

Bredemeier, Christian; Gravert, Jan; Juessen, Falko

Working Paper

Estimating Labor-Supply Elasticities with Joint Borrowing Constraints of Couples

IZA Discussion Papers, No. 10267

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Bredemeier, Christian; Gravert, Jan; Juessen, Falko (2016) : Estimating Labor-Supply Elasticities with Joint Borrowing Constraints of Couples, IZA Discussion Papers, No. 10267, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<http://hdl.handle.net/10419/147953>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 10267

Estimating Labor-Supply Elasticities with Joint Borrowing Constraints of Couples

Christian Bredemeier
Jan Gravert
Falko Juessen

October 2016

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*

Discussion Paper No. 10267

October 2016

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Estimating Labor-Supply Elasticities with Joint Borrowing Constraints of Couples*

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 than for primary earners. The reason is that, 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 than wage-rate changes of the primary earner. We illustrate the differential estimation bias in the framework of an incomplete-markets model with two-earner households and provide empirical support using PSID data. We show that Frisch elasticities can be estimated more consistently in samples of secondary earners. Our empirical results show that Frisch elasticities are larger than often reported in microeconomic studies. Further, we show that differences in labor-supply elasticities of men and women are overestimated when borrowing constraints are ignored.

JEL Classification: E24, J16, J22, E21

Keywords: labor-supply elasticity, incomplete markets, double-earner households

Corresponding author:

Falko Juessen
University of Wuppertal
Gaußstraße 20
42119 Wuppertal
Germany
E-mail: juessen@uni-wuppertal.de

* We would like to thank conference and seminar participants in Zurich, 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 labor supply to a one percent change in the net wage rate holding the marginal utility of wealth constant. This elasticity concept is important for understanding labor-supply reactions to, e.g., transitory tax or productivity shocks as well as to predictable life-cycle patterns in wage rates. It is a decisive determinant of, e.g., the fiscal multiplier and the costs of business cycles. The Frisch elasticity is also an upper bound for other (Marshall, Hicks) labor-supply elasticities which are important to evaluate the effects of, e.g., active labor-market policies and the design of optimal tax systems, see, e.g., Keane (2011).

Yet, there is no consensus on how elastic labor supply is and the micro and macro view on labor-supply elasticities differ markedly, see, e.g., Keane and Rogerson (2015). The literature has pointed to a number of econometric problems when estimating the Frisch elasticity, see, e.g., Blomquist (1985), Blomquist (1988), Alogoskoufis (1987), Heckman (1993), Rupert, Rogerson, and Wright (2000), or Imai and Keane (2004). An important estimation problem results from borrowing constraints. In the absence of borrowing constraints, the labor-supply reaction to expected wage changes is governed by the Frisch elasticity and hence identifies this elasticity. However, if households are borrowing constrained, they tend to raise, rather than to reduce, labor supply in reaction to negative wage-rate changes since they can no longer smooth consumption through borrowing. Domeij and Flodén (2006) have shown that, in a pooled sample of constrained and unconstrained households, the negative relation between wage-rate changes and labor supply in borrowing-constrained households biases the estimate of the Frisch elasticity downwards. 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 expected wage changes. 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 consider the case of a two-earner household and show that the bias due to borrowing constraints is less pronounced for secondary earners than it is for primary earners.¹ As a consequence, more consistent – and larger – estimates of labor-supply elasticities can be obtained from samples of secondary earners. Our estimation results can thus help reconcile micro and macro evidence on labor-supply elasticities.

¹In 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). In a recent paper, Blundell, Pistaferri, and Saporta-Eksten (2016) stress the importance of the double-earner perspective for understanding labor-supply dynamics.

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 We show that household consumption is more strongly correlated with wage-rate changes of the primary earner. The same holds for the household’s willingness to borrow which is key for the bias in estimated labor-supply elasticities. 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.³ As negative wage-rate shocks for the secondary earner can be relatively easily smoothed out, such shocks play a less important role for the household’s total wealth, and hence cause less pronounced surges in labor supply in borrowing constrained households. As a consequence, the downward bias of estimated Frisch labor-supply elasticities is less pronounced for secondary earners than it is for primary earners.

We first show this relation analytically in a simplified version of our model. Specifically, we show that the covariance between (expected) wage changes and the household’s willingness to borrow—and, hence, the resulting bias in the estimated Frisch elasticity—is smaller for secondary than for primary earners. We then solve the calibrated model numerically and use it to simulate a synthetic household panel data set on which we run standard labor-supply regressions. Quantitatively, we find that the estimate of the Frisch elasticity for male primary earners is about 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 male secondary earners is only about 20% below its true value. Hence, our analysis shows that more consistent estimates of labor-supply elasticities can be obtained from samples of secondary earners. We obtain almost unbiased estimates when a selection on earner status is combined with a mild selection on household assets.

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 expected wage changes as regressors that would yield

²This 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 spouse’s pre-marriage debt. Further, spouses’ credit ratings are mutually dependent.

³Blundell, Pistaferri, and Saporta-Eksten (2016) provide direct evidence that household consumption reacts more strongly to husbands’ wage shocks than to wives’ wage shocks.

unbiased estimates of the Frisch elasticity in the absence of borrowing constraints. Our theoretical 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 consist of only men or only women, respectively. Similarly, we address potential differences in true labor-supply elasticities of primary and secondary earners which may stem from, e.g., different roles in home production. To do so, we exploit information on household liquid assets which is available for a subsample of our data. 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 theory 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 estimate a Frisch elasticity of 0.47 for primary earners while the less biased estimate for secondary earners is 0.86 and hence about twice as large. These results can help reconcile micro and macro evidence on labor-supply elasticities, see, e.g. Keane (2011) and Keane and Rogerson (2015). The majority of microeconomic studies has estimated relatively small values for labor-supply elasticities which appears puzzling as estimates based on macro data are typically larger.⁴ Our analysis makes a case for larger elasticities: one reason why micro studies estimate relatively small elasticities is that these studies are usually based on samples of male household heads. These are individuals who tend to be primary earners in their households. We show that, in such samples, the downwards estimation bias is most pronounced. Looking at secondary earners suggests that true labor-supply elasticities are considerably larger.

A further implication of our analysis is that differences in labor-supply elasticities between population groups that differ in their usual earner roles in the household are in fact smaller than suggested by standard estimation techniques. An important example are gender differences. While standard techniques yield an estimate for the female elasticity which is about 85% larger than the one for men, our model can replicate the gender differences in estimated elasticities with true labor-supply elasticities of 0.72 for men and 0.95 for women,

⁴Small labor-supply elasticities would also suggest that efficiency costs of distortionary income taxation were small.

respectively, hence a difference of only about 32%. Similarly, when one pools individuals independent of earner status, we find that highly educated individuals seem to supply labor less elastically than less educated individuals - but this finding disappears when we restrict the sample to secondary earners.

The selection on earner status that we propose is often more feasible in empirical applications than a restriction on asset holdings, which can in principle also reduce the bias due to borrowing constraints. It is well known in the incomplete-markets literature that 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 earnings and working time. The distinction between primary and secondary earners is a practically feasible way to improve estimates of Frisch elasticities.⁵

The remainder of this paper is organized as follows. In Section 2, we develop an incomplete markets model with two earners. 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, distinguishing between primary and secondary earners. 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

2.1 Model set-up

The economy is populated by heterogeneous households consisting of two household members who live forever. 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, so that 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. This behavior is similar as in the model of Ortigueira and Siassi (2013).

Decision problem. The index i identifies 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

⁵In our empirical analysis, we also present estimations for subsamples around the years 1984, 1989, and 1994 for which we can exploit available information on household wealth, next to discriminating between primary and secondary earners. From 1999 onwards, regular information on asset holdings has been added to the PSID, but this is also the period where the PSID changed to biennial interviews. We will discuss later that the biennial frequency is problematic for labor-supply regressions.

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 $i = 1, 2$, the wage process is given by

$$\begin{aligned}\ln w_i &= \psi_i + z_i, \\ z'_i &= \rho_i \cdot z_i + \varepsilon'_i,\end{aligned}\tag{1}$$

where ψ_i is a fixed effect, $\psi_1 > \psi_2$. The stochastic wage component z_i follows an AR(1) process with innovations ε_i and autocorrelation ρ_i .⁶ Throughout, a variable with a prime denotes next period values.

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

$$u_i(c_i, n_i) = \frac{c_i^{1-\sigma_i} - 1}{1 - \sigma_i} - \alpha_i \cdot \frac{(n_i)^{1+1/\eta_i}}{1 + 1/\eta_i},\tag{2}$$

where σ_i , α_i , and η_i denote an individual's risk aversion, taste for leisure, and the (true) Frisch elasticity of labor supply. These preference parameters are allowed to differ between household members and, in our quantitative evaluations, we will account for potential gender differences in these parameters. In our baseline model, we assume that preferences are additively separable, which allows for a straightforward calibration of the Frisch elasticity through the parameter η_i . Wu und Krueger (2015) show that a model with separable preferences (2) explains more than 90% of the consumption insurance documented by Blundell, Pistaferri, and Saporta-Eksten (2016). For completeness, we consider an alternative model specification with non-separable preferences in Appendix C.1. This model yields similar results as our baseline model.

Since cooperative intra-household decision making leads to Pareto optimal allocations, the decision process can be represented by the decisions of a household planner. The planner 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,

$$v(c_1, c_2, n_1, n_2) = \mu \cdot u_1(c_1, n_1) + (1 - \mu) \cdot u_2(c_2, n_2).\tag{3}$$

Agents discount future utility with the discount factor $\beta \in (0, 1)$ and so does the household planner future values of v . We assume full commitment between the two household members. This means that, at any point in time, decisions reflect individuals' initial bargaining posi-

⁶Blundell, 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.

tions (spouses' Pareto weights in a household μ are constant over time). The weights could easily be endogenized as a solution to the couple's initial bargaining problem. Empirically, determinants of spouses' decision-making power can be identified by exploiting estimated household demand and supply functions as in, e.g., Chiappori, Fortin, and Lacroix (2002) and Cherchye, De Rock, Lewbel, and Vermeulen (2015) or, alternatively, by exploiting stability conditions for the marriage market as in Cherchye, Demuynck, De Rock, and Vermeulen (2016). For simplicity, we treat the Pareto weights as parameters, so that our model is de facto a unitary model. This assumption does not affect our results because the Frisch elasticity is identified in labor-supply regressions in first log differences where the Pareto weights will cancel out. Even allowing for endogenous time-varying Pareto weights in the spirit of a limited-commitment model 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 wages. We will come back to this point when we outline the estimation procedure for recovering the Frisch elasticity.

