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

Conventional estimates of Frisch labor-supply elasticities are biased in presence of borrowing constraints. We develop an incomplete-markets model with two-earner households and derive a new estimation approach for the Frisch elasticity that yields unbiased estimates even in samples that include borrowing-constrained households. Our approach exploits that the strength of the estimation bias depends on individuals' relative contribution to household earnings. It takes the form of a simple interaction-term model with minimum data requirements. Using PSID data, we estimate Frisch elasticities of about 0.7 for men and rather homogeneous Frisch elasticities across the population. Our results help to better understand the micro/macro puzzle on the labor-supply elasticity and have rich implications for optimal taxation.

JEL classification: E24, J16, J22, E21

Keywords: Labor-supply elasticity; incomplete markets; double-earner households

We would like to thank Mike Waldman (Editor), two anonymous referees, and conference and seminar participants in Zurich, Antwerp, Cologne, Bochum, Wuppertal, Munich, the 69th European Meeting of the Econometric Society 2016 in Geneva, the IAAE 2016 Annual Conference in Milan, the 22nd International Conference Computing in Economics and Finance 2016 in Bordeaux, and the Annual Meeting of the Verein fuer Socialpolitik in Muenster 2015 for helpful comments and suggestions.

1 Introduction

The Frisch elasticity of labor supply measures the percentage reaction of hours worked to a one percent change in the net wage rate holding the marginal utility of wealth constant. Thus, the Frisch elasticity determines adjustments in labor supply to wage-rate changes that trigger pure intertemporal substitution effects but no accompanying income effects. There are various examples for wage-rate changes that have this property. First, under perfect capital markets, purely transitory wage-rate changes should have no impact on the marginal utility of wealth. Second, if agents are forward-looking and can borrow freely, wage-rate changes that can be expected by the agent in advance should also leave the marginal utility of wealth unchanged. Accordingly, the Frisch elasticity is important for reactions to transitory tax or productivity shocks and to predictable life-cycle patterns in wage rates.[1]

In the literature, there is no consensus on the size of the Frisch elasticity. In fact, the micro and macro view on labor-supply elasticities differ markedly, see, e.g., Keane and Rogerson (2015). While quantitative macroeconomic models tend to require a relatively large value for the Frisch elasticity to match the data well, existing microeconometric studies on the Frisch elasticity typically estimate smaller values for this parameter. The micro/macro puzzle on the Frisch elasticity may be due to a number of estimation biases discussed in the literature, see, e.g., Blomquist (1985), Blomquist (1988), Alogoskoufis (1987), Heckman (1993), Rupert, Rogerson, and Wright (2000), or Imai and Keane (2004).

A particular estimation problem has been highlighted by Domeij and Flodén (2006), who have shown that, in presence of borrowing constraints, conventional methods to estimate the Frisch elasticity are subject to a downward bias. This bias is important since borrowing constraints are a substantial restriction to many households in the U.S. (see, e.g., Diaz-Gimenez, Glover, and Rios-Rull 2011). In this paper, we derive a new estimation approach for the Frisch elasticity that yields unbiased estimates even in samples of potentially borrowing-constrained households. Our approach critically exploits the couple structure of households, i.e., we exploit information from households with two potential earners.[2] Our approach is appealing to the applied researcher as it takes the form of a simple interaction-term regression with minimum data requirements. When we apply our method to household data from the

[1] In macroeconomics, the Frisch elasticity is a key determinant of the size of the fiscal multiplier and the costs of business cycles. The Frisch elasticity is also important in microeconomic applications, where often other elasticity concepts, such as Marshall and Hicks elasticities, are relevant, as these other elasticity concepts can be deducted from the Frisch elasticity and the Frisch can be shown to be an upper bound for these other elasticities.

[2] In 2015, 70% of all men aged 35-55 in the U.S. were married or lived together with a partner as an unmarried couple (Census Bureau).

Panel Study of Income Dynamics (PSID), we estimate relatively large values for the Frisch elasticity in comparison to previous studies.

The point of departure for our analysis is the conventional approach for estimating the Frisch elasticity using microeconomic panel data going back to Altonji (1986). He has shown that, in a world without borrowing constraints, the Frisch elasticity can be identified from the covariance of hours changes and expected wage-rate changes. In this approach, using expected wage-rate changes as regressor is key as expected wage-rate changes have the property of leaving the marginal utility of wealth unchanged (which is the Frisch concept). Thus, the Frisch elasticity can be recovered from a simple regression of hours growth on expected wage growth when there are no borrowing constraints. In the following, we will refer to such regressions as "Altonji (1986) regressions".

To understand the bias in Altonji (1986) regressions that occurs when borrowing constraints are occasionally binding, i.e., when capital markets are incomplete (see, e.g., Deaton 1991, Aiyagari 1994), consider a situation where an individual's current wage rate is lower than the future wage rate —either due to a negative transitory wage shock or a predictable life-cycle pattern. Without restrictions on borrowing, the individual would work less today and smooth consumption through borrowing, so that the hours change between now and the future is only determined by the Frisch elasticity which is thereby identified. However, if borrowing is not possible, the household's marginal valuation of borrowing and with it the marginal utility of wealth is affected—which violates the Frisch concept. In a borrowing-constrained household, a negative wage-rate shock then tends to increase (rather than decrease) labor supply today since households cannot smooth consumption through borrowing. As a consequence, the hours change is not only determined by the Frisch elasticity in these households. Put differently, the intertemporal-substitution effect of expected wage changes is confounded by a willingness-to-borrow effect which impedes identification in an Altonji (1986) regression. Domeij and Flodén (2006) have shown that, in a pooled sample of constrained and unconstrained households, the negative relation between changes in wage rates and changes in labor supply in borrowing-constrained households biases the estimate of the Frisch elasticity downward. A main result of their analysis is that, without conditioning correctly on household asset holdings—for instance, by eliminating wealth-poor households from the estimation sample— the Frisch elasticity cannot be estimated correctly in an Altonji (1986) set-up. Yet, from a practical point of view, reliable household panel data on assets are hardly available, and even if they are, such data are often not observed in the same panel as labor earnings and working time.

We contribute by extending the analysis of Domeij and Flodén (2006) to a two-person household set-up and by deriving an unbiased estimator of the Frisch elasticity. Our approach critically exploits the couple structure of the model and the data but does not require information on household assets. Intuitively, in a double-earner household, also the partner can react to one's own wage-rate shocks, i.e., also the partner's labor supply can be used to smooth consumption. This is particularly important if the partner earns relatively much. Then, a given negative wage-rate shock can be smoothed relatively easily as the partner's hours have to be raised by only relatively little. Importantly, this relation holds even when the household is borrowing constrained. And, when it is predominantly the partner's labor supply that smooths consumption, one's own hours change is again mostly (in the limit, only) determined by the Frisch elasticity. Accordingly, to derive an unbiased estimator of the Frisch elasticity in presence of borrowing constraints, we can exploit the relation that the household's desire to borrow against wage growth is the less important the less an individual contributes to total household earnings.[3]

In an analytical part of the paper, we make this relation explicit and show that, in borrowing-constrained households, the slope of the decision rule for hours growth is linear in a spouse's usual percentage contribution to household earnings. This relation allows us to derive a regression framework where an interaction term between expected wage growth and this earnings contribution takes up the willingness-to-borrow effect and the non-interacted coefficient on expected wage growth is an unbiased estimate of the Frisch elasticity. Intuitively, expected wage growth multiplied with the relative earnings contribution measures the expected earnings growth (in percent) associated with the expected wage growth if labor supply was unchanged. And it is earnings growth that borrowing-constrained households would want to borrow against and thus causes the estimation bias in the first place.

We then evaluate our estimator in Monte Carlo experiments using a calibrated incomplete-markets model populated by double-earner households, and finally we use the method for estimations using PSID data. Importantly, our approach critically exploits the couple structure of our model and the data, the key issue being that, only in a population of double-earner households, there is variation in individuals' percentage contribution to household earnings

[3]Blundell, Pistaferri, and Saporta-Eksten (2016) provide direct empirical evidence that household consumption reacts more strongly to husbands' wage shocks than to wives' wage shocks which is line with our model since husbands on average contribute larger shares to household earnings. In Guner, Kaygusz, and Ventura (2012a, 2012b), Domeij and Klein (2013), and Bick (2016), similar mechanisms to ours affect labor-supply reactions to permanent wage-rate changes. In these studies, income effects are weaker for women (who are often secondary earners) such that their reactions to permanent wage-rate changes mostly reflect substitution effects governed by the Frisch elasticity.

which we use to identify the Frisch elasticity.[4] Our estimations using PSID data suggest Frisch elasticities for men of about 0.7. We also take into account modifications of our interaction-term approach to cope with challenges that arise when estimating Frisch elasticities for women. For women, we find Frisch elasticities of around one.

