and 1994 PSID waves contain a supplement on household wealth.[15] For the remaining analysis, we therefore have to restrict the sample to these years. Specifically, we follow Domeij and Flodén (2006) and pool three subsamples of the PSID around the years where the data on household wealth is available: 1983-1985, 1988-1990, and 1993-1995.[16] We also follow Domeij and Flodén (2006) in restricting the sample to male household heads. Within this sample, we discriminate between primary and secondary earners by comparing the average wage rate of the household head to the wife's average wage rate.

In line with the estimation results reported by Domeij and Flodén (2006), column (1) in Table 7 shows that the estimated Frisch elasticity in the full sample (i.e., where we do not control for borrowing constraints) is relatively small. Yet, columns (2) and (3) show that the pooled estimation hides substantial differences in the point estimates between primary and secondary earners, in line with the mechanism discussed in this paper: The point estimate

[15] From 1999 onwards, some of the required asset and liabilities components have been added to the regular set of survey items in the PSID, but this is also the period where the PSID changed to biennial interviews.

[16] For details on the subsamples we refer to Domeij and Flodén (2006).

Table 7: Empirical labor-supply regressions for male household heads, PSID data (conditional on liquid assets).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	full sample			assets > 0			assets > 1000		assets > 2000	
	all	prim.	sec.	all	prim.	sec.	prim.	sec.	prim.	sec.
constant	0.01	0.01	-0.01	0.00	0.00	-0.02	0.01	0.00	0.00	0.00
	(0.01)	(0.00)	(0.03)	(0.01)	(0.01)	(0.05)	(0.01)	(0.02)	(0.01)	(0.02)
log wage change	0.29	0.09	1.02	0.60	0.32	1.77	0.21	0.53	0.26	0.33
	(0.21)	(0.15)	(1.36)	(0.38)	(0.27)	(2.74)	(0.25)	(0.66)	(0.30)	(0.67)
Observations	7746	5866	1880	5665	4324	1341	3826	1197	3534	1108
F statistic	4.81	6.73	0.30	2.11	2.67	0.16	2.52	0.69	1.93	0.49
Sargan test (pval)	0.00	0.00	0.94	0.30	0.13	0.99	0.19	0.83	0.16	0.71
Ratio $\hat{\eta}_2/\hat{\eta}_1$		10.95			5.52		2.53		1.23	

from the sample of primary earners is substantially smaller than the point estimate from the sample of secondary earners. However, it should be acknowledged that the estimates are very imprecise and the F statistic points towards instruments being weak in this sample.[17]

As discussed before, our theory predicts that the difference in the point estimates between primary and secondary earners gets smaller, the less relevant are borrowing constraints. To investigate this hypothesis, we follow Domeij and Flodén (2006) and exclude households that are likely to be borrowing constrained. In a first experiment, we restrict the analysis to households having non-negative liquid assets. The results for the pooled sample of primary and secondary earners with non-negative liquid assets, see column (4), are in line with the results of Domeij and Flodén (2006): In a sample where borrowing constraints are less binding, the estimated Frisch elasticity is about twice as large compared to the full sample (compare column 1), arguably due to a less pronounced downard bias in the estimation. What is more important for our analysis is, however, that the *difference* between the estimated Frisch elasticities for primary and secondary earners becomes smaller when we condition on asset holdings. In the full sample without controlling for borrowing constraints, the ratio $\hat{\eta}_2/\hat{\eta}_1$ was about 11, while it is only 5.5 when we exclude households with negative liquid assets. To further illustrate this relation, we restrict the sample by requiring that household liquid assets are larger than \$1000 (in 1983 dollars). In line with our mechanism, the difference between both estimates decreases to $\hat{\eta}_2/\hat{\eta}_1 = 2.53$. Finally, for an even more restricted sample

[17]Domeij and Flodén (2006) find a similar pattern. They estimate Frisch elasticities from a similar sample as we use in Table 7 and also find that the estimates tend to be insignificant.

where we exclude all households with less than $2000 liquid assets, this ratio drops to 1.23. Hence, the relation between difference in estimated labor-supply elasticities and the level of household wealth is in line with our theoretical predictions in general and, in particular, the quantitative evaluations in Table 4 where we analyzed the role of assets in our model.