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.⁷ There is a restriction on borrowing, i.e., assets cannot fall below an exogenous minimum level denoted by a_{\min} :

$$a' \geq a_{\min}. \quad (4)$$

The household budget constraint is given by

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

where a' are next period asset holdings to be determined and purchased at price $1/(1+r)$ this period. In our baseline model, we assume that consumption is a household public good, i.e., there is no rivalry between spouses, $c_1 = c_2 = c$. 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 a model with private consumption, commitment

⁷A partial-equilibrium set-up is sufficient because we neither analyze policy nor parameter changes. As Domeij and Flodén (2006), we assume $\beta(1+r) < 1$.

implies that the ratio of spouses's marginal utilities is constant such that a constant marginal utility of wealth for one partner also implies constant marginal utility of wealth for the spouse. Thus, in our context, models with public or private consumption yield very similar results, which we document in Appendix C.2.

In our baseline model, we abstract from non-linear labor-income taxation. This assumption allows to recover the true Frisch elasticities consistently if borrowing constraints were absent. For completeness, we also consider a model version with progressive joint taxation of spouses in Appendix C.4. In this model, we obtain similar results as in our baseline economy.

Recursive formulation. The state space of the recursive maximization problem is $\Omega = [a_{\min}, a_{\max}] \times S$, 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 $\omega = (w_1, w_2)$.⁸

The household problem in recursive formulation is given by

$$V(a, \omega) = \max_{a', c, n_1, n_2} \mu \cdot u_1(c, n_1) + (1 - \mu) \cdot u_2(c, n_2) + \beta \mathbb{E} [V(a', \omega') | \omega], \quad (6)$$

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

$$x = x(a, \omega), \quad (7)$$

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

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

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

$$\phi = \lambda - (1 + r) \beta \mathbb{E} [\lambda' | \omega], \quad (9)$$

$$\lambda \cdot w_1 = \mu \cdot \alpha_1 \cdot n_1^{1/\eta_1}, \quad (10)$$

$$\lambda \cdot w_2 = (1 - \mu) \cdot \alpha_2 \cdot n_2^{1/\eta_2}, \quad (11)$$

$$\phi \geq 0, \quad (12)$$

$$a' \geq a_{\min}, \quad (13)$$

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

⁸The 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 be never reached.

together with the budget constraint (5), given the 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 (4) and λ is the Lagrange multiplier on the budget constraint (5). Condition (8) reflects that the household equalizes marginal utility of consumption and marginal utility of wealth. Condition (9) is the consumption Euler equation of the household which takes its standard form if the borrowing constraint does not bind, $\phi = 0$, and otherwise determines the household's willingness to borrow. 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. However, the weights do not impact on *changes* in labor supply (the same holds for α_1 and α_2). Conditions (12)-(14) are the Kuhn-Tucker conditions associated with the borrowing constraint (4).

From conditions (10) and (11), 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.

2.2 Estimation equation and estimation bias

Empirical studies estimate the relation between expected 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 $i = 1, 2$)⁹

$$\Delta \ln w'_i = \frac{1}{\eta} \cdot \Delta \ln n_i - \Delta \ln \lambda'. \quad (15)$$

A log-linear approximation of the Euler equation (9) implies

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

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

$$\Delta \ln n'_i = \eta \cdot \Delta \ln w'_i - \eta \cdot \ln \beta - \eta \cdot \ln(1+r) - \eta \cdot \frac{\phi}{\lambda} + \eta \cdot \xi'. \quad (17)$$

There are two econometric problems when estimating the labor-supply relation (17). First, independent of borrowing constraints, the expectation error ξ' which is part of the residual in

⁹Through first-differencing, the constant terms $\ln \mu$, $\ln(1-\mu)$, and $\ln \alpha_i$ drop out. When we estimate the labor-supply regression from real-world data, we come back to the role of these terms.

(17) is correlated with the regressor $\Delta \ln w'_i$: unexpected wage changes lead to expectation errors and hence $\Delta \ln w'$ and ξ' are correlated. Second, the multiplier ratio ϕ/λ which measures the household's willingness to borrow in terms of consumption goods¹⁰ affects labor supply but is not observable in empirical data and is therefore part of the combined residual. This causes an estimation bias because the willingness to borrow is correlated with wage growth.

The first problem is well-known in the literature (see, e.g., Blundell and MaCurdy 1999 or Keane 2011) and can be resolved by estimating the labor-supply reaction to *expected* wage changes which are obviously uncorrelated with the expectation error. The labor-supply reaction to expected wage changes identifies the Frisch elasticity when borrowing is unconstrained such that expected wage changes have no impact on the marginal utility of wealth, see, e.g., Keane (2011).¹¹ Another advantage of using expected wage changes is that, under limited commitment (e.g., Ligon, Thomas, and Worrall 2002), the Pareto weights change mostly in response to unexpected changes. Expectations as of period t , such as $E_t \Delta \ln w_{t+1}$, are already incorporated in the bargaining weights of period t . In an empirical application, using expected wage changes as regressor also solves problems of measurement error in the wage data, see, e.g., Keane (2011). In addition, expected wage changes are likely uncorrelated to transitory taste-shifters that may survive first-differencing.

When using synthetic data, we can calculate the expected wage change $E \Delta \ln w'_i$ as

$$E \Delta \ln w'_i = E(z'_i - z_i) = E((\rho - 1) \cdot z_i + \varepsilon'_i) = (\rho - 1) \cdot z_i, \quad (18)$$

while in real-world data, expected wage changes can be obtained through a first-stage regression using variables as regressors which are known to the agent in advance, see MaCurdy (1981), Altonji (1986), and Keane (2011).¹²

The focus of our study is on addressing the second estimation problem that originates from the presence of borrowing constraints. Formally, the multiplier ratio ϕ/λ which is part of the combined residual in (17) is positively correlated with the regressor $\Delta \ln w'$, leading to a biased estimate of the Frisch elasticity. This problem is not resolved by using expected wage changes as regressors. The willingness to borrow is correlated with precisely the expected component of wage changes as expected wage growth induces households to wish to front-

¹⁰ ϕ expresses the willingness to borrow in terms of utils and $1/\lambda$ converts utils into goods.

¹¹An alternative is to identify the Frisch elasticity from the covariance between labor income and transitory wage components in a GMM approach, see Blundell, Pistaferri, and Saporta-Eksten (2016). Yet, when a household is borrowing constrained, also transitory wage fluctuations impact on consumption and the marginal utility which biases the estimated Frisch elasticity.

¹²Such first-stage regression does not identify all expected wage changes but the identified wage changes can surely be expected, which is critical for the identification of the Frisch elasticity.

load consumption through borrowing.¹³ We will now show that the associated estimation bias differs between primary and secondary earners.

2.3 Analytical results

We now demonstrate analytically that the bias in estimated labor-supply elasticities which results from borrowing constraints is more pronounced for primary earners than it is for secondary earners. First, we show that the percentage estimation bias depends on the covariance between the multiplier ratio ϕ/λ on the one hand and the expected wage change or the actual wage change on the other hand, depending on whether the former is known to the econometrician. Then, we show that these covariances are larger for the primary earner than for the secondary earner.

Proposition 1 *When the expected wage change $E \Delta \ln w'$ is known, the percentage bias in the estimate for the Frisch elasticity is*

$$\frac{\hat{\eta} - \eta}{\eta} = -\frac{1 + \rho}{(1 - \rho) \cdot \sigma_{\varepsilon}^2} \cdot \text{cov}(\phi/\lambda, E \Delta \ln w').$$

If expected wage changes are identified using an instrument x , the bias is

$$\frac{\hat{\eta} - \eta}{\eta} = -\frac{1}{\hat{\gamma}^2 \cdot \text{var}(x)} \cdot (\text{cov}(\phi/\lambda, \Delta \ln w') - \text{cov}(\phi/\lambda, \nu)),$$

where $\hat{\gamma}$ is the estimated coefficient on the instrument in the first-stage regression and ν is the residual from the first-stage regression.

Proof. See appendix. ■

For the following analytical results, we consider a simplified version of the model, where we assume $\eta_1 = \eta_2 = \eta$, $\alpha_1 = \alpha_2 = 2$, $\mu = 1/2$, $a_{\min} = 0$, $\rho = 0$, and $\text{var}(\varepsilon_1) = \text{var}(\varepsilon_2)$. First, we can state that the multiplier ratio in the true labor-supply relation (17) is weakly decreasing in a symmetric, increasing function of spouses' wage levels:

Lemma 1 *Define $\Lambda = w_1^{1+\eta} + w_2^{1+\eta}$. The multiplier ratio ϕ/λ is weakly decreasing in Λ ,*

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

Proof. See appendix. ■

¹³The bias due to borrowing constraints is closely connected to an approximation error which stems from the approximation of the Euler equation and non-linear decision rules. Domeij and Flóden (2006) show that the resulting bias is small compared to the one due to borrowing constraints. They also show that, when households are not borrowing constrained, the approximation bias is negligible in yearly data.

Second, the symmetric function Λ of the spouses' wage rates has a lower covariance with the secondary earner's stochastic wage component which translates to the covariance of the multiplier ratio with the stochastic wage component:

Lemma 2 *The multiplier ratio ϕ/λ has, in absolute terms, a larger covariance with the primary earner's stochastic wage component than it has with the secondary earner's stochastic wage component,*

$$\text{cov}(\phi/\lambda, z_1) < \text{cov}(\phi/\lambda, z_2) < 0.$$

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 $E \Delta \ln w'_i$, respectively, of the two spouses on the other hand.