A direct implication of our analysis is that conventional methods tend to overestimate differences in labor-supply elasticities between population groups that tend to have different earner roles in the household. One example is the often-discussed difference in labor-supply elasticities between men and women, with women usually being attributed a substantially larger value for the Frisch elasticity than men. Another example is the difference in labor-supply elasticities between individuals with high and low levels of earnings. Our analysis suggests that potential differences in the true elasticities are magnified by the differential importance of the estimation bias so that differences in true elasticities are in fact smaller than suggested by previous studies. This way, our analysis has implications for, e.g., the taxation of couples (Kleven, Kreiner, and Saez 2009), genders (Alesina, Ichino, and Karabarbounis 2011), and top-income earners (Saez 2001). Further, our analysis shows that the negative estimation bias is of particular importance in samples where individuals contribute large shares to total household income—a sample of prime-age male household heads being a prominent example. When we correct for the downward bias, we estimate a Frisch elasticity for men of about 0.7 which is larger than the majority of previous microeconometric estimates, see, e.g., Keane and Rogerson (2015).

The remainder of this paper is organized as follows. In Section 2, we develop an incomplete markets model with two earners. In Section 3, we derive an unbiased estimator of the Frisch elasticity in presence of borrowing constraints exploiting the couple structure of our model. In Section 4, we perform Monte Carlo experiments where we test our estimator on synthetic data from a realistically calibrated version of our model. Section 5 provides an empirical application using PSID data. In Section 6, we discuss the implications of our results for estimated differences in labor-supply elasticities between population groups. Section 7 concludes.

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

The model is a partial-equilibrium incomplete-markets model with two household members. Households differ from one another by asset holdings and wage rates. Members of a household are subject to joint budget and borrowing constraints and take decisions cooperatively

[4]By contrast, single earners by definition always contribute 100% to household earnings.

under full commitment, so that the resulting allocations are Pareto optimal. Households are potentially borrowing constrained and use precautionary savings in a non-state contingent asset and labor supply of both household members to insure against bad wage-rate realizations. This behavior is similar as in the model of Ortigueira and Siassi (2013) and extends the model of Domeij and Flodén (2006) to a two-person setup.

2.1 Decision problem

The decision problem can be represented by the decisions of a household planner. The planner maximizes a weighted sum of members' utilities with weights μ and $1-\mu$ for the two household members $i = 1, 2$, respectively. The household problem in recursive formulation is given by

$$V(a, \omega) = \max_{a', c, n_1, n_2} \mu \cdot u_1(c, n_1) + (1-\mu) \cdot u_2(c, n_2) + \beta \, \mathrm{E}\left[V\left(a', \omega'\right) |\omega\right] \tag{1}$$

subject to the household budget constraint

$$c + a' = w_1 n_1 + w_2 n_2 + (1+r) \cdot a, \tag{2}$$

and the borrowing constraint

$$a' \geq 0, \tag{3}$$

where u_i is the instantaneous utility function, c is household consumption, n_i is hours worked by household member i, β is the rate of time preference, E is the expectation operator, a denotes the household's asset holdings, ω is the vector of wage rates of both household members, $\omega = (w_1, w_2)$, and r is the exogenous interest rate.[5] A prime ($'$) denotes next period values.

In our baseline model, we consider the standard additively separable utility function

$$u_i(c, n_i) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \alpha_i \cdot \frac{(n_i)^{1+1/\eta_i}}{1+1/\eta_i}, \tag{4}$$

where σ denotes risk aversion and α_i is the taste for leisure. This utility function has the property that the true value of the Frisch elasticity is given by the curvature parameter η_i. In Appendix E.1, we consider an alternative model specification with non-separable preferences which yields similar results as our baseline model. The preference parameters are indexed by i as, in our quantitative evaluations, we will account for potential differences in these

[5] A partial-equilibrium set-up is sufficient for our purposes because we neither analyze policy nor parameter changes. We assume $\beta(1+r) < 1$.

parameters between household members.

Wage rates are stochastic and exogenous. Our analytical results depend on wage differences within the household and on variation in expected wage growth but not on the particular specification of the wage process. In our calibrated model, we will assume that wage rates follow stationary first-order autoregressive processes with constant terms that differ between members of the same household as well as across different households ("fixed effects"). Intra-household differences in these constant wage components lead to long-run differences in earner roles among spouses. Transitory wage-rate fluctuations may induce borrowing constraints to bind. Since wage rates are mean-reverting, low wage-rate realizations lead to positive expected wage growth which induces workers to wish to substitute working time intertemporally and to work less in the current period. At the same time, households wish to smooth consumption. For households who do not hold sufficient assets, the borrowing constraint is then binding.

The solution to the maximization problem is described by the policy functions

$$x = x\left(a, \omega\right), \tag{5}$$

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

In our baseline model, we assume that consumption is a household public good, i.e., there is no consumption rivalry between spouses. Along with additive separability, the public good assumption allows a simple notion of Frisch elasticities in a context with two earners. Specifically, the issue of whose marginal utility of wealth is held constant (husband's, wife's, or household's) does not arise, since, if one of them is constant, the other two are constant as well, independent of bargaining. For completeness, we also considered a model version with private instead of public consumption. In this version, we obtain almost identical results, see Appendix E.2. Even allowing for endogenous time-varying Pareto weights in the spirit of a limited-commitment model (see, e.g., Ligon, Thomas, and Worrall 2002) would have no substantial impact on our results since the weights would mostly react to unexpected changes in wage rates while the Frisch elasticity is identified through changes in expected wage rates. In our baseline model, we further abstract from non-linear taxation. This assumption allows to recover the true Frisch elasticity consistently in absence of borrowing constraints. We also consider a model version with progressive joint taxation of spouses in Appendix E.3. Also in this version, we obtain similar results as in our baseline economy.

In a further model extension, we follow the literature (Guner, Kaygusuz, and Ventura 2012a, 2012b, Bick 2016) and take into account the possibility that labor supply of women is also affected by fluctuations in the disutility of work originating from taste-for-work shocks, e.g., capturing shocks to home production or child care, see Section 6.2 for details. This model extension delivers important insights for our empirical investigation of labor-supply elasticities for women.

2.2 Equilibrium conditions

The first-order conditions of the household problem are

$$\mu \cdot \frac{\partial u\left(c, n_1\right)}{\partial c} + (1-\mu) \cdot \frac{\partial u\left(c, n_2\right)}{\partial c} = \frac{\partial V\left(a, \omega\right)}{\partial a} = \lambda\,, \tag{6}$$

$$\phi = \lambda - (1+r)\,\beta\,\mathrm{E}\left[\lambda' | \omega\right]\,, \tag{7}$$

$$\lambda \cdot w_1 = \mu \cdot \alpha_1 \cdot n_1^{1/\eta_1}\,, \tag{8}$$

$$\lambda \cdot w_2 = (1-\mu) \cdot \alpha_2 \cdot n_2^{1/\eta_2}\,, \tag{9}$$

$$\phi \geq 0\,, \tag{10}$$

$$a' \geq 0\,, \tag{11}$$

$$\phi \cdot a' = 0, \tag{12}$$

together with the budget constraint (2), given exogenous wage rates w_1 and w_2 and the initial asset stock a_0. ϕ is the Kuhn-Tucker multiplier on the borrowing constraint (3) and λ is the Lagrange multiplier on the budget constraint (2). Condition (6) reflects that the household equalizes marginal utility of consumption and marginal utility of wealth. Condition (7) is the household's consumption Euler equation which takes its standard form if the borrowing constraint does not bind, $\phi = 0$, and otherwise determines the household's willingness to borrow. Conditions (8) and (9) 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. However, the weights do not impact on *changes* in labor supply, which is the dependent variable in Altonji (1986) regressions (the same holds for α_1 and α_2). Conditions (10)-(12) are the Kuhn-Tucker conditions associated with the borrowing constraint (3). From conditions (8) and (9), it can be seen that the Frisch labor-supply elasticities are equal to the parameters η_1 and η_2, independent of whether the household is borrowing constrained or not. With more general preferences, the true Frisch elasticities would depend on the form of the labor-disutility function but not on the bindingness of the borrowing constraint.