6 Conclusion

Estimates of Frisch labor-supply elasticities are biased in the presence of borrowing constraints. We have shown that this estimation bias is of differential importance among household members. In couples with joint borrowing constraints, wage-rate fluctuations of secondary earners are less important for the couples' willingness to borrow which results in smaller estimation biases for secondary earners.

We have presented an incomplete markets model with two earners to make this point explicit. A direct empirical implication of our model is that, other things equal, better estimates of the Frisch elasticity can be obtained from samples consisting of secondary earners, while the downward estimation bias is most substantial in samples of primary earners. In an empirical application using PSID data, we have found support for these predictions.

More generally, our results imply that labor-supply reactions differ by household wealth, relative wage positions within the household, and their interaction. Shocks and policies which change net real wage rates have relatively similar effects on primary and secondary earners in rich households but effects are more different among household members in wealth-poor households. In other words, in wealth-poor households, reactions to wage rate fluctuations are not solely governed by Frisch elasticities. Our study also points towards policy interdependencies. Policies that affect the importance of borrowing constraints, e.g. saving subsidies or public loan guarantees, affect the distribution of labor-supply reactions to policies that affect net wage rates, such as temporary tax cuts.

References

Alogoskoufis, G. S. (1987). On intertemporal substitution and aggregate labor supply. *Journal of Political Economy 95*(5), 938–60.

Altonji, J. G. (1986). Intertemporal substitution in labor supply: Evidence from micro data. *Journal of Political Economy 94*(3), s176–S215.

Blomquist, N. S. (1985). Labour supply in a two-period model: The effect of a nonlinear progressive income tax. *Review of Economic Studies 52*(3), 515–24.

Blomquist, N. S. (1988). Nonlinear taxes and labor supply. *European Economic Review 32*(6), 1213–1226.

Blundell, R. and T. MaCurdy (1999). *Labor Supply: A Review of Alternative Approaches*, Volume 3A of *Handbook of Labor Economics*, Chapter 27, pp. 1559–1695. Elsevier.

Bredemeier, C. (2015). Household Specialization and the Labor-Supply Elasticities of Women and Men. University of Cologne, Working Paper in Economics 81.

Diaz-Gimenez, J., A. Glover, and J.-V. Rios-Rull (2011). Facts on the distributions of earnings, income, and wealth in the United States: 2007 update. *Federal Reserve Bank of Minneapolis Quartely Review 34*(1), 2–31.

Domeij, D. and M. Flodén (2006). The labor-supply elasticity and borrowing constraints: Why estimates are biased. *Review of Economic Dynamics 9*, 242262.

Heckman, J. J. (1993). What has been learned about labor supply in the past twenty years? *American Economic Review 83*(2), 116–21.

Imai, S. and M. P. Keane (2004). Intertemporal labor supply and human capital accumulation. *International Economic Review 45*(2), 601–641.

Keane, M. P. (2011). Labor Supply and Taxes: A Survey. *Journal of Economic Literature 49*(4), 961–1075.

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

Ortigueira, S. and N. Siassi (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics 60*, 650–666.

Poi, B. P. (2006). Jackknife instrumental variables estimation in Stata. *Stata Journal 6*(3), 364–376.

Rupert, P., R. Rogerson, and R. Wright (2000). Homework in labor economics: Household production and intertemporal substitution. *Journal of Monetary Economics 46*(3), 557–579.

Stock, J. H., J. H. Wright, and M. Yogo (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics 20*(4), 518–29.

A Appendix to Section 2

For notational convenience we define $Z_i = \exp z_i$ and $\Psi_i = \exp \psi_i$, $\Psi_1 = 1$, $\Psi_2 < 1$. Without loss of generality, we normalize $\psi_1 = 0$, hence $\psi_2 < 0$, and $\alpha = 1$.