Proposition 2 *The multiplier ratio ϕ/λ has a larger covariance with the primary earner's next-period wage-rate change than it has with the secondary earner's next-period wage-rate change,*

$$\text{cov}(\phi/\lambda, \Delta \ln w'_1) > \text{cov}(\phi/\lambda, \Delta \ln w'_2) > 0.$$

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

$$\text{cov}(\phi/\lambda, E \Delta \ln w'_1) > \text{cov}(\phi/\lambda, E \Delta \ln w'_2) > 0.$$

Proof. See appendix. ■

Propositions 1 and 2 can easily be combined to the following statements about the relative biases in the labor-supply estimations for primary and secondary earners, respectively:

Proposition 3 *If the expected wage change is known, the percentage bias in $\hat{\eta}$ is, in absolute terms, unambiguously smaller for the secondary earner than for the primary earner. If expected wage changes are identified using a first-stage regression, the percentage downward bias in $\hat{\eta}$ is smaller for the secondary earner than for the primary earner if the variance of the instrument, the first-stage coefficient, and the covariance of the first-stage residuals with the multiplier ratio do not differ too strongly between primary and secondary earner.*

Proof. Follows directly from Proposition 1 together with Proposition 2. ■

3 Numerical analysis

In this section, we relax the simplifying assumptions from the previous analysis and solve the full model numerically. We describe the calibration of the model in Section 3.1 and in Section 3.2, we present policy functions to demonstrate the mechanisms behind our results. Thereafter, in Section 4, we will use the model to simulate a synthetic panel data set of labor-supply decisions which we use for labor-supply regressions.

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 microeconomic literature on labor-supply elasticities, see, e.g., MaCurdy (1981) and Altonji (1986), and has also been used by Domeij and Flodén (2006). Our baseline sample uses observations from the years 1972-1997.¹⁴ Note that our baseline sample is relatively large in comparison to related studies. For instance, Domeij and Flodén (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. In our specifications where we exploit asset holdings to control for borrowing constraints, we will also consider this subsample of the PSID data, see Section 5.

Our sample selection is as follows. Due to our focus on double-earner households, we consider household heads and wives for whom both partners' wage rates are observed.¹⁵ 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. Further, we apply similar sample selection criteria as Altonji (1986) and Domeij and Flodén (2006).¹⁶

Table 1 summarizes our calibration. We estimate the parameters of the wage processes separately for men and women, using a Generalized Method of Moments (GMM) estimator. We first obtain residual wages by filtering out deterministic cross-sectional variation using a first-stage OLS regression. We then identify auto-correlations, ρ_m for men and ρ_f for women, and innovation variances $\sigma_{m,\epsilon}^2$, $\sigma_{f,\epsilon}^2$ from gender-specific GMM estimations where we take into account a permanent component capturing individual-specific fixed characteristics and a

¹⁴Before 1972, there is no information on wives' education. After 1997, the PSID switched from annual to biennial interviews. Note that the frequency of the data set is relevant for the quality of the labor-supply estimation. At annual frequency, expectation errors are comparatively small such that the approximation bias arising from the approximation of the Euler equation is small. A second advantage of using annual instead of biennial or lower-frequency data is that expected wage changes are small which mitigates problems due to limited commitment and progressive taxation, respectively. In a limited commitment model in the style of Ligon, Thomas, and Worrall (2002), Pareto weights remain constant if changes in wages are sufficiently small. Under a progressive tax code, changes in marginal tax rates are small when wage changes are small and marginal tax rates may even be constant if the wage change is small enough not to induce a change in tax brackets.

¹⁵We do not restrict the analysis to married couples, i.e., we take into account cohabitation. In the PSID, a "wife" can also be the household head's girlfriend. Note that a household head is not necessarily also the primary earner. The PSID identifies the former mostly through gender while we identify the latter through wage rates relative to one's partner.

¹⁶We 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 drop household-years where individuals' reported annual hours of work are larger than 4860 (more than 92 average weekly hours) and where hours worked or the wage rate fall by more than 40 percent or increase by more than 250 percent between two consecutive years. To eliminate the influence of extreme observations and data errors, we drop observations falling in the top 0.5 percentiles of male and female wages, respectively.

Table 1: Parameter values.

Description	Parameter	Value	Determination
Autocorrelation wage process	ρ_m	0.84	estimated
	ρ_f	0.81	estimated
Standard deviation wage innovation	$\sigma_{m,\epsilon}$	0.24	estimated
	$\sigma_{f,\epsilon}$	0.44	estimated
Wage fixed effects	$\psi_{m,1}$	-0.03	calibrated
	$\psi_{m,2}$	-0.41	calibrated
	$\psi_{f,1}$	-0.30	calibrated
	$\psi_{f,2}$	-0.78	calibrated
Share of hh's with male primary earners		0.81	observed
Borrowing limit	a_{\min}	0	set
Discount factor	β	0.95	set
Relative risk aversion	σ	1.5	set
Specification 1			
Frisch labor-supply elasticity	η_m	1	set
	η_f	1	set
Labor disutility weight	α_m	14.7	calibrated
	α_f	13.0	calibrated
Risk-free interest rate	r	0.0115	calibrated
Specification 2			
Frisch labor-supply elasticity	η_m	0.72	calibrated
	η_f	0.95	calibrated
Labor disutility weight	α_m	21.5	calibrated
	α_f	13.3	calibrated
Risk-free interest rate	r	0.010	calibrated

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 estimated annual autocorrelations of $\rho_m = 0.84$ and $\rho_f = 0.81$, respectively. For men, the estimated standard deviation of idiosyncratic wage shocks is $\sigma_{m,\epsilon} = 0.24$ which is very similar to the estimate used by Domeij and Flodén (2006). For women, we estimate that idiosyncratic labor market risk is about twice that large. For completeness, we also simulated the model for identical stochastic wage processes z , in order to make sure that differences in estimation biases between population groups do not solely reflect differences in their respective wage dynamics. We find that this is not the case. This is also confirmed by our estimations within gender discussed in Section 5, where differences in wage dynamics are absent by construction. In 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 allow for off-grid choices by using spline interpolation of the value functions.

We simulate our model with households consisting of a man and a woman, each, and two types of households. In one type of households, the man is the primary earner while the woman is the primary earner in the other type. We set the population shares of the two household types to 81% and 19%, respectively, which accords to the frequencies of households with male primary earners and female primary earners, respectively, in our PSID sample.¹⁷ The wage fixed effects of the four types of agents (male primary earners, female primary earners, male secondary earners, female secondary earners) are calibrated to match relative conditional mean wage rates in the PSID data.¹⁸

In a first specification (Specification 1), we assume that the true Frisch elasticity is identical for both household members, and given by the preference parameter $\eta = 1$ which allows to interpret differences in estimated elasticities as differential biases. In Specification 2, which is our preferred specification, we assume that the true Frisch elasticities differ by gender, with women having a higher Frisch elasticity than men. Specifically, we calibrate the gender-specific values for the Frisch elasticities so that the *estimated* Frisch elasticities in our Monte

¹⁷For completeness, we also considered a simulation with one household type. Also there, we obtain smaller estimation biases for secondary than for primary earners, as in our main specification with two household types.

¹⁸For example, the average real wage rate of men who are primary earners expressed in units of the average real wage of all men is 1.058 in our PSID sample. The ergodic variance of male wages is given by $\frac{1}{1-\rho_m^2} \cdot \sigma_m^2 = 0.1886$, so that matching this target requires the wage fixed effect of male primary earners, $\psi_{m,1}$, to satisfy $1.058 = \exp(\psi_{m,1} + 0.1886/2)$. This gives $\psi_{m,1} = \ln 1.058 - 0.1886/2 = -0.0379$. We proceed analogously for male secondary earners, female primary earners, and female secondary earners.

Carlo study coincide with the *estimated* Frisch elasticities for men and women that we obtain using PSID data (see Section 5). We will discuss in Section 6 that one needs only relatively small differences in the true gender-specific Frisch elasticities ($\eta_m = 0.72$ and $\eta_f = 0.95$) 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, respectively.

The utility parameters measuring the taste for leisure, α_m and α_f , are set to match gender-specific mean hours worked. In our PSID sample, men’s average yearly hours amount to 2260 and women work 1574 hours on average. Targeting $1/3$ as average working time for men implies $0.7 \cdot 1/3 = 0.23$ as the target for female hours.

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. Relative risk aversion is set to $\sigma = 1.5$ for both household members.¹⁹ Following Domeij and Flodén (2006), we set $\beta = 0.95$ (understanding a period as one year), assume a borrowing limit of $a_{\min} = 0$, and calibrate the interest rate so that the bottom 40% of the wealth distribution own 1.4% of total wealth. When we vary model parameters in alternative specifications, e.g., when we vary η_m and η_f , we re-calibrate the parameters α_m , α_f , and r so that the targeted statistics are matched in all specifications.

3.2 Policy functions


We now present policy functions to illustrate the mechanisms leading to the key property of our model that wage-rate fluctuations of the secondary earner affect the household’s willingness to borrow less strongly than those of the primary earner. To save on space, we show policy functions only for the household type where the man is the primary earner and for the case of homogeneous true Frisch elasticities (policy functions for the other household type and for the case of heterogeneous Frisch elasticities are qualitatively similar).²⁰

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, for a given wage rate of the primary earner. In the figures, we choose the level of the partner’s wage rate in order to highlight the role of earner status. Specifically, for the primary earner’s labor-supply curve, we hold the partner’s wage constant

¹⁹We also considered a specification with differences in risk aversion between household members. The results for this model version are very similar to the ones obtained from our baseline specification.

²⁰We also solved the model for the case of non-separable preferences. Appendix C.1 shows the policy functions for this model version.

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




Notes: 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.

at its lowest value on the wage grid. Vice versa, for the secondary earner's labor-supply curve, 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 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. From the left panel, it can be seen that, for primary earners, household asset holdings are relevant also for the shape of the labor-supply curves. The labor-supply curve is globally upward-sloping if the household is wealth-rich (dashed line). By contrast, 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) because, at the borrowing constraint, consumption can no longer be smoothed through dissaving. This gives rise to a non-standard shape of the labor-supply function which is similar to the one in the model of Domeij and Flodén (2006), where households consist 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. To see how this affects labor-supply estimations, consider expectations when low wage rates are realized at the borrowing constraint. When the primary

Figure 2: Cross-wage labor-supply curves.




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

earner's wage rate is below average, the borrowing-constrained household expects the primary earner's wage rate to rise and the primary earner's labor supply to fall. This 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.

By contrast, 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, 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 different shapes of the labor-supply curves for primary and secondary earners are 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 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. Thus, in particular for primary earners, a negative wage-rate shock can actually lead to an increase (rather than to a decrease) of

Figure 3: Consumption policy functions.



Notes: 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.