3 Exploiting the couple structure to derive an unbiased estimator of the Frisch elasticity

We now derive a procedure for obtaining an unbiased estimate of the Frisch elasticity in presence of borrowing constraints. We first derive our approach analytically and then evaluate it numerically using Monte Carlo experiments. To derive the estimator analytically, we apply a simplifying assumption on data frequency, which will be relaxed in the Monte Carlo experiments where we will consider a realistic, i.e., annual, data frequency. Specifically, to obtain closed-form solutions, we assume an arbitrarily small period length. Due to this assumption, it is sufficient to consider the group of borrowing-constrained households and the group of unconstrained households and, in first differences, one can neglect households that move from one group to the other.[6] For both groups, we can derive the relation between hours changes and expected wage growth analytically, and then we can pool both groups to derive the population estimate. In the Online Appendix, we derive analytical results which are independent of the period length. While the derivations are more cumbersome, the main results presented here under the assumption of an arbitrarily small period length carry over to the more general case. For simplicity, we assume in the analytical part that spouses' Frisch elasticities are identical. In the quantitative model analysis in Section 4, we account for potential heterogeneity in the true Frisch elasticity to capture gender differences. Further, we assume in the analytical part that the process for stochastic wage-rate components is homogenous across the population. In the quantitative model analysis in Section 4, we take into account gender differences in these processes.

3.1 Households unaffected by borrowing constraints

For households unaffected by borrowing constraints, a regression of hours growth on expected wage growth yields an unbiased estimate of the Frisch elasticity. For bachelor households, this has been shown in the seminal paper by Altonji (1986). Our case of a double-earner household is a straightforward extension. In Appendix A.1, we show that, after taking logs and first differences, the Frisch elasticity can be recovered through regressions of the form

$$\Delta \ln n_i' = \eta \cdot \Delta \operatorname{E} \ln w_i' - \eta \cdot \ln (1+r) - \eta \cdot \ln \beta - \eta \cdot \left(\xi' - \omega_i'\right), \tag{13}$$

for household members $i = 1, 2$, where $\xi' = \ln \lambda' - \operatorname{E} \ln \lambda'$ is an expectation error which results from using the Euler equation to substitute marginal utility of consumption from the

[6] As Altonji (1986) and Domeij and Flodén (2006), we estimate labor-supply regressions in first differences, i.e., these regressions use data from periods $t+1$ and t. Our simplifying assumption ensures that the number of households that are borrowing constrained in one but not both periods is infinitely small.

labor-supply conditions. The terms ω_i', $i = 1, 2$, are unexpected components of wage growth which result from a decomposition of observed wage growth in an expected and unexpected component. As shown by Altonji (1986), the combined residual $\eta \cdot (\xi' - \omega_i')$ is uncorrelated with the regressor expected wage growth, see Appendix A.1 for an intuitive explanation. The terms $\eta \cdot \ln(1 + r)$ and $\eta \cdot \ln \beta$, $i = 1, 2$, can be captured by time fixed effects and a constant, respectively. Thus, when borrowing constraints are not binding, a simple regression of hours growth on expected wage growth ("Altonji (1986) regression") identifies the Frisch elasticity.

3.2 Borrowing-constrained households

When borrowing constraints are binding, a standard Altonji (1986) regression does not yield an unbiased estimate of the Frisch elasticity. This has been shown by Domeij and Flodén (2006) who consider bachelor households and directly translates to our double-earner set-up. Other than Domeij and Flodén (2006), we obtain closed-form expressions for the estimates and biases due to our simplifying assumption of an arbitrarily small period length.

For borrowing-constrained households, for which $a = a' = 0$, we can log-linearize and summarize the first-order conditions (2), (6), (8), and (9),

$$\ln(n_1/\overline{n}_1) = \eta \cdot \ln(w_1/\overline{w}_1) + \eta \cdot \ln(\lambda/\overline{\lambda}), \tag{14}$$

$$\ln(n_2/\overline{n}_2) = \eta \cdot \ln(w_2/\overline{w}_2) + \eta \cdot \ln(\lambda/\overline{\lambda}), \tag{15}$$

$$\ln(\lambda/\overline{\lambda}) = -\sigma \cdot (\overline{s}_1 \cdot (\ln(w_1/\overline{w}_1) + \ln(n_1/\overline{n}_1)) + \overline{s}_2 \cdot (\ln(w_2/\overline{w}_2) + \ln(n_2/\overline{n}_2))), \tag{16}$$

where variables with a bar refer to the point of approximation and $\overline{s}_i$ is individual i's percentage contribution to household earnings at this point, i.e.,

$$\overline{s}_i = \overline{w}_i\overline{n}_i / (\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2),$$

see Appendix A.2 for a derivation.[7] We measure the earnings contribution in the point of approximation $\overline{s}_i$ by the individual's average contribution to household earnings during the sample period. Put differently, the point of approximation is the situation where both spouses contribute their usual shares to household income.[8]

In Appendix A.2, we solve this system to obtain

$$\Delta \ln n_i' = \left(\eta - \frac{\sigma\eta(\eta + 1)}{\sigma\eta + 1} \cdot \overline{s}_i\right) \cdot \mathrm{E}\,\Delta \ln w_i' + \kappa', \tag{17}$$

[7]Equation (16) is a first-order approximation of the budget constraint (in logs). In Appendix C, we evaluate the importance of the approximation for our results and find that the approximation has a negligible effect.

[8]Other variables referring to the point of approximation will drop out in the following due to taking first differences.

where the combined residual $\kappa' = \left(\eta - \frac{\sigma\eta(\eta+1)}{\sigma\eta+1} \cdot \overline{s}_i\right) \cdot (\ln w_i' - \mathrm{E} \ln w_i') - \frac{\sigma\eta(\eta+1)}{\sigma\eta+1} \cdot (1 - \overline{s}_i) \cdot \Delta \ln w_{-i}'$. The first term in this residual stems from a decomposition of observed wage growth in an expected and unexpected component. The second term reflects the cross-reaction to the partner's wage-rate changes.

Equation (17) separates the two effects of expected wage growth on hours growth in a borrowing-constrained household. First, as in an unconstrained household, expected wage growth induces the wish to substitute labor into periods where it is paid more. This intertemporal-substitution effect is governed by the Frisch elasticity, η. Second, expected wage growth induces the willingness to borrow against expected future earnings in order to smooth consumption. However, a borrowing-constrained household can only smooth consumption by supplying more labor which counteracts the intertemporal-substitution effect. The strength of this willingness-to-borrow effect depends on the individual's usual contribution to household earnings $\overline{s}_i$. Expected wage growth for individuals with low earnings contributions induces only relatively small expected changes in total household earnings and can more easily be smoothed through labor-supply adjustments of the partner. Hence, for these individuals, the willingness-to-borrow effect is weak and hours growth is mostly determined by intertemporal substitution and, thus, the Frisch elasticity.[9]

Assuming that $\mathrm{E}\,\Delta \ln w_i'$ has a homogenous variance across the population, the estimated coefficient in a regression of hours growth on expected wage growth in a sample of individuals from borrowing-constrained households with usual earnings contribution $\overline{s}$ equals

$$\eta - \frac{\sigma\eta\,(\eta+1)}{\sigma\eta+1} \cdot \overline{s},$$

which shows that the estimate does not generally recover the true Frisch elasticity.

3.3 Mixed population

We now consider a sample that includes individuals from both borrowing-constrained and unconstrained households. We denote the sample shares of the constrained and unconstrained households by p and $1 - p$, respectively. As an intermediate step, we consider a group of

[9] A similar mechanism applies to permanent changes in wage rates for which borrowing constraints are less important but classical income effects play an important role. For individuals who contribute little to household earnings, changes in hourly wage rates induce a small change in household earnings which mutes the income effect of changes in wage rates. As a consequence, the labor-supply response to these changes is mostly driven by substitution effects and, hence, tends to be stronger than for individuals who contribute larger shares to household earnings or who are the sole earners in their households. This mechanism can help to understand findings reported by Guner, Kaygusz, and Ventura (2012a, 2012b), Domeij and Klein (2013), and Bick (2016) who all document that, in quantitative macro models with double-earner households, labor supply is particularly responsive for groups that can be expected to contribute small shares to household earnings.