Proof of Proposition 1. Consider a period is entered with state variables a, $w_1 = Z_1$, and $w_2 = \Psi_2 \cdot Z_2$. By the FOC for consumption (9), the multiplier ratio ϕ/λ is given by

$$\frac{\phi}{\lambda} = \frac{\lambda/(1+r) - \beta E\lambda'}{\lambda} = \frac{1}{1+r} - \frac{\beta \cdot E\lambda'}{\lambda}.$$

Obviously, this is zero if the household is not borrowing constrained in this period, i.e. if $a' > 0$. Since z_2 and z_1 are i.i.d., a borrowing constrained household expect to enter the next period with state variables $a' = z_1' = z_2' = 0$. Hence, expectations for the next period do not depend on current states in case of being borrowing constrained and we can consider $E\lambda'$ in case of being borrowing constrained as a constant, which we denote as $\overline{\lambda}$.

Taken together, we can express the multiplier ratio as

$$\frac{\phi}{\lambda} = \max\left[\frac{1}{1+r} - \beta \cdot \frac{\overline{\lambda}}{\widetilde{\lambda}}, 0\right]$$

and it is hence weakly increasing in $\widetilde{\lambda}$, $\partial\left(\phi/\lambda\right)/\partial\widetilde{\lambda} \leq 0$.

In case the household is borrowing constrained, the Lagrange multiplier on being borrowing constrained can be obtained from the remaining first order conditions for the current period:

$$\begin{aligned}
n_2^{1/\eta} &= \widetilde{\lambda} \cdot w_2 = \widetilde{\lambda} \cdot Z_2 \cdot \Psi_2, \\
n_1^{1/\eta} &= \widetilde{\lambda} \cdot w_1 = \widetilde{\lambda} \cdot Z_1, \\
\widetilde{\lambda} &= c^{-\sigma}, \\
c &= w_2 \cdot n_2 + w_1 \cdot n_1 + a \\
&= Z_2 \cdot \Psi_2 \cdot n_2 + Z_1 \cdot n_1 + a,
\end{aligned}$$

where the final condition uses $a' = 0$. Combining all conditions gives

$$\left(\widetilde{\lambda}\right)^{-1/\sigma} - \Lambda \cdot \widetilde{\lambda}^{\eta} - a = 0,$$

where $\Lambda = W_1^{1+\eta} + W_2^{1+\eta} = Z_2^{1+\eta} \cdot \Psi_2^{1+\eta} + Z_1^{1+\eta}$. Defining the left-hand sind of the above as F and applying the implicit-function theorem gives

$$\frac{\partial \widetilde{\lambda}}{\partial \Lambda} = -\frac{\partial F/\partial \Lambda}{\partial F/\partial \widetilde{\lambda}} = -\frac{-\widetilde{\lambda}^{\eta}}{-\left(\frac{1}{\sigma}\left(\widetilde{\lambda}\right)^{-1/\sigma-1} + \Lambda \eta \widetilde{\lambda}^{\eta-1}\right)} < 0.$$

Together with $\partial\left(\phi/\lambda\right)/\partial\widetilde{\lambda} \leq 0$ from above, this gives $\partial\left(\phi/\lambda\right)/\partial\Lambda \leq 0$. ■

Proof of Proposition 2. Combining $\partial\left(\phi/\lambda\right)/\partial\Lambda \leq 0$, and $\partial\Lambda/\partial z_i > 0$ gives $\partial\left(\phi/\lambda\right)/\partial z_i \leq 0$ and $corr\left(\phi/\lambda, z_i\right) \leq 0$.