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 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 closely related to the household's willingness to borrow. Figure 3 shows household 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 . 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. 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. Blundell, Pistaferri, and Saporta-Eksten (2016) provide direct evidence for this key piece of our mechanism. They show that, in couple households, consumption reacts more strongly to wage fluctuations of husbands than to wage fluctuations of wives. The distinction between husbands and wives is similar to our distinction between primary and secondary earners since most husbands are primary earners.²¹

Thus, in wealth-poor households, consumption is most strongly correlated with the primary earner’s wage-rate realizations. Since it is the marginal utility of consumption that determines the household’s willingness to borrow, see equation (9), the strongest covariance between the shadow price on the borrowing constraint and wage rates occurs for primary earners which also translates into a stronger covariance between the household’s willingness to borrow and the primary earner’s (expected) wage growth, as we have shown analytically in the simplified model version in Section 2.3 (Lemma 2 and Proposition 2). 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 it is 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 which 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

$$\Delta \ln n_{ijt+1} = \text{const.} + \eta \cdot \mathbb{E}_t \Delta \ln w_{ijt+1} + u_{ijt+1}, \quad (19)$$

²¹Blundell, Pistaferri, and Saporta-Eksten (2016) focus on the response to permanent wage shocks while the wage shocks we analyze are not fully persistent. However, when a household is borrowing constrained, even transitory wage shocks cannot be smoothed through borrowing which makes them similar to permanent wage shocks in this respect. Blundell, Pistaferri, and Saporta-Eksten (2016) also report that the cross-wage reaction is stronger for wives than for husbands, in line with our cross-wage labor-supply functions in Figure 2.

where we introduced household and time indices to clarify the panel dimension of the estimation. The index $ijt + 1$ refers to member i of household j in period $t + 1$. As explained before, we use the expected wage change as regressor to eliminate the correlation with the expectation error which is part of u_{ijt+1} . In the simulations, we determine the expected wage change as $E_t \Delta \ln w_{ijt+1} = (\rho_j - 1) \cdot z_{ijt}$. The constant term in (19) captures the effects of the rate of time preference, interest rates, and the mean willingness to borrow while u_{ijt+1} is the combined residual capturing the effect of expectation errors (in wage growth and the marginal utility of wealth) and the deviation of the willingness to borrow from its mean, see (17) for the theoretical counterpart to (19).

4.1 Simulation setup

We simulate a synthetic panel data set with similar size as our baseline PSID sample. Specifically, we simulate $N = 2,500$ households for a long period of time and keep the final $T = 7$ periods of this simulation, corresponding to the average number of years for which we observe a household in the PSID. In total, this gives $2,500 \text{ households} \times 2 \text{ persons} \times 6 \text{ periods}$ (for which first differences can be calculated) = $30,000$ observations. For completeness, we also considered simulations with a substantially larger number of households and time periods. Results are similar to the ones discussed in the main text. We identify an individual's earner status by comparing spouses' relative mean realized wage rates in our final estimation sample. This has the advantage that we can proceed analogously in our estimations based on real-world data from the PSID presented in Section 5.²²

4.2 Estimation results from synthetic data

We first consider the case where there are no differences in the true Frisch elasticity between household members, $\eta_m = \eta_f = 1$. 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 model population must be due to estimation biases. Column (1) of Table 2 shows the results for individuals identified as primary earners. In this sample, the estimate for the Frisch elasticity is 0.68 and is hence subject to a downward bias of 32%. In the sample of secondary earners, see column (2) of Table 2, the downward bias is about 16% and, hence, substantially smaller compared to the sample of primary earners.

²²In the simulated data, an alternative way of identifying primary and secondary earners is to condition on the fixed effects in the wage process. We do not choose this identification as our baseline because, in real-world data, there are no directly observable counterparts to the fixed effects. For completeness, we also performed the Monte Carlo estimations conditioning on the fixed effects. The results are similar to the ones discussed in the main text.

Table 2: Estimation results from synthetic household panel data.

	(1)	(2)	(3)	(4)	(5)	(6)
	men and women ($\eta = 1$)		men ($\eta = 0.72$)		women ($\eta = 0.95$)	
	primary earners	secondary earners	primary earners	secondary earners	primary earners	secondary earners
expected log wage change	0.68 (0.01)	0.84 (0.01)	0.38 (0.01)	0.58 (0.04)	0.71 (0.01)	0.78 (0.01)
true Frisch bias	1.00 -31.7%	1.00 -16.5%	0.72 -46.7%	0.72 -19.5%	0.95 -25.1%	0.95 -17.9%
observations	15,000	15,000	11,652	3,348	3,348	11,652

Notes: Regression results for equation (19). Constant included but not shown. $E \Delta \ln w' = (\rho - 1) \cdot z$. Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise.

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.

In Appendix C, we repeat our baseline estimations for alternative specifications of the model. Specifically, we consider the cases of private consumption, identical wage dynamics for men and women, and non-separable preferences. In addition, we consider a model version with progressive labor-income taxation. For all these model versions, we find that the estimation bias is less pronounced for secondary than for primary earners.

Yet, looking at estimated differences between primary and secondary earners may be misleading since, in the real world, these groups may differ in their true Frisch elasticities. In our PSID sample, 81% of secondary earners are women for whom the Frisch elasticity is often found to be larger than for men. Thus, when we obtain a higher estimate for the Frisch elasticity in an *empirical* sample of secondary earners, this might simply reflect this difference in the true mean elasticity and does not necessarily prove a differential estimation bias. We address this point by distinguishing between primary and secondary earners *within genders*.²³ Our model predicts that the estimated labor-supply elasticities differ also between primary and secondary earners of the same gender. Other than differences between primary and secondary earners of both genders, these within-gender differences cannot be driven by gender differences in the true elasticities. With this specification, we can also rule out that differences in estimated elasticities are driven by biases which are specific to gender, for

²³In our model, men and women differ from one another in terms of the parameters of the stochastic wage process and, potentially, preference parameters.

example due to different wage processes, rather than earner status. In the next section, we test this within-gender prediction using PSID data. Here, we first demonstrate it within our model.

Columns (3)-(6) in Table 2 summarize the results for the within-gender estimations, where we assumed in the data generating process that the true Frisch elasticities differ across gender ($\eta_m = 0.72$ and $\eta_f = 0.95$). For both men and women, we find that the estimated elasticities are larger for secondary than for primary earners although the true elasticities differ only by gender but not within gender. Put differently, since there is no within-gender variation in the true elasticities, the within-gender variation in the estimated elasticities solely reflects that the estimation bias is smaller for secondary than for primary earners. The most pronounced bias, -47% , is found for male primary earners due to their particularly strong relative wage position within the household. When we look at secondary earners, by contrast, we find that the bias is only -20% for male secondary earners and -18% for female secondary earners. Note that many previous studies have restricted the sample to males working full-time, thereby de facto using a sample of mostly primary earners. Our results help understand why these studies tend to find relatively small estimates for the labor-supply elasticity. Estimates from samples of secondary earners are likely larger which helps reconcile the micro and macro view on labor-supply elasticities, see, e.g. Keane and Rogerson (2015).

Still, in real-world data, one may argue that, even within gender, the true labor-supply elasticities of secondary earners might be larger than the ones of primary earners. For instance, due to household specialization, secondary earners can be expected to spend more time in home production which impacts on the true labor-supply elasticities, see, e.g., Alesina, Ichino, and Karabarbounis (2011). In the estimations based on PSID data discussed in the next section, we will address this issue. Here, we foreshadow and demonstrate our approach using simulated data from our model. In the spirit of a difference-in-differences approach, we exploit the prediction that the *differences* in estimation biases between primary and secondary earners become smaller in samples where borrowing constraints are less important. To illustrate this, we reconsider our previous within-gender estimations but additionally apply sample splits conditional on household asset holdings. We then investigate the ratio of estimated labor-supply elasticities of secondary and primary earners, $\hat{\eta}_{sec}/\hat{\eta}_{prim}$ within the group of men and conditional on assets.

Table 3 summarizes the results of this exercise. To be consistent with our estimations using PSID data, we focus on samples of men (results for women are similar). In the simulated data, the Frisch elasticities differ by gender but are identical within gender, so that the

Table 3: Estimation results from synthetic household panel data (men, conditional on assets).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	all men		assets > 0		assets > 0.01 · \bar{a}		assets > 0.1 · \bar{a}	
	primary earners	secondary earners	primary earners	secondary earners	primary earners	secondary earners	primary earners	secondary earners
expected log wage change	0.38 (0.01)	0.58 (0.04)	0.52 (0.02)	0.66 (0.03)	0.59 (0.02)	0.70 (0.03)	0.69 (0.02)	0.71 (0.03)
true Frisch bias	0.72 -46.7%	0.72 -19.5%	0.72 -27.8%	0.72 -8.9%	0.72 -17.8%	0.72 -2.6%	0.72 -4.2%	0.72 -0.8%
$\hat{\eta}_{sec}/\hat{\eta}_{prim}$	1.51		1.26		1.19		1.04	
observations	11,652	3,348	9,679	3,134	8,867	3,039	6,869	2,697

Notes: Regression results for equation (19). Constant included but not shown. $E \Delta \ln w' = (\rho - 1) \cdot z$. Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise. \bar{a} denotes mean asset holdings in the economy.

true secondary-primary earner ratio within men is $\eta_{sec}/\eta_{prim} = 1$. In the full sample of all men, columns (1) and (2), differences in estimated labor-supply elasticities for primary and secondary earners are pronounced. The ratio in the estimates amounts to $\hat{\eta}_{sec}/\hat{\eta}_{prim} = 1.51$. When we restrict the sample to households with positive asset holdings, see columns (3) and (4), point estimates for the Frisch elasticities increase which reflects the lower importance of borrowing constraints and of the associated downward bias. Importantly, the ratio of the estimated elasticities decreases because the downward bias was larger for primary earners in the first place. When we further increase the asset requirement in the sample selection, see columns (5) and (6) as well as (7) and (8), the ratio $\hat{\eta}_{sec}/\hat{\eta}_{prim}$ gets closer to its true value of 1.

In line with Domeij and Flodén (2006), we obtain almost unbiased estimates when we condition on asset holdings being sufficiently large. Note that Domeij and Flodén (2006) report a similarly unbiased estimate only when applying a considerably stronger asset selection (they condition on asset holdings being larger than $1 \cdot \bar{a}$). 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 issues of sample size. Thus, the distinction between primary and secondary earners which we propose is a practically feasible alternative to improve estimates of the Frisch elasticity.

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 expected wage changes 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 is likely to 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. To identify primary and secondary earners, we apply the same restriction on realized wage rates as in the Monte Carlo study, i.e., the person with the higher average wage over the sample period is classified as the primary earner in the household.