individuals with usual earnings contribution $\overline{s}$ but which include both, unconstrained and constrained households. In such a sample, a standard Altonji (1986) regression of hours growth on expected wage growth yields the following estimate for the Frisch elasticity:

$$\begin{aligned}\left.\frac{\operatorname{cov}\left(\Delta\ln n', \mathrm{E}\,\Delta\ln w'\right)}{\operatorname{var}\left(\mathrm{E}\,\Delta\ln w'\right)}\right|_{\overline{s}} &= \frac{\mathrm{E}\left(\Delta\ln n'\cdot\mathrm{E}\,\Delta\ln w'\right)}{\operatorname{var}\left(\mathrm{E}\,\Delta\ln w'\right)} - \frac{\mathrm{E}\left(\Delta\ln n'\right)\cdot\mathrm{E}\left(\mathrm{E}\,\Delta\ln w'\right)}{\operatorname{var}\left(\mathrm{E}\,\Delta\ln w'\right)}\\ &= p\cdot\left(\eta - \frac{\sigma\eta\left(\eta+1\right)}{\sigma\eta+1}\cdot\overline{s}\right) + (1-p)\cdot\eta \\ &= \eta - p\cdot\frac{\sigma\eta\left(\eta+1\right)}{\sigma\eta+1}\cdot\overline{s},\end{aligned} \tag{18}$$

which uses that $\mathrm{E}\left(\Delta\ln n'\right)\cdot\mathrm{E}\left(\mathrm{E}\,\Delta\ln w'\right)=0$.[10] The final step is to consider a sample where individuals differ in their usual contributions to household earnings. In such a sample, the OLS estimate averages over the different $\overline{s}$ such that the coefficient on expected wage growth is

$$\widehat{\eta} = \frac{\operatorname{cov}\left(\Delta\ln n', \mathrm{E}\,\Delta\ln w'\right)}{\operatorname{var}\left(\mathrm{E}\,\Delta\ln w'\right)} = \eta - p\cdot\frac{\sigma\eta\left(\eta+1\right)}{\sigma\eta+1}\cdot\overline{\overline{s}}, \tag{19}$$

where $\overline{\overline{s}}$ is the sample average of the usual earnings contribution of individuals from borrowing-constrained households.

The bias term in equation (19), $-p\cdot\frac{\sigma\eta(\eta+1)}{\sigma\eta+1}\cdot\overline{\overline{s}}$, has three important properties. First, as pointed out by Domeij and Flodén (2006), borrowing constraints lead to a downward biased estimate $\widehat{\eta}$ as the term that is subtracted from the true Frisch elasticity is unambiguously positive. Second, we also see Domeij and Flodén (2006)'s result that an unbiased estimate can in principle be obtained in a sample of individuals from households which are unaffected by borrowing constraints as, in such sample, $p=0$. The third property is of utmost importance from a practical point of view. Standard Altonji (1986) regressions yield less strongly biased estimates of the Frisch elasticity in samples of individuals that usually contribute only little to household earnings as, in such samples, $\overline{\overline{s}}$ is small (for example, in a sample of secondary earners). In empirical applications using PSID data, we will provide evidence supporting this relation.

3.4 Deriving an unbiased estimator

The relation between the earnings contribution and the covariance term $\operatorname{cov}\left(\Delta\ln n', \mathrm{E}\,\Delta\ln w'\right)$ in (19) holds two key insights for deriving a regression specification that yields an unbiased estimate of the Frisch elasticity, even in samples that include

[10] Since the wage process is assumed to be identical for both groups, also the variance of the regressor expected wage growth is identical for both groups. Consequently, the OLS estimator weighs both groups according to their respective sample shares.

borrowing-constrained households. First, in a population of double-earner households, there is variation in individuals' contribution to household earnings which can be used to identify the Frisch elasticity. Second, the covariance $\operatorname{cov}(\Delta \ln n', \mathrm{E}\, \Delta \ln w')$ is a *linear* function of $\overline{s}$, see (18).

This implies that, in a sample that consists of individuals with different usual contributions to household earnings, an interaction-term regression of the type (introducing household and time indices to clarify the panel dimension of the estimation)

$$\Delta \ln n_{ijt+1} = const. + \delta_1 \cdot \mathrm{E}_t\, \Delta \ln w_{ijt+1} + \delta_2 \cdot \mathrm{E}_t\, \Delta \ln w_{ijt+1} \cdot \overline{s}_{ij} + u_{ijt+1}, \tag{20}$$

gives

$$\widehat{\delta}_2 = -\frac{\sigma\eta\,(\eta+1)}{\sigma\eta+1} \cdot p$$

and

$$\widehat{\delta}_1 = \eta,$$

where $\overline{s}_{ij}$ is the average percentage contribution of individual i to labor earnings of household j and the index $ijt+1$ refers to member i of household j in period $t+1$. Thus, in a regression that controls for the interaction between expected wage growth and the individual's average contribution to household earnings, the coefficient on expected wage growth is an unbiased estimate of the Frisch elasticity. Note that, in our approach, the estimated coefficient on the interaction term is not of interest per se but the interaction term needs to be included as a control variable to correctly identify the Frisch elasticity as the coefficient on expected wage growth.

Intuitively, our interaction-term regression controls for the product of expected wage growth and the individual's average percentage earnings contribution. This product measures the expected earnings growth (in percent) which is implied by the expected wage growth. For a borrowing-constrained household, income growth is tightly connected to earnings growth. And it is expected income growth a household would like to borrow against. Thus, we control for the expected income growth caused by the individual's expected wage growth and hence we control for the change in the willingness to borrow. This takes out the willingness-to-borrow effect from the coefficient on the non-interacted regressor $\mathrm{E}\, \Delta \ln w_{ijt+1}$ and what remains is the pure intertemporal-substitution effect governed by the Frisch elasticity.

Note that the double-earner framework is incremental for this method to recover the Frisch elasticity. For singles or single earners, $\overline{s} = 1$ so that the two regressors in the regression above are the same and it is impossible to identify δ_1 and δ_2 separately.

Figure 1: Policy functions for labor supply.

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

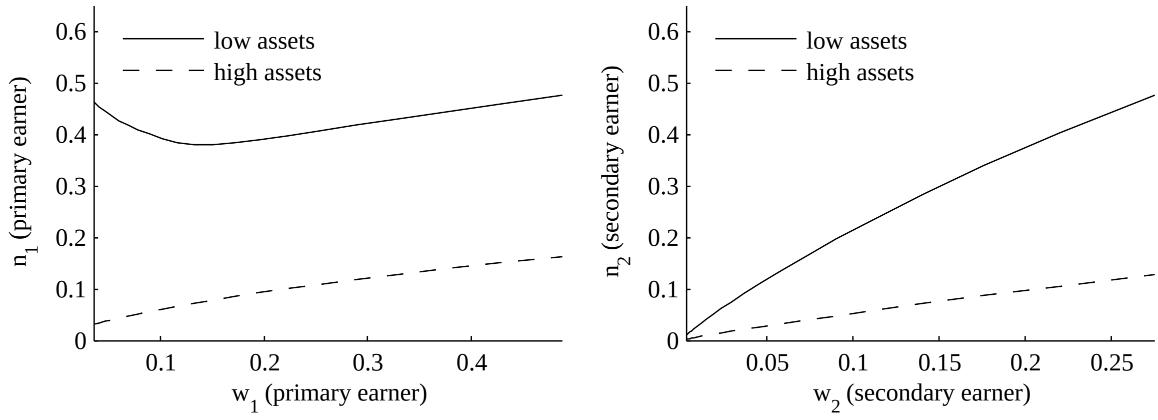

Notes: Policy functions refer to household type X, which is a household type with pronounced long-run intra-household wage differences. In the left panel, the wage rate of the secondary earner is at its lowest possible grid value. In the right panel, the wage rate of the primary earner is at its highest possible grid value. Solid lines refer to zero asset holdings. Dashed lines refer to an unconstrained household.

3.5 Graphical illustration

Figure 1 shows policy functions from a numerical solution of our calibrated full model.[11] For the graphs, we compare the household member who contributes, in the long run, more to household earnings (the primary earner) and the member who contributes less (the secondary earner). For illustration, we consider a household with strong wage-rate differences between household members such that the willingness-to-borrow effect is strong for the primary earner and weak for the secondary earner.

The labor-supply curve of the primary earner (left panel) is globally upward-sloping if the household is wealth-rich (dashed line), reflecting the intertemporal-substitution effect governed by the Frisch elasticity, due to the household's ability to smooth consumption through desaving when wage rates are low. The standard Altonji (1986) regression identifies the Frisch elasticity from this upward-sloping shape of the labor-supply curve. By contrast, for a household with low asset holdings (solid lines), the borrowing constraint is binding when wage rates are low. Then, the labor-supply curve of the primary earner has a downward-sloping range where a further wage decrease triggers an increase in labor supply (rather than a decrease), because consumption cannot be smoothed through borrowing but only through an increase in labor supply. This negative relation between wage-rate changes and labor-supply

[11] The model calibration is discussed in Section 4.1.

changes leads to a downward estimation bias.

In contrast to the primary earner, the labor-supply curves of the secondary earner (right panel) are globally upward-sloping, independent of whether the household is wealth-rich or borrowing constrained. Also at the borrowing constraint, low wage rates of the secondary earner can be compensated relatively easily by a relatively small increase in the primary earner's hours. Thus, for secondary earners, the labor-supply reaction to transitory wage-rate changes is mostly governed by the Frisch elasticity, so that, everything else equal, 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 larger the intra-couple wage gap, the stronger is this effect. Our interaction-term approach given by (20) generalizes this to the case where we exploit variation in individuals' usual contribution to household earnings and its continuous effect on the slope of the labor-supply curve.