Now, consider the covariance between Λ and individual wage shocks. For each household member, $cov\left(\Lambda, Z_i\right) = cov\left(\Psi_2^{1+\eta} Z_2^{1+\eta} + Z_1^{1+\eta}, Z_i\right) = \Psi_2^{1+\eta} \cdot cov\left(Z_2^{1+\eta}, Z_i\right) + cov\left(Z_1^{1+\eta}, Z_i\right)$. Therefore, $cov\left(\Lambda, Z_2\right) = \Psi_2^{1+\eta} \cdot cov\left(Z_2^{1+\eta}, Z_2\right) + cov\left(Z_1^{1+\eta}, Z_2\right) = \Psi_2^{1+\eta} \cdot cov\left(Z_i^{1+\eta}, Z_i\right)$ and $cov\left(\Lambda, Z_1\right) = cov\left(Z_i^{1+\eta}, Z_i\right)$. Since $Z_i > 0$, $cov\left(Z_i^{1+\eta}, Z_i\right) > 0$ and, hence, $cov\left(\Lambda, Z_1\right) > cov\left(\Lambda, Z_2\right)$ since $\Psi_2^{1+\eta} < 1$.

Since wage shocks are i.i.d., there is a zero correlation between the current wages and the beginning-of-period asset level a which was determined in the previous period. The following is performed conditional on any a.

Generally, it holds that $cov\left(f\left(x\right), y|a\right) = E\left(\partial f/\partial x|a\right) \cdot cov\left(X, Y|a\right)$. Define f as the multiplier ratio ϕ/λ as a function of Λ.

$$cov\left(f\left(\Lambda\right), Z_i|a\right) = E\left(\partial\left(\phi/\lambda\right)/\partial\Lambda|a\right) \cdot cov\left(\Lambda, Z_i|a\right).$$

The fact that $\partial\left(\phi/\lambda\right)/\partial\Lambda \leq 0$ implies that the first factor on the right hand side is negative but it does not depend on i. Since $cov\left(\Lambda, Z_1|a\right) \geq cov\left(\Lambda, Z_2|a\right)$ as $\Psi_2 < 1$, it follows that $cov\left(\phi/\lambda, Z_2|a\right) \geq cov\left(\phi/\lambda, Z_1|a\right) \iff \Psi_2$. Together with $cov\left(\phi/\lambda, Z_i|a\right) \leq 0$ it follows that

$$cov\left(\phi/\lambda, Z_1|a\right) \leq cov\left(\phi/\lambda, Z_2|a\right) \leq 0$$

and this holds with strict inequality for those a that do not rule out being borrowing constrained at the end of the period.

Now, aggregating over the different a gives $cov\left(\phi/\lambda, Z_g\right) = E\left(\phi/\lambda \cdot Z_g\right) - E\left(\phi/\lambda\right) E\left(Z_g\right) = \int E\left(\phi/\lambda \cdot Z_g|a\right) \quad df\left(a\right) - \int E\left(\phi/\lambda|a\right) E\left(Z_g|a\right) \quad df\left(a\right) =$

$\int cov\,(\phi/\lambda, Z_g|)\, a\ df\ (a)$, such that

$$cov\,(\phi/\lambda, Z_1) < cov\,(\phi/\lambda, Z_2) < 0$$

Due to $var\,(Z_1) = var\,(Z_2)$, the same holds for the correlations. $Z_i \approx 1 + z_i$ and, hence, $corr\,(\phi/\lambda, Z_g) \approx corr\,(\phi/\lambda, z_g)$, completes the proof. ■

Proof of Proposition 3. Since the z_i are i.i.d.,

$$corr(\phi/\lambda, \Delta \ln w_i') = corr(\phi/\lambda, z_i' - z_i) = corr(\phi/\lambda, z_i') - corr(\phi/\lambda, z_i)$$

and

$$corr(\phi/\lambda, z_i') = 0,$$

hence,

$$corr(\phi/\lambda, \Delta \ln w_i') = -corr(\phi/\lambda, z_i).$$

Also,

$$E\Delta \ln w_i' = E \ln w_i' - \ln w_i = Ez_i' - z_i = -z_i.$$

Therefore,

$$corr\left(\phi/\lambda, E\Delta \ln w_i'\right) = corr\,(\phi/\lambda, -z_i) = -corr\,(\phi/\lambda, z_i)\,.$$