5.1 Econometric aspects

We estimate the labor-supply regression (19) from PSID data. As in our theoretical model, we analyze the choice of hours worked at the intensive margin in double-earner households. In our sample, the standard deviation of annual hours growth, which is the left-hand side variable in (19), is 17.3% for primary earners and 24.5% for secondary earners, respectively. Thus, the data show that there is substantial variation at the intensive margin for both earner types.

We use individual characteristics to identify expected future wage changes, $E_t \Delta w_{ijt+1}$, in gender-specific first-stage 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.²⁴ As pointed out by Keane (2011, p. 1012), the purpose of these instruments is to obtain predicted wage growth over the life cycle which is needed to identify how hours respond to predictable variation of wages. If there were no borrowing constraints, such changes would leave the marginal utility of wealth unchanged and would hence identify the Frisch elasticity. Using individual characteristics

²⁴In 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 cannot 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).

as instruments has the advantage that measurement error in these variables is uncorrelated with measurement error in wage rates.²⁵ As discussed by Altonji (1986), Domeij and Flodén (2006) and Keane (2011), the instruments in the labor-supply regression are potentially weak, one reason being that most wage changes may simply be unexpected.²⁶ 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.²⁷

In empirical data, individual labor supply may also be affected by taste shifters.²⁸ 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. Yet, in Appendix C.4, we present Monte-Carlo estimations for a model version with progressive income taxation which show that, in our context, the biases due to progressive taxes 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 does not change substantially from year to year, see, e.g., Altonji (1986).

5.2 Estimation results

Table 4 summarizes the results of a first set of labor-supply regressions using PSID data. In line with our theoretical analysis, we find that the estimated Frisch elasticity is larger in the sample of secondary earners compared to the sample of primary earners, see columns (1) and (2). Quantitatively, the estimated elasticity for secondary earners is about twice as large as the one for primary earners. As an alternative regression specification, we include an

²⁵Measurement error in hours on the left-hand side of the regression reduces the R^2 of the regression but does not impact on the consistency of the estimate for η .

²⁶A 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 (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.

²⁷In 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.

²⁸Technically, the labor-supply relation in our model, (17), also includes 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 4: Empirical labor-supply regressions, PSID data.

	(1)	(2)	(3)
	primary earners	secondary earners	all
expected log wage change	0.47 (0.11)	0.86 (0.16)	0.93 (0.10)
expected log wage change × primary earner			-0.52 (0.05)
$\hat{\eta}_{sec}/\hat{\eta}_{prim}$		1.83	2.28
time effects	yes	yes	yes
observations	14,340	14,340	28,680

Notes: Regression results for equation (19). $E_t \Delta \ln w_{ijt+1}$ determined in a first-stage regression using individual characteristics. Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise. In column (3), $\hat{\eta}_{sec}$ corresponds to the coefficient on expected wage growth while $\hat{\eta}_{prim}$ corresponds to the sum of the coefficients on expected wage growth and the interaction term.

interaction term between expected wage changes and earner status in a pooled sample of both groups.²⁹ Specifically, column (3) shows the results from a specification where we interact expected wage growth with a dummy variable for being primary earner. As expected, the coefficient on the interaction term is negative and statistically significant. This corroborates that labor-supply elasticities are estimated to be substantially smaller for primary than for secondary earners when borrowing constraints are ignored.

Our results suggest that the usual sample restriction to, e.g., male household heads working full-time is potentially problematic in microeconomic estimations of the labor-supply elasticity. In fact, such samples consist mostly of primary earners, 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 perhaps 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. Notably, we also obtain this finding in a sample of *male* secondary earners that we consider below. Distinguishing between primary and secondary earners can thus help to reconcile micro and macro estimates of labor-supply

²⁹In the Monte Carlo study, this regression approach yields virtually identical results as applying sample splits.

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

	(1)	(2)	(3)	(4)	(5)	(6)
	men			women		
	primary earners	secondary earners	all	primary earners	secondary earners	all
expected log wage change	0.35 (0.11)	0.54 (0.27)	0.52 (0.12)	0.25 (0.39)	0.90 (0.19)	0.87 (0.17)
expected log wage change × primary earner			-0.15 (0.09)			-0.45 (0.10)
$\hat{\eta}_{sec}/\hat{\eta}_{prim}$		1.54	1.42		3.55	2.07
time effects	yes	yes	yes	yes	yes	yes
observations	11,632	2,708	14,340	2,708	11,632	14,340

Notes: Regression results for equation (19). $E_t \Delta \ln w_{ijt+1}$ determined in a first-stage regression using individual characteristics. Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise. In columns (3) and (6), $\hat{\eta}_{sec}$ corresponds to the coefficient on expected wage growth while $\hat{\eta}_{prim}$ corresponds to the sum of the coefficients on expected wage growth and the interaction term.

elasticities (Keane and Rogerson 2015).

We now provide additional evidence in order to show that the estimated differences in the 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. To address deep gender differences, we perform the labor-supply regressions separately for men and women, as we did in the Monte Carlo study. Differences in estimated coefficients between primary and secondary earners of the same gender cannot result from gender differences in true elasticities or any gender-specific (rather than earner-role-specific) biases. Table 5 shows 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. The estimate for male secondary earners is 0.54, thus more than 50% larger than the one for male primary earners, see columns (1) and (2) of Table 5. When we repeat the within-gender evaluation for women, we find even larger differences in the estimates for primary and secondary earners, respectively, see columns (4) and (5) of Table 5. These results are confirmed by the alternative specification with interaction terms between expected wage growth and being primary earner, see columns (3) and (6).

As a final step, and to corroborate the differential role of borrowing constraints for labor

supply of primary and secondary earners, respectively, we address the possibility that there are differences in the true Frisch elasticities between primary and secondary earners even within the same gender. As argued before, an example is that a male secondary earner may contribute more to home production than a male primary earner, leading to a higher true Frisch elasticity for male secondary earners. To address this point, we extend our empirical analysis to take into account information on household wealth, which can serve as an inverse proxy of the bindingness of the borrowing constraints. Our theory implies that, even if there are differences in the true Frisch elasticities between primary and secondary earners, the ratio $\hat{\eta}_{sec}/\hat{\eta}_{prim}$ of the estimated elasticities should be larger in wealth-poor households where estimates for primary earners are reduced relative to those for secondary earners due to the stronger estimation bias.

As discussed by Domeij and Flodén (2006), the 1984, 1989, and 1994 PSID waves contain a supplement on household wealth.³⁰ 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. We also follow Domeij and Flodén (2006) in restricting the sample to male individuals. Within this sample, we distinguish between primary and secondary earners.

In line with the literature, we use liquid assets as the relevant indicator of a household's ability to smooth consumption. In the PSID, we measure liquid assets as the sum of checking and savings account balances, bonds, and stocks minus other debts, such as credit card debt, medical or legal bills, or loans from relatives. We apply three different selections based on household liquid assets, requiring asset holdings to be positive, above half an average full-time male monthly income, or above one average full-time male monthly income, respectively.

Table 6 summarizes the estimation results from regressions using interaction terms, conditional on liquid assets.³¹ When we do not control for borrowing constraints, the results imply that the estimated elasticity is about twice as large for secondary earners, see column (1) of Table 6. When we select wealthier households, see columns (2)-(4), estimated labor-supply elasticities for secondary earners decrease slightly which may be due to non-separable preferences or home production. Importantly, the estimated coefficient on the interaction term, measuring the difference between estimated elasticities between primary and secondary earn-

³⁰From 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.

³¹For completeness, Table 11 in Appendix D reports the results for the regressions using sample splits by assets and earner status. Both approaches yield similar results.

Table 6: Empirical labor-supply regressions for men, PSID data (conditional on liquid assets).

	(1)	(2)	(3)	(4)
	all men	assets > 0	assets > 0.5 · y^m	assets > y^m
expected log wage change	0.60 (0.20)	0.63 (0.23)	0.57 (0.23)	0.52 (0.24)
expected log wage change × primary earner	-0.35 (0.19)	-0.31 (0.22)	-0.25 (0.22)	-0.19 (0.23)
$\hat{\eta}_{sec}/\hat{\eta}_{prim}$	2.41	1.95	1.80	1.59
time effects	yes	yes	yes	yes
observations	9,780	6,886	6,023	5,511

Notes: Regression results for equation (19). $E_t \Delta \ln w_{ijt+1}$ determined in a first-stage regression using individual characteristics. y^m is an average full-time male monthly income, $y^m = \$2014$ (1983 dollars). Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise. $\hat{\eta}_{sec}$ corresponds to the coefficient on expected wage growth while $\hat{\eta}_{prim}$ corresponds to the sum of the coefficients on expected wage growth and the interaction term.

ers, is the smaller, the wealthier is the considered group of households. This corroborates that differences in the estimated labor-supply elasticities of primary and secondary earners stem from differential estimation biases due to borrowing constraints. Borrowing constraints are less relevant in wealthier households such that the implied biases shrink and estimated labor-supply elasticities for primary earners rise relative to those of secondary earners.

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. One example is the often-discussed difference in labor-supply elasticities between men and women. Another example is the difference in labor-supply elasticities between high skilled and low skilled individuals. 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.

To illustrate this, Table 7 considers estimated labor-supply elasticities by gender. When we do not distinguish between primary and secondary earners in our baseline PSID sample

Table 7: Empirical labor-supply regressions, PSID data (gender perspective).

	(1)	(2)	(3)	(4)
	PSID		model	
	men	women	men	women
expected log wage change	0.41 (0.10)	0.78 (0.17)	0.41 (0.01)	0.78 (0.01)
$\hat{\eta}_f/\hat{\eta}_m$	1.91		1.91	
true Frisch bias	—	—	0.72 -42.8%	0.95 -18.1%
η_f/η_m	—	—	1.32	
time effects	yes	yes	—	—
observations	14,340	14,340	15,000	15,000

Notes: Regression results for equation (19). $E_t \Delta \ln w_{ijt+1}$ determined in a first-stage regression using individual characteristics. Standard errors in parentheses.