4 Estimating labor-supply elasticities from synthetic data

In this section, we use our model which is calibrated to a period length of one year to quantify how successful our interaction-term approach is to recover the true Frisch elasticity in data sets that have realistic properties and where households are occasionally borrowing constrained, i.e., where households move from being borrowing constrained to being unconstrained between periods. We solve the full model globally using numerical techniques.

4.1 Calibration

Our baseline PSID sample used for the calibration covers the period 1972-1997, see Appendix B.1 for details on the sample selection. Due to our focus on double-earner households, we consider household heads and their partners for whom both partners' wage rates are observed. Further, we apply similar sample selection criteria as Altonji (1986) and Domeij and Flodén (2006). In particular, we consider individuals between age 25 and 60.

In the numerical evaluations, we assume that the wage process consists of a stochastic component z_i which follows an AR(1) process with autocorrelation ρ_i and innovations ε_i, and we account for constant terms ψ_i leading to long-run wage differences between individuals within and across households (fixed effects),

$$\begin{aligned} \ln w_i &= \psi_i + z_i, \\ z_i' &= \rho_i \cdot z_i + \varepsilon_i'. \end{aligned} \tag{21}$$

We estimate the parameters of the stochastic wage processes, i.e., autocorrelations ρ_m, ρ_f and innovation variances $\sigma^2_{m,\varepsilon}$, $\sigma^2_{f,\varepsilon}$, separately for men (m) and women (f).[12] We first obtain residual wages by filtering deterministic cross-sectional variation using an OLS regression. We then identify autocorrelations and innovation variances from gender-specific Generalized Method of Moments (GMM) estimations, see Appendix B.2 for details.[13]

Our interaction-term approach that corrects for the bias due to borrowing constraints exploits variation in individuals' usual percentage contribution to household earnings. In order to assess our method in the model, the simulated economy has to feature sufficient and realistic variation in individuals' contribution to household earnings. We therefore solve and simulate our model with ten household types. Household types differ in the constant (=permanent) wage components ψ_i of its members which we set to match average male and female wage rates in the ten deciles of the empirical distribution of relative wage rates of spouses in couple households in our PSID sample. We then calibrate household-type specific preference weights α_m and α_f to match average hours worked by gender and group, and, as a result, our calibrated model displays a realistic distribution of relative labor earnings within households.[14]

We calibrate the gender-specific values for the Frisch elasticities so that the *estimated* Frisch elasticities in our Monte Carlo study coincide with the *estimated* Frisch elasticities for men and women that we estimate from the PSID data (see Section 5), both using a standard Altonji (1986) regression. We will discuss in Section 6 that one needs only relatively small differences in the true gender-specific Frisch elasticities ($\eta_m = 0.65$ and $\eta_f = 0.90$) to rationalize the relatively strong difference in empirically estimated Frisch elasticities (roughly factor 2), as the difference in the true elasticities is magnified by the differential importance of the estimation bias for men and women.

For the remaining preference parameters we use standard values from the literature. Relative risk aversion is set to $\sigma = 1.5$.[15] Following Domeij and Flodén (2006), we set $\beta = 0.95$ (annual model frequency), and calibrate the interest rate so that the bottom 40%

[12] Blundell, Pistaferri, and Saporta-Eksten (2016) document that alternatively using a combination of permanent and transitory shocks leads to similar estimation results for preference parameters such as Frisch elasticities.

[13] For the numerical solution of the model, the joint wage process is discretized using Tauchen's (1986) algorithm with 21 grid points per household member, i.e., 441 husband-wife wage combinations. We solve the model using the endogenous grid point method of Kabukcuoglu and Martinez-Garcia (2016) who extend Carroll (2006)'s method to an infinite horizon model with an arbitrary number of control variables.

[14] An alternative approach would be to target the estimated variance of fixed effects from the microeconometric wage process estimation. While this would capture the gender-specific *across*-household variance of (residual) wage rates appropriately, we implement the former approach to obtain a realistic distribution of *within*-household wage differences.

[15] We also considered a model specification with differences in risk aversion between household members. The results are very similar to the ones obtained from our baseline model.

of the wealth distribution own 1.4% of total wealth. Table 7 in Appendix B.3 summarizes all parameter values of our baseline model.

4.2 Simulation set-up

We simulate a synthetic panel data set with similar size as our baseline PSID sample. Specifically, we simulate households for a long period of time and calculate hours growth, expected wage growth, average wage rates, and average contributions to household earnings. We then draw 10,000 samples of 15,000 household-year observations which we use for the regressions and report mean point estimates and mean standard deviations. In the estimations, we consider separate samples of men and women to take into account gender differences in both, true Frisch elasticities and usual earner roles. In the main text, we report the estimation results for men while results for women are similar and can be found in Appendix D. To determine the regressor expected wage growth, we exploit the properties of the wage process (21), i.e., we calculate $\mathrm{E}_t \, \Delta \ln w_{ijt+1} = (\rho_i - 1) \cdot z_{ijt}$.

4.3 Monte Carlo results

Table 1 summarizes the estimation results from various regression specifications using the simulated model data. To begin with, column (1) shows the results from a standard Altonji (1986) regression, i.e., a regression of hours growth on expected wage growth without the interaction term that we proposed in Section 3. This auxiliary regression reflects our calibration target, as we calibrated the true Frisch elasticity for men (0.65) so that the standard Altonji (1986) regression yields an estimated value of 0.41 (which we obtain in our estimations from PSID data) and illustrates the negative estimation bias.

Before applying our preferred interaction-term approach to simulated model data, we illustrate two main implications of our model for standard Altonji (1986) regressions. Later, we will test both implications empirically using PSID data. The first implication is that estimates of the Frisch elasticity obtained by Altonji (1986) regressions should, ceteris paribus, be smaller in samples of individuals that contribute larger shares to household earnings, i.e., in samples of individuals with large $\overline{\overline{s}}$ in equation (19). The second implication is that differences in estimated Frisch elasticities between groups with different contributions to household earnings should become smaller when the samples are less affected by borrowing constraints, i.e., in samples of wealthier households where p in equation (19) tends to be small.

To illustrate both implications, we estimate otherwise standard labor-supply regressions

Table 1: Estimation results for men, from synthetic household panel data.

	(1)	(2)	(3)	(4)	(5)
expected wage growth	0.41 (0.01)	0.50 (0.05)	0.67 (0.07)	0.63 (0.10)	0.62 (0.12)
expected wage growth × primary earner		-0.10 (0.05)	-0.01 (0.07)		
expected wage growth × earnings contribution (%)				-0.32 (0.15)	-0.33 (0.17)
bias	-38%	—	—	-3%	-5%
sample	all	all	$a > \overline{\overline{a}}$	all	$\overline{w}_m > \overline{w}_f$
observations	15,000	15,000	4,600	15,000	13,518

Notes: Estimation results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Primary-earner dummy d_{ij} is one when individual i is the primary earner in household j and zero otherwise. Individuals identified as primary earners if the mean realized wage rate in the simulation $\overline{w}_{ij}$ exceeds the mean realized wage rate of the spouse $\overline{w}_{-ij}$. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the simulation. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses. In columns (3) and (5), we first draw a sample of 15,000 observations in each Monte-Carlo repetition and then only keep the observations which satisfy the respective sample selection criterion (see second to last row). Reported sample sizes in columns (3) and (5) are average sample sizes.

but include an interaction between expected wage growth and a dummy variable that indicates whether the individual is the primary earner in the household, i.e., has a higher average wage rate than the spouse in our simulation. When we estimate this specification from the simulated data, we obtain a negative estimate for this interaction term, see column (2). This indicates that standard Altonji (1986) regressions would assign a smaller estimate of the Frisch elasticity to primary earners although the true Frisch elasticity η is the same across the male population. The reason is that, almost by definition, primary earners have a high contribution to household earnings.

Column (3) relates our analysis to Domeij and Flodén (2006) and shows results for samples where we condition on household assets. Specifically, we restrict the sample to households whose asset holdings exceed the average asset holdings in the simulated economy. As expected, the estimated coefficient on the primary-earner interaction becomes substantially smaller in absolute value than the one in column (2), reflecting that earner roles tend to become irrelevant when borrowing constraints are not relevant in the estimation sample. In line with Domeij and Flodén (2006), we find that the estimated coefficient on expected wage growth is rather close to the true Frisch elasticity when the sample is restricted to above-average wealth. However, in an estimation based on real-world data, a sufficiently strong restriction on assets can be practically problematic, for example due to data availability or small sample sizes due to missing information on wealth components.