Together with Proposition 2, this completes the proof. ■

B Appendix to Section 3

We estimate the parameters of the wage processes separately for men and women. To save on notation, we do not account for a gender index $g = m, f$ in the following. We assume that the observed log wage rate of individual i, w^*_{it}, is composed of three components: a deterministic component $f(o_{it})$, a permanent component α_i and a transitory component $\tilde{w}_{it}$. By including a deterministic component we filter out observable cross-sectional variation in wage rates. The permanent component captures individual-specific fixed characteristics. The transitory component captures temporary shocks affecting individual's wage rates. We allow for persistence in the transitory shocks and model the transitory component as an ARMA(1,1) process, which has been shown to be an accurate description of earnings dynamics. In line with the empirical literature on earnings processes, we account for time effects in both the permanent and the transitory wage component. In sum, the wage process (for gender $g = m, f$) is given by:

$$\begin{aligned} w^*_{it} &= f(o_{it}) + \tilde{w}_{it}, \\ \tilde{w}_{it} &= p_t \alpha_i + \lambda_t w_{it}, \\ w_{it} &= \rho w_{it-1} + \theta \varepsilon_{it-1} + \varepsilon_{it}, \end{aligned} \tag{20}$$

where $E(\alpha_i) = E(\tilde{w}_{it}) = E(w_{it}) = 0$ and o_{it} are observable characteristics of individual i at time t. $\tilde{w}_{it}$ is the stochastic component of individual's wages (residual wages). p_t and λ_t are factor loadings that allow the permanent and transitory components, respectively, to change over time in a way that is common across individuals. Serial correlation in the transitory shocks $\tilde{w}_{it}$ is modelled using an ARMA(1,1) process with autocorrrelation ρ and MA parameter θ. ε_{it} is a random variable with variance σ^2_ε. Concerning initial conditions, we follow the approach by MaCurdy (1982) and treat the variance at the start of the sample period, σ^2_{w1} as an additional parameter to be estimated.

In our theoretical model, the wage processes for men and women are given by simple autoregressive processes, taking into account a gender gap in mean wages. To parameterize these processes (gender-specific autocorrelation and gender-specific innovation variances), we use the estimated values for ρ and σ^2_ε from the gender-specific estimation of the time-series model described above. In the PSID data, the mean wage of women amounts to 67% of the mean wage of men. We use this statistic to calibrate the wage gap in our model by setting $\psi_1 > \psi_2$.

We estimate (20) separately for men and women using a two step procedure. First, we estimate the deterministic component, $f(o_{it})$, by running an OLS regression of gender-specific log wage rates on time dummies, age dummies, and schooling dummies interacted up to a quadratic age trend. We drop all observations where the residual of this regression belongs to the bottom or top 0.5 percent of all residuals for an age group. We then re-estimate the first-stage regression to obtain the final vector of residual wage rates, $\tilde{w}_{it}$, for which we estimate the covariance structure as described above using a Generalized Method of Moments (GMM) procedure.

The true variance-covariance matrix of residual wages $\tilde{w}_{it}$ has diagonal elements

$$\sigma_1^2 = p_1^2\sigma_\alpha^2 + \lambda_t^2\sigma_{w1}^2$$
$$\sigma_t^2 = p_t^2\sigma_\alpha^2 + \left\{\lambda_t^2\left(\rho^{2t-2}\sigma_{w1}^2 + K\sum_{j=0}^{t-2}\rho^{2j}\right)\right\}, t > 1$$

and off-diagonal elements

$$Cov(\tilde{w}_t, \tilde{w}_{t+s}) = p_t p_{t+s}\sigma_\alpha^2 + \lambda_t\lambda_{t+s}\left(\rho^s\sigma_{w1}^2 + \rho^{s-1}\theta\sigma_\varepsilon^2\right), t = 1 \text{ and } s > 0$$
$$Cov(\tilde{w}_t, \tilde{w}_{t+s}) = p_t p_{t+s}\sigma_\alpha^2 + \lambda_t\lambda_{t+s}\left(\rho^{2t+s-2}\sigma_{w1}^2 + \rho^2 K\sum_{j=0}^{t-1}\rho^{2j} + \rho^{s-1}\theta\sigma_\varepsilon^2\right), t > 1 \text{ and } s > 0,$$

where $K = \sigma_\varepsilon^2\left(1 + \theta^2 + 2\rho\theta\right)$.