(1972-1997), the estimated female Frisch elasticity (0.78) exceeds the male one (0.41) by 90%, see columns (1) and (2) of Table 7. The implied gender ratio $\hat{\eta}_f/\hat{\eta}_m$ of about 2 is very similar to the ratio of estimates reported in the survey articles by Keane (2011) and Browning, Hansen, and Heckman (1999). In estimations comparable to ours, Blundell, Pistaferri, and Saporta-Eksten (2016) report elasticities which imply a gender ratio of about 1.6 for a more recent sample. We can use our model to decompose the estimated gender difference into the true difference in labor-supply elasticities between women and men and the differential estimation bias which arises because women are more often secondary earners than men. Note that we have calibrated our model to match the gender-specific estimates in our PSID sample. Thus, the gender-specific estimates based on synthetic model data are identical to the estimates from PSID data, see Table 7. The calibrated Frisch elasticities amount to $\eta_m = 0.72$ and $\eta_f = 0.95$, so that the difference in the *true* Frisch elasticities is only about 32%. This difference is magnified by a differential estimation bias leading to the targeted difference in *estimated* Frisch elasticities.

For similar reasons, conventional methods overestimate the differences in labor-supply elasticities between education groups. When we distinguish between individuals with at least some college education and those without, see columns (1) and (2) of Table 8, conventional labor-supply estimates suggest that high-skill labor is in considerably less elastic supply. Accordingly, one might draw the conclusion that it is efficient to tax the highly educated at

Table 8: Empirical labor-supply regressions, PSID data (by education).

	(1)	(2)	(3)	(4)
	all		secondary earners	
	low educ.	high educ.	low educ.	high educ.
expected log wage change	1.05 (0.26)	0.60 (0.16)	1.22 (0.32)	1.41 (0.43)
$\hat{\eta}_{low}/\hat{\eta}_{high}$		1.75		0.87
observations	11,809	16,871	6,559	7,781

Notes: Regression results for equation (19). $E_t \Delta \ln w_{ijt+1}$ determined in a first-stage regression using individual characteristics. Individuals identified as secondary earners if the mean realized wage rate in the sample \bar{w}_i is below the mean realized wage rate of the spouse \bar{w}_{-i} . High education is some college and above. We exclude education from the set of explanatory variables in the first-stage regression and use the number of children instead.

higher rates.³² However, our analysis suggests that this difference in estimated labor-supply elasticities is over-estimated as highly educated individuals are more often primary earners than low educated individuals are. In fact, when we restrict the sample to secondary earners, where estimates can be expected to be more consistent, estimated labor-supply elasticities are found to be similar for both education groups, see columns (3) and (4) of Table 8.

7 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 between primary and secondary earners. 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. As a consequence, differences in labor-supply elasticities between groups which differ in their earner roles in the household, such as genders, are overestimated by conventional estimation methods which ignore borrowing constraints.

³²There is a close relation between Frisch elasticities and the long-run (Marshall) elasticities which are relevant for optimal tax rates. In our model, Marshall elasticities are monotonically increasing in the parameter η . Independent of the specific form of preferences, the Frisch elasticity is an upper bound for the Marshall elasticity.

In an empirical application using PSID data, we have found support for these predictions. For primary earners, conventional methods yield relatively small estimates for Frisch labor-supply elasticities, in line with the literature. For secondary earners, we obtain substantially larger estimates, in line with our theoretical arguments. These results can help reconcile micro and macro evidence of the elasticity of labor supply. Regarding gender differences in labor-supply elasticities, we obtain estimates for women which are about twice as large as those for men, also in line with the literature. However, our quantitative model evaluations suggest that the true gender differences in labor-supply elasticities are substantially smaller, amounting only to about one third. Our model matches the empirical estimation results with true Frisch elasticities of 0.72 for men and 0.95 for women.

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.
- Alesina, A., A. Ichino, and L. Karabarbounis (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy* 3.
- 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.
- 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. and T. MaCurdy (1999). Labor Supply: A Review of Alternative Approaches. Volume 3A of *Handbook of Labor Economics*, Chapter 27, pp. 1559–1695. Elsevier.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.

- Browning, M., L. P. Hansen, and J. J. Heckman (1999). Micro data and general equilibrium models. In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Volume 1 of *Handbook of Macroeconomics*, Chapter 8, pp. 543–633. Elsevier.
- Cherchye, L., B. De Rock, A. Lewbel, and F. Vermeulen (2015). Sharing rule identification for general collective consumption models. *Econometrica* 83(5), 2001–2041.
- Cherchye, L., T. Demuynck, B. De Rock, and F. Vermeulen (2016). Household consumption when the marriage is stable. Working paper, University of Leuven.
- Chiappori, P.-A., B. Fortin, and G. Lacroix (2002). Marriage market, divorce legislation, and household labor supply. *Journal of Political Economy* 110(1), 37–72.
- Clark, T. (1996). Small-sample properties of estimators of nonlinear models of covariance structure. *Journal of Business & Economic Statistics* 14(3), 367–373.
- 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 Quarterly 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.
- 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., 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.
- 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.

- 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.
- Moffitt, R. A. and P. Gottschalk (1995). Trends in the covariance structure of earnings in the U.S., 1969-1987. Working paper, Johns Hopkins University.
- 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.
- 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.
- 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.

A Proofs

Proof of Proposition 1. If expected wage changes are known to the econometrician, the labor-supply regression is

$$\Delta \ln n' = \text{const} + \eta \cdot \text{E} \Delta \ln w' + u$$

and the estimate $\hat{\eta}$ is

$$\hat{\eta} = \frac{\text{cov}(\Delta \ln n', \text{E} \Delta \ln w')}{\text{var}(\text{E} \Delta \ln w')} = \frac{\text{cov}(\Delta \ln n', (\rho - 1)z)}{\text{var}((\rho - 1)z)} = \frac{\text{cov}(\Delta \ln n', z)}{(\rho - 1) \cdot \text{var}(z)}. \quad (20)$$

We rearrange the true labor-supply relation (17) to

$$\begin{aligned} \Delta \ln n' &= \kappa + \eta \cdot \Delta \ln w' - \eta \cdot \frac{\phi}{\lambda} + \tilde{\xi}' = \kappa + \eta \cdot (z' - z) - \eta \cdot \frac{\phi}{\lambda} + \tilde{\xi}' \\ &= \kappa + \eta \cdot (\rho z + \varepsilon - z) - \eta \cdot \frac{\phi}{\lambda} + \tilde{\xi}' = \kappa + \eta \cdot ((\rho - 1)z + \varepsilon) - \eta \cdot \frac{\phi}{\lambda} + \tilde{\xi}', \end{aligned}$$

where $\kappa = -\eta \cdot \ln \beta - \eta \cdot \ln(1 + r)$ and $\tilde{\xi}' = \eta \cdot \xi'$. Hence, the covariance between $\Delta \ln n'$ and z is

$$\begin{aligned} \text{cov}(\Delta \ln n', z) &= \text{cov}\left(\eta \cdot (\rho - 1)z - \eta \cdot \frac{\phi}{\lambda}, z\right) \\ &= \eta \cdot (\rho - 1) \cdot \text{var}(z) - \eta \cdot \text{cov}\left(\frac{\phi}{\lambda}, z\right). \end{aligned}$$

Inserting this into (20) gives the estimate as

$$\hat{\eta} = \frac{\eta \cdot (\rho - 1) \cdot \text{var}(z) - \eta \cdot \text{cov}\left(\frac{\phi}{\lambda}, z\right)}{(\rho - 1) \cdot \text{var}(z)} = \eta - \frac{\eta \cdot \text{cov}\left(\frac{\phi}{\lambda}, z\right)}{(\rho - 1) \text{var}(z)}.$$

The percentage bias then is

$$\begin{aligned} \frac{\hat{\eta} - \eta}{\eta} &= -\frac{\text{cov}\left(\frac{\phi}{\lambda}, z\right)}{(\rho - 1) \text{var}(z)} = -\frac{\text{cov}(\phi/\lambda, \text{E} \Delta \ln w' / (\rho - 1))}{(\rho - 1) \text{var}(z)} \\ &= -\frac{\text{cov}(\phi/\lambda, \text{E} \Delta \ln w')}{(\rho - 1)^2 \cdot \text{var}(z)} = -\frac{1 + \rho}{(1 - \rho) \cdot \sigma_\varepsilon^2} \cdot \text{cov}(\phi/\lambda, \text{E} \Delta \ln w'), \end{aligned}$$

where the last step uses $\text{var}(z) = 1/(1 - \rho^2) \cdot \sigma_\varepsilon^2$ and $1 - \rho^2 = (1 + \rho) \cdot (1 - \rho)$.

If expected wage changes are identified based on a first-stage regression, $\Delta \ln w' = const_1 + \gamma \cdot x + \nu$, where x is an instrument, the first-stage results are

$$\widehat{\gamma} = \frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)}, \quad \widehat{\Delta \ln w'} = const_1 + \frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)} \cdot x.$$

Then, the second-stage regression is

$$\Delta \ln n' = const + \eta \cdot \frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)} \cdot x + u,$$

where $const$ includes $\eta \cdot const_1$ and the estimated coefficient is

$$\widehat{\eta} = \frac{\text{cov}\left(\frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)} \cdot x, \Delta \ln n'\right)}{\text{var}\left(\frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)} \cdot x\right)} = \frac{\frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)} \cdot \text{cov}(x, \Delta \ln n')}{\left(\frac{\text{cov}(\Delta \ln w', x)}{\text{var}(x)}\right)^2 \cdot \text{var}(x)} = \frac{\text{cov}(x, \Delta \ln n')}{\text{cov}(\Delta \ln w', x)}. \quad (21)$$

We rearrange the true relation (17) to

$$\Delta \ln n' = \kappa + \eta \cdot \Delta \ln w' - \eta \cdot \frac{\phi}{\lambda} + \tilde{\xi}' = \kappa + \eta \cdot (const_1 + \gamma \cdot x + \nu) - \eta \cdot \frac{\phi}{\lambda} + \tilde{\xi}',$$

which uses the notation from above. Hence, the covariance between the dependent variable and the instrument is

$$\text{cov}(\Delta \ln n', x) = \eta\gamma \cdot \text{var}(x) - \eta \cdot \text{cov}(x, \phi/\lambda).$$

Expressing the instrument as $x = (\Delta \ln w' - const_1 - \nu)/\widehat{\gamma}$, we can state that

$$\text{cov}(x, \phi/\lambda) = \frac{1}{\widehat{\gamma}} \cdot \text{cov}(\Delta \ln w', \phi/\lambda) - \frac{1}{\widehat{\gamma}} \cdot \text{cov}(\nu, \phi/\lambda).$$

We can use the first-stage results to determine the covariance in the denominator of the estimate $\widehat{\eta}$ in (21) as

$$\text{cov}(\Delta \ln w', x) = \widehat{\gamma} \cdot \text{var}(x).$$

Inserting this into (21) gives

$$\widehat{\eta} = \frac{\eta\gamma \text{var}(x) - \eta \cdot \text{cov}(x, \phi/\lambda)}{\widehat{\gamma} \cdot \text{var}(x)} = \eta - \eta \cdot \frac{\text{cov}(\Delta \ln w', \phi/\lambda) - \text{cov}(\nu, \phi/\lambda)}{\widehat{\gamma}^2 \cdot \text{var}(x)}.$$

Hence, the percentage bias is

$$\frac{\widehat{\eta} - \eta}{\eta} = -\frac{1}{\widehat{\gamma}^2 \cdot \text{var}(x)} \cdot (\text{cov}(\Delta \ln w', \phi/\lambda) - \text{cov}(\nu, \phi/\lambda))$$

as stated in the proposition. ■

To simplify the notation in the following proofs, we define $Z_i = \exp z_i$ and $\Psi_i = \exp \psi_i$. Without loss of generality, we normalize $\Psi_1 = 1$, hence $\Psi_2 < 1$, and $\alpha_i = 1$.