Column (4) shows the estimation results for our preferred interaction-term model summarized in equation (20), estimated from the unrestricted sample. In this model, we extend the standard Altonji (1986) regression by an interaction term between expected wage growth and an individual's average earnings contribution as a control variable. We find that our interaction-term regression works well in samples with annual data frequency. The estimated Frisch elasticity is 0.63 which is very close to the true value of 0.65. Hence, our approach that exploits the couple structure of the data yields almost unbiased estimates of the Frisch elasticity in data sets that have realistic properties in terms of sample size and data frequency. Note that we estimate our interaction-term approach on the unrestricted sample of individuals, i.e., without using any information on household wealth. Thus, our Monte Carlo experiments show that our interaction-term approach yields almost unbiased estimates even in samples of potentially borrowing-constrained individuals.

While we have shown that the bias due to borrowing constraints in Altonji (1986) regressions is smaller for individuals who contribute little to household income, in real-world data, men often tend to be primary earners in the household. Accordingly, one might be concerned that, in an application using empirical data, a group of male secondary earners has specific characteristics which might cause additional problems when inferring the Frisch elasticity. We therefore perform an additional Monte Carlo experiment to corroborate that, while men with low earnings contributions are less subject to the bias due to borrowing constraints in Altonji (1986) regressions, they are not necessarily needed for identification in the regression framework we propose. To do so, we estimate equation (20) on a restricted sample that includes only men who are *primary* earners in their respective households. Column (5) shows that also in such a sample, we obtain an estimate very close to the true Frisch elasticity when we account for our interaction term. Put differently, also variation in the upper part of the distribution of earnings contributions can be exploited to successfully recover the Frisch elasticity through our method. To understand this result, recall that we have shown in Section 3 that the covariance between expected wage growth and hours growth for different groups of individuals with earnings contribution $\overline{s}$ is a linear function of $\overline{s}$, see (18). Also in the full model, the relation seems to be close to linear.

We have considered several extensions of our baseline model and have investigated the performance of our interaction-term approach in these extended model environments. Specifically, we have incorporated non-separable preferences, progressive income taxes, and private instead of public consumption, see Appendices E.1-E.3 for details. In all model extensions, we find that our preferred interaction-term approach delivers an estimate of the Frisch elasticity which is close to its true value while standard Altonji (1986) regressions underestimate it considerably. Finally, we have developed modifications of our interaction-term approach to cope with challenges when estimating the Frisch elasticity for women instead of men. In Appendix E.4, we show that our modified approaches are robust in presence of taste-for-work shocks originating from, e.g., child care or home production, in particular our approach using predicted instead of actually observed earnings contributions.

5 Estimating labor-supply elasticities from PSID data

In this section, we present empirical results for our interaction-term approach using PSID data. As shown above, our approach corrects for the bias due to borrowing constraints and is able to deliver an almost unbiased estimate of the Frisch elasticity. The outline of the empirical analysis closely follows our Monte Carlo experiments, i.e., we first investigate two key implications of our model and then apply our preferred interaction-term approach to the data. Before presenting estimation results, we discuss some econometric aspects that are relevant in an empirical application of our approach.

5.1 Econometric aspects

As in our theoretical model, we analyze the choice of hours worked at the intensive margin in double-earner households.[16] When estimating labor-supply regressions from PSID data, we use individual characteristics to determine expected future wage changes, $\mathrm{E}_t\, \Delta w_{ijt+1}$, in gender-specific OLS regressions. Specifically, we follow, e.g., MaCurdy (1981) and Domeij and Flodén (2006) and use as predictors age, age squared, years of schooling, and an interaction term between age and years of schooling. If there were no borrowing constraints, predictable wage growth would leave the marginal utility of wealth unchanged and would hence identify the Frisch elasticity. Using individual characteristics as predictors has the advantage that measurement error in these variables is uncorrelated with measurement error in wage rates.[17]

[16] In our sample, the standard deviation of annual hours growth, which is the left-hand side variable in our regressions, is 16.9% for men (and 24.8% for women). Thus, the data show that there is substantial variation in hours at the intensive margin.

[17] Measurement error in hours on the left-hand side of the regression reduces the R^2 of the regression but the estimate for η is consistent. As discussed by Altonji (1986), Domeij and Flodén (2006) and Keane (2011), the instruments used to determine expected wage growth are potentially weak, one reason being that most wage changes may simply be unexpected. A potentially strong instrument is the lagged wage rate but this instrument should be avoided because it magnifies biases stemming from measurement error in the wage data

In empirical data, individual labor supply may also be affected by taste shifters.[18] Using first-differenced data is helpful in addressing this aspect. First-differencing eliminates the need to control for permanent taste shifters which are likely correlated with wages, such as education. In turn, this means that only transitory taste shifters may still be present in the first-differenced regression. As argued by, e.g., Keane (2011), transitory taste shifters are less likely to be correlated with expected wage changes.

Empirical estimates of labor-supply elasticities can also be affected by non-linear taxation (e.g., Aaronson and French 2009). When taxes are progressive, changes in gross wage rates overstate changes in net wage rates. Even if marginal net wages were observable, they would be endogenous as changes in hours affect marginal tax rates under progressive income taxation. In Appendix E.3, we present Monte-Carlo estimations for a model version with progressive income taxation which show that, in our context, the biases due to progressive taxation are small compared to the biases arising from borrowing constraints. Relatedly, it may be argued that taxes will largely drop out of a labor-supply condition in yearly differences as the household's marginal tax rate typically does not change substantially from year to year, see, e.g., Altonji (1986).

5.2 Testing two key implications of our model

As in the Monte Carlo experiments, we begin with comparing estimates for primary and secondary earners. For this, we use the same regression specification as in the Monte Carlo experiments, i.e., an otherwise standard Altonji (1986) labor-supply regression which we augment by an interaction between expected wage growth and a dummy variable indicating whether the individual is the primary earner in the household. In our baseline specification, we classify the spouse with the higher average wage rate over the sample period as the primary earner, as we did in the Monte Carlo experiments.

Our theory predicts a negative coefficient for the interaction term. This is confirmed in our estimations using PSID data, see Table 2. For both, men and women, the incremental effect of being a primary earner on the estimated Frisch elasticity is significantly negative. This corroborates that labor-supply elasticities are estimated to be substantially smaller for primary than for secondary earners when borrowing constraints are ignored. Considering

(see Altonji, 1986, for further discussion). Using higher lags of the wage rate would mitigate this problem, but, in our sample, such instruments are barely informative for future wage changes. In our gender-specific first-stage regressions to obtain expected wage growth, the F statistics are 18.59 for men and 11.69 for women. In the IV literature (see Stock, Wright, and Yogo 2002), instruments are regarded as reliable if, in the case of one endogenous regressor, the F statistic exceeds 10.

[18] Technically, hours growth in our model is also affected by changes in preferences and bargaining weights, $\Delta \ln \alpha$ and $\Delta \ln \mu$, which are both equal to zero in our model but need not be in empirical data.

Table 2: Empirical labor-supply regressions, PSID data, distinction by binary earner status.

	(1)	(2)
	men	women
expected wage growth	0.52 (0.12)	0.87 (0.17)
expected wage growth $\times$ primary earner	-0.15 (0.09)	-0.45 (0.10)
$\widehat{\eta}_{prim}/\widehat{\eta}_{sec}$	0.70	0.48
time effects	yes	yes
observations	14,340	14,340

Notes: Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Individuals identified as primary earners if the mean realized wage rate in the sample $\overline{w}_{ij}$ exceeds the mean realized wage rate of the spouse $\overline{w}_{-ij}$. Standard errors in parentheses.

gender-specific regressions is important to make this point as they show that the estimated differences in labor-supply elasticities of primary and secondary earners are indeed related to differences in earner status and do not primarily pick up gender differences in the true Frisch elasticities. In Appendix F, we present additional evaluations corroborating this point. Specifically, we consider alternative definitions of primary and secondary earners and we compare Altonji (1986) regressions in several ranges of the relative contribution to household earnings.

We find a particularly large Altonji (1986) estimate for female secondary earners which is in line with our argumentation as this group of women contributes particularly little to household earnings (29% on average in our sample). Of course, estimated gender differences may also reflect differences in the true Frisch elasticities. We come back to the issue of gender differences in true labor-supply elasticities in Section 6.2.

The second testable implication of our analysis is that differences in estimated Frisch elasticities from Altonji (1986) regressions should become smaller when the samples are less affected by borrowing constraints. As in our Monte Carlo experiments, we test this prediction by comparing male primary and secondary earners in samples of households with different liquid wealth. In particular, we repeat the estimations including an interaction term with the primary-earner dummy but only consider households with liquid wealth above a certain threshold which we increase step by step. For this evaluation, we build on Domeij and Flodén

(2006) and restrict the PSID data to three 3-year panels for which detailed asset data are available. Our theoretical model predicts the coefficient on the interaction term to decrease with an increasing wealth cut-off. We find that this pattern is confirmed in the PSID data, see Appendix F.3 for details.