GMM estimation is carried out by replacing population moment conditions by their sample analogues. For both genders, the parameter vector to be estimated is

$$\varphi = \left\{\sigma_\alpha^2, \rho, \sigma_{w1}^2, \sigma_\varepsilon^2, \theta, \lambda_2 ... \lambda_T, p_2 ... p_T\right\}.$$

For identification, the factor loadings p_1, λ_1 are set equal to one. In our unbalanced panel data set, each sample moment is constructed using all available observations for that moment. Following Haider (2001), we adjust the standard erros of the parameter estimates to take into account the number of observations used in the computation for each moment. We follow Altonji and Clark (1996) and Clark (1996) and use the identity matrix as the weighting matrix, which has been shown to lead to better small sample performance than the optimal weighting matrix.

Table 8 summarizes the parameter estimates.[18] For men, the autocorrelated shocks to idiosyncratic wages have an annual autocorrelation of $\rho_m = 0.82$ and a standard deviation of $\sigma_{m,\varepsilon} = 0.22$. These results are well in line with the literature, see, e.g., Domeij and Flodén

[18] To save on space, the point estimates for the time-varying factor loadings are not shown. In line with the literature, we find that idiosyncratic labor market risk tends to increase over time. Results for the first-stage filter regression are available on request.

(2006) and the references cited therein. For women, the estimated autocorrelation is with $\rho_f = 0.82$ similar compared to men's, and the estimated degree of idiosyncratic labor market risk is with $\sigma_{f,\varepsilon} = 0.22$ about twice as large as for men.

Table 8: Estimated wage processes for men and women.

	(1) men	(2) women
σ^2_α	0.06	0.05
	(0.02)	(0.02)
ρ	0.84	0.82
	(0.03)	(0.03)
σ^2_{w1}	0.13	0.23
	(0.02)	(0.02)
σ^2_ε	0.05	0.21
	(0.01)	(0.03)
# moments	351	351

C Appendix to Section 4

Tables 9 and 10 are the counterparts to tables 5 and 6, estimated with Limited Information Maximum Likelihood instead of 2SLS. Overall, we obtain similar results. Note that, under the LIML estimator, the estimates for secondary earners increase by more than the ones for primary earners. Thus, the LIML estimator which is more robust to instruments being potentially weak suggests even larger differences in the point estimates between primary and secondary earners than conventional 2SLS.

Table 9: Empirical labor-supply regressions, PSID data (Limited Information Maximum Likelihood).

	(1)	(2)	(3)	(4)	(5)
	pooled	men	women	prim.	sec.
constant	-0.01	-0.01	-0.00	-0.01	-0.01
	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)
log wage change	0.73	0.45	0.84	0.43	0.9
	(-0.14)	(-0.14)	(-0.25)	(-0.14)	(-0.26)
Observations	28,700	14,350	14,350	14,350	14,350
F statistic	30.01	11.11	18.41	19.58	19.58
Sargan test (pval)	0.0234	0.228	0.281	0.365	0.331

Table 10: Empirical labor-supply regressions, PSID data (within gender, Limited Information Maximum Likelihood).

	(1)	(2)	(3)	(4)
	prim.	sec.	prim.	sec.
	(within men)		(within women)	
constant	-0.01	-0.03	0.01	-0.01
	(0.00)	(-0.02)	(0.02)	(0.01)
log wage change	0.39	0.78	0.32	0.93
	(-0.14)	(-0.48)	(-0.45)	(-0.28)
Observations	11,640	2,710	2,710	11,640
F statistic	11.77	16.48	2.198	2.167
Sargan test (pval)	0.265	0.856	0.732	0.402