Proof of Lemma 1. Consider a period with state variables a , $w_1 = Z_1$, and $w_2 = \Psi_2 \cdot Z_2$. Using the first-order condition for consumption (9), the multiplier ratio ϕ/λ is

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

Obviously, this ratio is zero if the household is not borrowing constrained in the current period, i.e., if $a' > 0 \Leftrightarrow \phi = 0$. Since z_2 and z_1 are i.i.d., a borrowing constrained household expects to enter the next period with state variables $a' = z'_1 = z'_2 = 0$. Hence, we can consider $E\lambda'$ in case of being borrowing constrained as a constant, which we denote by $\bar{\lambda}$.

Taken together, we can express the multiplier ratio as

$$\frac{\phi}{\lambda} = \max \left[1 - \beta \cdot (1+r) \cdot \frac{\bar{\lambda}}{\lambda}, 0 \right].$$

It is weakly increasing in the Lagrange multiplier on the current period's budget constraint $\tilde{\lambda}$, $\partial(\phi/\lambda)/\partial\tilde{\lambda} \leq 0$.

When the household is borrowing constrained, the Lagrange multiplier on the borrowing constraint can be determined from the remaining first-order conditions for the current period:

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

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

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

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 side of this expression as F and applying the implicit-function theorem gives

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

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

Proof of Lemma 2. Combining $\partial(\phi/\lambda)/\partial\Lambda \leq 0$ and $\partial\Lambda/\partial z_i > 0$ gives $\partial(\phi/\lambda)/\partial z_i \leq 0$ and $\text{cov}(\phi/\lambda, z_i) \leq 0$.

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

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

Generally, for a function f , it holds that $\text{cov}(f(x), y|a) = \text{E}(\partial f/\partial x|a) \cdot \text{cov}(x, y|a)$. Define $f(\Lambda) = \phi/\lambda$ and consider the covariance

$$\text{cov}(f(\Lambda), Z_i|a) = \text{E}(\partial(\phi/\lambda)/\partial\Lambda|a) \cdot \text{cov}(\Lambda, Z_i|a).$$

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

$$\text{cov}(\phi/\lambda, Z_1|a) \leq \text{cov}(\phi/\lambda, Z_2|a) \leq 0$$

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 $\text{cov}(\phi/\lambda, Z_i) = \text{E}(\phi/\lambda \cdot Z_i) - \text{E}(\phi/\lambda) \text{E}(Z_i) = \int \text{E}(\phi/\lambda \cdot Z_i|a) \, dh(a) - \int \text{E}(\phi/\lambda|a) \text{E}(Z_i|a) \, dh(a) = \int \text{cov}(\phi/\lambda, Z_i|a) \, dh(a)$, where h is the distribution function of a , such that

$$\text{cov}(\phi/\lambda, Z_1) < \text{cov}(\phi/\lambda, Z_2) < 0.$$

$Z_i \approx 1 + z_i$ and, hence, $\text{cov}(\phi/\lambda, Z_g) \approx \text{cov}(\phi/\lambda, z_g)$, completes the proof. ■

Proof of Proposition 2. Since the z_i are i.i.d.,

$$\text{cov}(\phi/\lambda, \Delta \ln w'_i) = \text{cov}(\phi/\lambda, z'_i - z_i) = \text{cov}(\phi/\lambda, z'_i) - \text{cov}(\phi/\lambda, z_i)$$

and

$$\text{cov}(\phi/\lambda, z'_i) = 0,$$

hence,

$$\text{cov}(\phi/\lambda, \Delta \ln w'_i) = -\text{cov}(\phi/\lambda, z_i).$$

Also,

$$\mathbb{E} \Delta \ln w'_i = \mathbb{E} \ln w'_i - \ln w_i = \mathbb{E} z'_i - z_i = -z_i.$$

Therefore,

$$\text{cov}(\phi/\lambda, \mathbb{E} \Delta \ln w'_i) = \text{cov}(\phi/\lambda, -z_i) = -\text{cov}(\phi/\lambda, z_i).$$

Together with Lemma 2, this completes the proof. ■

B Wage processes

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 $q(o_{it})$, a permanent component χ_i , and a transitory component \tilde{w}_{it} . The deterministic component filters out observable cross-sectional variation in wage rates. The permanent component captures individual-specific fixed characteristics. The transitory component captures temporary shocks affecting individuals' 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, see, e.g., Moffitt and Gottschalk (1995). 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:

$$\begin{aligned} w_{it}^* &= q(o_{it}) + \tilde{w}_{it}, \\ \tilde{w}_{it} &= p_t \cdot \chi_i + \zeta_t \cdot \tilde{z}_{it}, \\ \tilde{z}_{it} &= \rho \cdot \tilde{z}_{it-1} + \theta \cdot \varepsilon_{it-1} + \varepsilon_{it}, \end{aligned} \tag{22}$$

where $E(\chi_i) = E(\tilde{w}_{it}) = E(w_{it}) = 0$ and o_{it} is a vector of observable characteristics of individual i at time t . \tilde{w}_{it} is the stochastic component of individual's wages (residual wages). p_t and ζ_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 autocorrelation ρ and MA parameter θ . ε_{it} is a random variable with variance σ_ε^2 . Concerning initial conditions, we follow the approach by MaCurdy (1982) and treat the variance at the start of the sample period, $\sigma_{w_1}^2$, as an additional parameter to be estimated.

We estimate the wage process (22) separately for men and women using a two-step procedure. First, we estimate the deterministic component, $q(o_{it})$, 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 1 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 below using a Generalized Method of Moments (GMM) procedure.

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

$$\begin{aligned}\sigma_1^2 &= p_1^2 \sigma_\chi^2 + \zeta_1^2 \sigma_{w1}^2 \\ \sigma_t^2 &= p_t^2 \sigma_\chi^2 + \left\{ \zeta_t^2 \left(\rho^{2t-2} \sigma_{w1}^2 + K \sum_{j=0}^{t-2} \rho^{2j} \right) \right\}, t > 1\end{aligned}$$

and off-diagonal elements

$$\begin{aligned}\text{cov}(\tilde{w}_t, \tilde{w}_{t+s}) &= p_t p_{t+s} \sigma_\chi^2 + \zeta_t \zeta_{t+s} (\rho^s \sigma_{w1}^2 + \rho^{s-1} \theta \sigma_\varepsilon^2), t = 1 \text{ and } s > 0 \\ \text{cov}(\tilde{w}_t, \tilde{w}_{t+s}) &= p_t p_{t+s} \sigma_\chi^2 + \zeta_t \zeta_{t+s} \left(\rho^{2t+s-2} \sigma_{w1}^2 + \rho^2 K \sum_{j=0}^{t-1} \rho^{2j} + \rho^{s-1} \theta \sigma_\varepsilon^2 \right), t > 1 \text{ and } s > 0,\end{aligned}$$

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

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

$$\varphi = \{ \sigma_\chi^2, \rho, \sigma_{w1}^2, \sigma_\varepsilon^2, \theta, \zeta_2 \dots \zeta_T, p_2 \dots p_T \}.$$

For identification, the first-period factor loadings p_1, ζ_1 are set to one. In our unbalanced panel data set, each sample moment is constructed using all available observations covering the respective time span. Following Haider (2001), we adjust the standard errors of the parameter estimates to take into account the number of observations used in the computation for each moment. We follow Altonji and Segal (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 asymptotically optimal weighting matrix.

Table 9 summarizes the parameter estimates.³³ 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 (2006) and the references therein. For women, the estimated autocorrelation $\rho_f = 0.82$ is similar compared to men's, and the estimated degree of idiosyncratic labor market risk $\sigma_{f,\varepsilon} = 0.45$ is about twice as large as for men.

³³To 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 regressions are available on request.

Table 9: Estimated wage processes for men and women.

	(1)	(2)
	men	women
σ_x^2	0.06 (0.02)	0.05 (0.02)
ρ	0.84 (0.03)	0.82 (0.03)
σ_{w1}^2	0.13 (0.02)	0.23 (0.02)
σ_ε^2	0.05 (0.01)	0.21 (0.03)
θ	-0.53 (0.04)	-0.51 (0.04)
# moments	351	351

Notes: GMM estimation results for the covariance structure of residual wage rates. Residual wages obtained using a first-stage regression with time dummies, age dummies, and schooling dummies interacted up to a quadratic age trend. Standard errors in parentheses.

C Alternative model specifications

In this appendix, we present results from alternative model specifications. First, we consider the case of non-separable preferences. Second, we solve the model for private instead of public consumption. Third, we simulate our baseline model with identical wage dynamics for both genders. Fourth, we extend our baseline model by progressive taxation. All model versions are re-calibrated to match the targeted statistics as explained in Section 3.1.