5.3 Frisch-elasticity estimates for men

We now estimate Frisch elasticities for men, comparing results from a standard Altonji (1986) regression to results from our preferred interaction-term approach. Column (1) in Table 3 shows results for a standard Altonji (1986) regression of hours growth on expected wage growth that does not include an interaction term. This specification is subject to the negative borrowing-constraint bias and delivers an estimated Frisch elasticity is 0.41.

We estimate a substantially larger Frisch elasticity when we use our preferred interaction-term approach that exploits the couple structure. This is in line with our theoretical analysis where we have shown that our approach corrects for the negative bias due to borrowing constraints. Our interaction-term approach using the husband's average percentage earnings contribution yields an estimated Frisch elasticity (the coefficient on non-interacted expected wage growth) of 0.72, see column (2). Compared to the estimation without the interaction term in column (1), the bias-corrected estimate is hence about three quarters higher.[19] Comparing the estimation results in columns (1) and (2) suggests that the bias due to borrowing constraints in Altonji (1986) regressions amounts to more than 40% for men which is quantitatively in line with our Monte-Carlo experiments. The negative coefficient on the interaction term in column (2) reflects that men with higher earnings contributions have a weaker connection between expected wage growth and hours growth. This corroborates that their labor supply is particularly strongly exposed to the effects of borrowing constraints which induce a negative co-movement of expected wage growth and hours growth counteracting the positive co-movement induced by intertemporal substitution and governed by the Frisch elasticity.

As in our Monte Carlo experiments, we also investigate in how far our estimates are driven by male secondary earners in the sample. As discussed before, these individuals may be particular in various aspects and one may be sceptical when identification would largely depend on these individuals. In order to address this concern, we re-estimate our interaction-term specification for a restricted sample, where we only include men who are primary earners, see column (3) of Table 3. We find that the estimate for the primary-earner

[19] A one-sided test supports the hypothesis that the estimate in column (2) is significantly larger than the one in column (1) (alternative hypothesis rejected with p-value of 0.08).

Table 3: Empirical labor-supply regressions for men, PSID data, preferred approach exploiting variation in relative contributions to household earnings.

	(1)	(2)	(3)
expected	0.41	0.72	0.69
wage growth	(0.10)	(0.21)	(0.27)
expected wage growth		-0.52	-0.51
$\times$ earnings contribution (%)		(0.29)	(0.38)
time effects	yes	yes	yes
sample	all	all	$\overline{w}_m > \overline{w}_f$
observations	14,340	14,340	11,632

Notes: Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the sample. Standard errors in parentheses.

only sample is similar to the one obtained for the full sample, in line with our results from the Monte Carlo analysis. This corroborates that male secondary earners are not solely responsible for identification although we use the covariance between expected wage growth and hours growth if the male earnings contribution *were* small.

6 Implications for labor-supply elasticities of different population groups

A direct implication of our analysis is that conventional methods tend to overestimate differences in labor-supply elasticities between population groups that tend to have different earner roles in the household, e.g., between primary and secondary earners. Another example is the often-discussed difference in labor-supply elasticities between men and women, with women usually being attributed a substantially larger value for the Frisch elasticity than men. A third example is the difference in labor-supply elasticities between individuals with high and low earnings, respectively. Our analysis suggests that potential differences in the true elasticities are magnified by the differential importance of the estimation bias so that differences in true elasticities are in fact smaller than suggested by previous studies that ignore borrowing constraints and earner roles.

6.1 Primary and secondary earners

In our baseline Altonji (1986) estimations where we allowed the estimate for the Frisch elasticity to depend on earner status, we find substantial differences between primary and secondary earners, see Table 2. The results of our preferred interaction-term approach using an individual's earnings contribution, see Table 3, corroborate that the differences between primary and secondary earners suggested by standard Altonji (1986) regressions are mostly the result of differential estimation biases rather than of differences in true Frisch elasticities. In fact, when we apply our interaction-term approach that corrects for the borrowing-constraint bias to a sample of male primary earners only, we obtain a similar estimate (0.69, see column (3) of Table 3) compared to the total sample of men that also includes secondary earners (0.72, see column (2) of Table 3). Put differently, differences between primary and secondary earners are small once the borrowing-constraint bias is corrected for.

This suggests that the usual sample restriction to, e.g., male household heads working full-time is potentially problematic in microeconometric estimations of the labor-supply elasticity. Such samples consist mostly of primary earners and are hence subject to strong estimation biases, which may be one reason why previous studies have often obtained relatively small estimates for the Frisch elasticity. Keane (2011) explicitly makes the point that, even among men, labor-supply elasticities are likely larger than estimated by the majority of existing studies. Our study supports this view, as we obtain substantially larger estimates in samples where the bias due to borrowing constraints is expected to be less severe. Our analysis can thus help to reconcile micro and macro estimates of labor-supply elasticities (Keane and Rogerson 2015).

6.2 Men and women

Our study suggests that part of the often-discussed gender difference in labor-supply elasticities can be attributed to the fact that men, who are in most cases primary earners in the household, usually contribute larger shares to household income than women, so that everything else equal, the negative estimation bias in Altonji (1986) regressions is larger for men than for women. To address potential gender differences in true elasticities, we take into account that while our interaction-term approach corrects for the bias due to borrowing constraints, this does not necessarily imply that it yields an unbiased estimate of the Frisch elasticity when there are other important sources of biases. While the literature has discussed savings as the most important non-wage labor-supply determinant for men, issues like child care are of particular relevance for women (see Keane 2011). Moreover, these issues can be particularly important for those women for whom we also observe low contributions to household earnings. This could then confound with the correction for the borrowing-constraint bias.

For example, we could measure low contributions to household earnings for women who work only few hours and have a relatively elastic labor supply because of child-care obligations, as shown by Alesina, Ichino, and Karabarbounis (2011). Then, our derived estimator would put relatively much weight on a group of women whose labor-supply elasticity is not representative for the total population. We therefore develop modifications of our baseline interaction-term approach to address challenges when estimating the Frisch elasticity for women.[20]

Specifically, we first extend our theoretical model by shocks to wives' preferences for labor supply and then modify our interaction-term approach appropriately. In particular, we add a stochastic term h to the disutility of work of women such that women's preferences are described by

$$u\left(c, n_i\right) = \frac{c^{1-\sigma}}{1-\sigma} - \alpha_i \cdot \frac{\left(n_i + h_i\right)^{1+1/\eta_i}}{1+1/\eta_i}. \tag{22}$$

The shock h can be understood as a home production requirement, e.g., the presence of children without the availability of informal or affordable formal child care.[21] Similar to Guner, Kaygusz, and Ventura (2012a, 2012b), we model h as a two-state Markov process with states h_{low} and h_{high} and transition probabilities κ_1 from h_{low} to h_{high} and κ_2 from h_{high} to h_{low}.[22]

In Appendix E.4, we present a detailed analysis of the model extension with preference shocks and we show that our baseline interaction-term approach tends to over-estimate the true Frisch elasticity in this setting. We therefore suggest modifications of our baseline approach and we show in Monte-Carlo experiments that, with these modifications, we obtain almost unbiased estimates also in the model with preference shocks. First, we adopt an approach where we consider a sample restriction and only consider women who contribute at least 30% to household earnings. Second, we apply an approach where we replace the wife's actual earnings contribution by the predicted earnings contribution based on a regression with observable determinants as regressors.[23] This alternative measure of the earnings contribution

[20] In terms of descriptive statistics, we also find that women with low contributions to households earnings have particular characteristics. Women with low earnings contributions (below 30%) work fewer hours and have more children than other women. Further, they earn on average almost 50% less than predicted by their characteristics (prediction regression based on a full set of age, education, and year dummies for women). By contrast, men with low contributions to household earnings are rather similar to other men in terms of hours worked, children, and deviations from predicted earnings.

[21] Alesina, Ichino, and Karabarbounis (2011) use these preferences to rationalize gender differences in labor-supply elasticities as a result of the division of household chores. Also Guner, Kaygusz, and Ventura (2012a, 2012b) and Bick (2016) apply similar preferences when analyzing the responses of female labor supply to tax reforms and child care subsidies, respectively. Most relatedly, Guner, Kaygusz, and Ventura (2012a, 2012b) add a constant term to mothers' (but not fathers') working time while young children are present in the household, which happens exogenously in their model.