C.1 Non-separable preferences

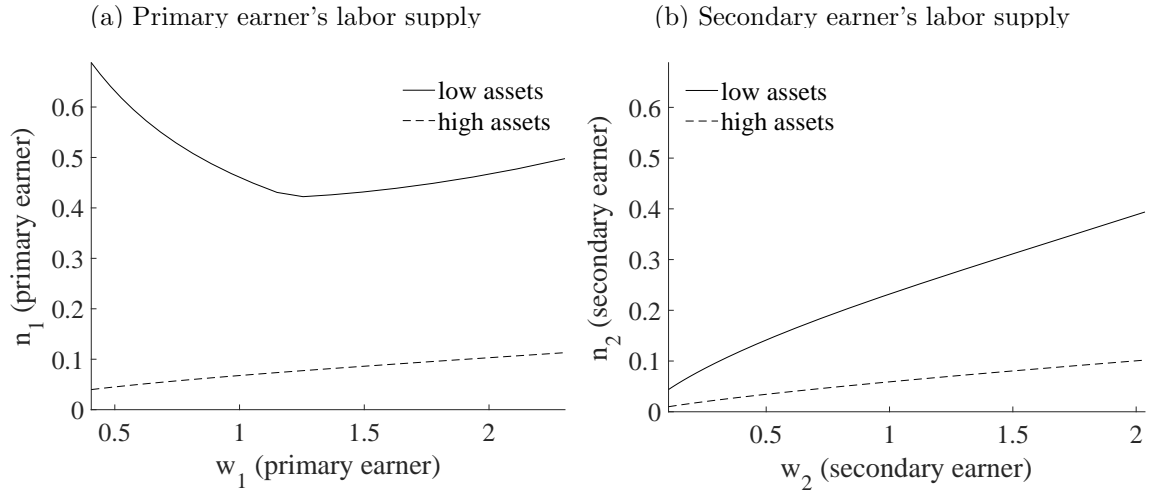
In a model with non-separable preferences, the true Frisch elasticities of labor supply are not determined by the curvatures of the disutility of working alone but also depend on household decisions. Thus, even at the individual level, the Frisch elasticity is not constant but varies over time. In our two-earner framework, non-separabilities in preferences might result from two potential sources. First, consumption and leisure may not be separable as individuals may enjoy consumption more in their leisure time. Second, in the context of a couple household, spouses may enjoy leisure more when they spend it together. To account for both dimensions, we follow Wu and Krueger (2015) and use the household target function

$$v = \frac{\left(\alpha \cdot c^\gamma + (1 - \alpha) \cdot \left(\zeta \cdot n_m^\theta + (1 - \zeta) \cdot n_f^\theta \right)^{-\gamma/\theta} \right)^{(1-\sigma)/\gamma} - 1}{1 - \sigma}, \quad (23)$$

where c , n_1 , and n_2 denote consumption, hours of the primary earner, and hours of the secondary earner, respectively. For simplicity, we directly specify a household target function (instead of individual preferences) as we have already discussed that a unitary model is sufficient for our purposes. We take the substitution elasticities γ and θ from Wu and Krueger (2015), $\sigma = 2.42$, $\gamma = -2.7$, and $\theta = 2.25$ which implies that there are complementarities between consumption and leisure as well as between the leisure times of the two spouses. We calibrate the preference weights α and ζ and the interest rate r to match the calibration targets described in the main text while all other parameters are left at their baseline values.

Policy functions. Figures 4 and 5 display the policy functions for hours worked and consumption for the case of non-separable utility. Importantly, the figures show that the mechanism discussed in the main text is also present under non-separable preferences. First, from panel (a) of Figure 4, it can be seen that labor supply of the primary earner is decreasing in his or her wage rate if the household is wealth poor and wage rates are not too high. This biases estimated Frisch elasticities downwards as discussed in the main text. Second, Figure 5 shows that, even in wealth-poor households, wage-rate fluctuations of the secondary earner

Figure 4: Labor supply as a function of own wage rate, non-separable preferences.




Notes: 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.

influence consumption only moderately while wage-rate fluctuations of the primary earner impact more strongly on consumption - as in Figure 3 in the main text, the consumption policy function has a kink. When the leisure times of both spouses are complements, it is particularly unattractive to raise the primary earner's hours when the secondary earner works relatively little. Nevertheless, wealth-poor households find it optimal to smooth negative wage shocks for the secondary earner relatively strongly through spousal labor supply. As a consequence, the estimation bias due to borrowing constraints will be relatively smaller for secondary earners, as in our baseline model with separable preferences.

Estimation. Columns (1) and (2) of Table 10 show the results of labor-supply regressions (19) for the model version with non-separable preferences. Again, we compare these estimates to the true Frisch labor-supply elasticities to determine estimation biases. Under the preferences (23), the Frisch elasticities of labor supply are not constant but depend on choices and, hence, differ even across individuals of the same gender. The true Frisch elasticity reported in the table corresponds to the average elasticity in the respective estimation sample.

Also under non-separable preferences, the estimation bias is more pronounced for primary earners than for secondary earners. Quantitatively, the effects are somewhat stronger under non-separable preferences. The estimation bias for primary earners is about 1.9 times as large as the one for secondary earners while, in our baseline model, this number is about 1.7.

Figure 5: Consumption policy functions, non-separable preferences.



Notes: 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.

Table 10: Estimation results from synthetic household panel data, alternative model specifications.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	non-separable pref.		private consumption		same wage dynamics		progr. taxes	
	primary earners	secondary earners	primary earners	secondary earners	primary earners	secondary earners	primary earners	secondary earners
expected log wage change	0.56 (0.01)	0.64 (0.01)	0.71 (0.01)	0.84 (0.01)	0.60 (0.01)	0.82 (0.02)	0.63 (0.01)	0.82 (0.01)
true Frisch bias	0.998 -44.2%	0.867 -26.3%	1 -29.2%	1 -15.7%	1 -39.8%	1 -18.0%	1 -37.1%	1 -17.8%
$\hat{\eta}_{sec}/\hat{\eta}_{prim}$	1.15		1.19		1.36		1.31	

Notes: Regression results for equation (19). Constant included but not shown. $E \Delta \ln w' = (\rho - 1) \cdot z$. Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise.

C.2 Private consumption

For the case where consumption is a private good within the household, the household budget constraint (5) is replaced by $c_1 + c_2 + \frac{a'}{1+r} \leq w_1 n_1 + w_2 n_2 + a$, where c_1 and c_2 denote consumption of the primary and secondary earner, respectively, and are choice variables. Further, instead of the first-order condition (8), it holds that $\frac{\partial v}{\partial c_1} = \mu \cdot c_1^{-\sigma} = \frac{\partial V(a, \omega)}{\partial a} = \lambda$ and $\frac{\partial v}{\partial c_2} = (1 - \mu) \cdot c_2^{-\sigma} = \frac{\partial V(a, \omega)}{\partial a} = \lambda$. Columns (3) and (4) of Table 10 show that the estimation results for this alternative model specification are virtually identical to the ones for the baseline model with private consumption.

C.3 Wage dynamics

Columns (5) and (6) of Table 10 present results for the case where we counterfactually simulate the model with identical wage dynamics for both genders. Specifically, we set the autocorrelation and standard deviation of the stochastic wage component for women, ρ_f and $\sigma_{\varepsilon, f}$, to the male parameter values. The estimation results for primary and secondary earners are similar to the ones in our full calibration with gender-specific wage dynamics, see columns (1) and (2) of Table 2.

C.4 Progressive taxation

In this subsection, we augment our baseline model by progressive taxation. We apply the parametric tax function used by Heathcote, Storesletten, and Violante (2014) and Blundell, Pistaferri, and Saporta-Eksten (2016). Specifically, the average tax rate is

$$1 - (1 - \xi) \cdot (w_1 n_1 + w_2 n_2)^{-\tau}, \quad (24)$$

such that the budget constraint of a household is $c + \frac{a'}{1+r} \leq (1 - \xi) \cdot (w_1 n_1 + w_2 n_2)^{1-\tau} + a$, instead of (5). Guner, Kaygusuz, and Ventura (2014) have estimated this tax function and we use their estimates for the "all married couples" sample to quantify the parameters in (24), $\xi = 0.1260$ and $\tau = 0.060$.³⁴

Columns (7) and (8) of Table 10 show the results of labor-supply regressions (19) using data from the model with progressive income taxation. As the estimation procedure ignores taxes (the right-hand side variable is, as before, expected growth in gross wages), the results are informative about the additional biases stemming from progressive taxation. The esti-

³⁴In Guner, Kaygusuz, and Ventura (2014), the counterpart to ξ is $0.087 \cdot \bar{y}^\tau$ where 0.087 is an estimate and \bar{y} is mean household income. In our model, this gives $\xi = 0.1260$. Guner, Kaygusuz, and Ventura (2014) use actual taxes paid (data from the Internal Revenue Service) rather than statutory tax rates and find that effective tax rates are substantially less progressive than statutory ones.

mates are somewhat smaller than those in Table 2 indicating that ignoring progressive taxes leads to an additional downward bias in the estimates in line with, e.g., Blundell, Pistaferri, and Saporta-Eksten (2016). However, we find that the estimates from the model with progressive taxes are quite similar to the ones obtained from our baseline model, especially when compared to the true Frisch elasticities, such that the *additional* biases due to progressive taxation are small compared to the biases stemming from borrowing constraints.

D Sample splits by asset holdings

Table 11 shows the results of the empirical labor-supply regressions when we apply sample splits by gender, earner status, and asset holdings. In the main text, we used a more parsimonious regression specification using interaction terms. The results presented here confirm our baseline results discussed in the main text, compare Table 6. We find that, in wealthier households, the estimated Frisch elasticities for primary earners become more similar to those of secondary earners. This corroborates that the estimation bias due to borrowing constraints is more important for primary earners.

Table 11: Empirical labor-supply regressions for men, PSID data (conditional on liquid assets).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	all men		assets > 0		assets > $0.5 \cdot y^m$		assets > y^m	
	primary	secondary	primary	secondary	primary	secondary	primary	secondary
	earners	earners	earners	earners	earners	earners	earners	earners
expected log	0.24	0.63	0.29	0.79	0.30	0.65	0.33	0.51
wage change	(0.14)	(0.35)	(0.16)	(0.41)	(0.17)	(0.42)	(0.18)	(0.44)
$\hat{\eta}_{sec}/\hat{\eta}_{prim}$	2.65		2.69		2.12		1.52	
time effects	yes	yes	yes	yes	yes	yes	yes	yes
observations	8,034	1,746	5,692	1,194	4,967	1,056	4,550	961

Notes: Regression results for equation (19). $E_t \Delta \ln w_{ijt+1}$ determined in a first-stage regression using individual characteristics. y^m is an average full-time male monthly income, $y^m = \$2014$ (1983 dollars). Standard errors in parentheses. Individuals identified as primary earners if the mean realized wage rate in the sample \bar{w}_i exceeds the mean realized wage rate of the spouse \bar{w}_{-i} and as secondary earners otherwise.