[22] We follow Guner, Kaygusuz, and Ventura (2012a) and Bick (2016) to calibrate these additional parameters, see Appendix E.4 for details.

[23] In the empirical application, we predict log earnings using a full set of age, education, and year dummies.

Table 4: Empirical labor-supply regressions for women, PSID data, modified approaches exploiting variation in relative contributions to household earnings.

	(1)	(2)	(3)
expected	0.78	1.08	1.05
wage growth	(0.17)	(0.27)	(0.23)
expected wage growth		-1.49	
$\times$ earnings contribution (%)		(0.46)	
expected wage growth			-0.70
$\times$ predicted contribution			(0.44)
time effects	yes	yes	yes
sample	all	$\overline{s}_{ij} \geq 0.3$	all
observations	14,340	8,966	14,340

Notes: Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution $\overline{s}_{ij}$ is the average percentage contribution of individual i to labor earnings of household j in the sample. Predicted earnings contribution based on a regression with a full set of age, education, and year dummies as regressors. Standard errors in parentheses.

is less affected by idiosyncratic determinants (such as child care needs).

When we apply these two approaches to the PSID data, we obtain an estimated Frisch elasticity for women of around one, see columns (2) and (3) of Table 4. These estimates for women are only about 45% larger than the one for men. By contrast, comparing gender-specific labor-supply elasticities on the basis of Altonji (1986) regressions, i.e., comparing the results in column (1) of Table 4 and column (1) of Table 3, would suggest considerably larger gender differences of about 90%.[24] Also the calibration of our theoretical model suggests rather small gender differences in labor-supply elasticities. In fact, a difference of only about 40% is needed to rationalize the substantially larger difference in Altonji (1986) estimates. In summary, our analysis for women suggests that potential gender differences in the true elasticities are magnified by the differential importance of the estimation bias so that differences in true elasticities are in fact smaller than suggested by previous studies. This way, our analysis has implications for, e.g., the taxation of couples (Kleven, Kreiner, and Saez 2009) or genders (Alesina, Ichino, and Karabarbounis 2011) where arguments often rely on gender differences in labor-supply elasticities.

Table 5: Empirical labor-supply regressions for men, PSID data, by earnings group.

	(1)	(2)	(3)	(4)
	top 25% earnings		bottom 75% earnings	
expected	0.18	0.62	0.52	0.87
wage growth	(0.16)	(0.41)	(0.13)	(0.25)
expected wage growth		-0.68		-0.59
× earnings contribution (%)		(0.56)		(0.36)
time effects	yes	yes	yes	yes
observations	3,735	3,735	10,605	10,605

Notes: Estimation results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the sample. Earnings groups defined using the distribution of individual labor earnings in the year of observation. Standard errors in parentheses.

6.3 High and low earnings

Our analysis also implies that conventional methods overestimate the differences in labor-supply elasticities between groups with high and low *levels* of earnings.[25] When we distinguish between men in the upper 25% of the earnings distribution and those in the bottom 75%, see columns (1) and (3) of Table 5, estimates from standard Altonji (1986) regressions suggest that the labor supply of individuals with high earnings is considerably less elastic. Accordingly, one might draw the conclusion that strong tax progressivity is efficient, see, e.g., Saez (2001) who relate optimal income tax rates to labor-supply elasticities.[26] However, our analysis suggests that this difference in labor-supply elasticities is over-estimated as individuals with high earnings on average also contribute larger *shares* to household earnings. In fact, when we apply our preferred interaction-term approach, estimated labor-supply elasticities are found to be more similar for both earnings groups, see columns (2) and (4) of Table 5.

[24] When we estimate an Altonji (1986) regression for women with $\overline{s} < 0.3$, we obtain a particularly large estimate in line with our extended model version with preference shocks.

[25] Due to assortative mating, men with high earnings also tend to have partners with above-average earnings. Nevertheless, men in the high-income group contribute larger average shares to household earnings (on average about 75% compared to 65%).

[26] The optimal tax rates derived by Saez (2001) use Marshall and Hicks labor-supply elasticities. In our model, Marshall and Hicks elasticities are monotonically increasing in the parameter η. Independent of the specific form of preferences, the Frisch elasticity is an upper bound for the other two elasticities. Saez (2001) considers both, preferences without income effects where the elasticities are identical as well as preferences with income effects. The shape of the optimal tax schedules is remarkably similar for both preference types and, hence, mostly determined by substitution effects.

7 Conclusion

Estimates of Frisch labor-supply elasticities are biased in presence of borrowing constraints. We have shown that the strength of this bias depends on individuals' relative contribution to household earnings. In couples with joint borrowing constraints, wage-rate fluctuations of secondary earners are less important for the couples' willingness to borrow and this relation is the stronger the more pronouced are intra-household wage differences. This results in smaller estimation biases for individuals who contribute little to household earnings. We have presented an incomplete-markets model with two earners to make this point explicit. We have used the model to develop a new method that corrects for the bias due to borrowing constraints. Specifically, we have extended standard Altonji (1986) regressions by the interaction between expected wage growth and the individual's usual contribution to household earnings. This estimation approach yields an unbiased estimate of the Frisch elasticity.

Empirically, we estimate a Frisch elasticity for men of about 0.7. This is larger than the majority of previous estimates from microeconometric studies. Further, we find rather homogenous labor-supply elasticities across the population compared to estimates from methods that neglect borrowing constraints and do not exploit the couple structure of the data.

References

Aaronson, D. and E. French (2009). The effects of progressive taxation on labor supply when hours and wages are jointly determined. *Journal of Human Resources 44*(2), 386–408.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Alesina, A., A. Ichino, and L. Karabarbounis (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy 3*, 1–40.

Alogoskoufis, G. S. (1987). On intertemporal substitution and aggregate labor supply. *Journal of Political Economy 95*(5), 938–960.

Altonji, J. and L. M. Segal (1996). Small-sample bias in GMM estimations of covariance structures. *Journal of Business & Economic Statistics 14*(3), 353–366.

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

Bayer, C. and F. Juessen (2012). On the dynamics of interstate migration: Migration costs and self-selection. *Review of Economic Dynamics 15*(3), 377–401.

Bick, A. (2016). The quantitative role of child care for female labor force participation and fertility. *Journal of the European Economic Association 14*(3), 639–668.

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), 515–524.

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

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters 91*(3), 312–320.

Clark, T. (1996). Small-sample properties of estimators of nonlinear models of covariance structure. *Journal of Business & Economic Statistics 14*(3), 367–373.

Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica 59*(5), 1221–48.

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*, 242–262.

Domeij, D. and P. Klein (2013). Should day care be subsidized? *Review of Economic Studies 80*(2), 568–595.

Guner, N., R. Kaygusuz, and G. Ventura (2012a). Taxation and household labour supply. *Review of Economic Studies 79*(3), 1113–1149.

Guner, N., R. Kaygusuz, and G. Ventura (2012b). Taxing women: A macroeconomic analysis. *Journal of Monetary Economics 59*(1), 111–128.

Guner, N., R. Kaygusuz, and G. Ventura (2014). Income taxation of U.S. households: Facts and parametric estimates. *Review of Economic Dynamics 17*(4), 559–581.

Haider, S. (2001). Earnings instability and earnings inequality of males in the United States: 1967-1991. *Journal of Labor Economics 19*(4), 799–836.

Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006. *Review of Economic Dynamics 13*(1), 15–51.

Heathcote, J., K. Storesletten, and G. L. Violante (2010). The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy 118*(4), 681–722.

Heathcote, J., K. Storesletten, and G. L. Violante (2014). Optimal tax progressivity: An analytical framework. NBER Working Paper 19899.

Heckman, J. J. (1993). What has been learned about labor supply in the past twenty years? *American Economic Review 83*(2), 116–121.

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

Kabukcuoglu, A. and E. Martinez-Garcia (2016). The market resources method for solving dynamic optimization problems. Globalization and Monetary Policy Institute Working Paper 274, Federal Reserve Bank of Dallas.

Keane, M. and R. Rogerson (2015). Reconciling micro and macro labor supply elasticities: A structural perspective. *Annual Review of Economics 7*(1), 89–117.

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

Kleven, H. J., C. T. Kreiner, and E. Saez (2009). The optimal income taxation of couples. *Econometrica 77*(2), 537–560.

Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? evidence and theory1. *The Review of Economic Studies 73*(1), 163–193.

Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies 69*(1), 209–244.

MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy 89*(6), 1059–1085.

MaCurdy, T. E. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics 18*(1), 83–114.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

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.

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.

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies 68*, 205–229.

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), 518–529.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics 51*(3), 609–633.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181.

Wu, C. and D. Krueger (2015). How much insurance in Bewley models with endogenous family labor supply? Working Paper, University of Pennsylvania.