Essays in Quantitative Family Economics - Labor Supply, Risk Sharing and Welfare

Inauguraldissertation zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaft (doctor rerum oeconomicarum)
an der Fakultät Wirtschaftswissenschaften
- Schumpeter School of Business and Economics -
der Bergischen Universität Wuppertal

vorgelegt von
Jan Hendrik Gravert, M. Sc.
aus Dortmund

Wuppertal, im Januar 2020

The PhD thesis can be quoted as follows:

urn:nbn:de:hbz:468-20201119-092417-2
[http://nbn-resolving.de/urn/resolver.pl?urn=urn%3Anbn%3Ade%3Ahbz%3A468-20201119-092417-2]

DOI: 10.25926/kkpb-bp94
[https://doi.org/10.25926/kkpb-bp94]

Acknowledgments

I have written this thesis during my time as PhD student at the Chair of International and Regional Economics at the University of Wuppertal. First and foremost, I have to thank my supervisor Falko Jüßen. More than anyone else he has helped me with support, advice, patience and his subtle ways to motivate me during this time. He has created a perfect environment for me and my research undertakings. I am also especially grateful to Christian Bredemeier. He opened the wonderful field of family economics to me during my Bachelor and Master studies at the TU Dortmund University. It is fair to say that without his initial efforts, this thesis would have been never written.

Additionally, I want to thank many colleagues from the University of Wuppertal. They have given plenty of suggestions and comments during several Brown Bag Seminar presentations and discussions. I want to thank Hendrik Jürges, André Betzer, Werner Bönte, and Hans Frambach for their willingness to become members of my dissertation committee. Further, I thank my long-time roommates in office Dorothee Averkamp, Monika Bajon, and Torben Engelmeyer. All of them have motivated me in sometimes hard and stressful times.

During the writing process of this thesis I have participated in several summer schools and conferences. I want to thank the Eberhard-Robke-Fonds for financial support of these undertakings. With his foundation, Eberhard Robke has given me the chance to broaden my mind both in a cultural and in a research dimension. I am grateful for these opportunities.

Finally, I have to thank my whole family. All of them have supported and motivated me in one or another way. Nevertheless, I want to mention some family members a bit more prominent. First and foremost I have to thank my wife Jana. No one else has suffered more under my spleens and mood swings in the last years. She gives me the mental strength which was often needed to continue the work even when things seemed hopeless to me. Without her support and encouragement, this thesis might have never been finished. I also want to thank my wonderful son Hendrik. Although becoming a father during the PhD studies has increased the difficulty of the whole undertaking, one smile of him motivates me even in the most stressful situations. Finally, I have to thank three important persons, wise and sober-minded people. First, my two parents Dagmar and Helmut. Without their love and their support, I would not be the same person. They have proven to me that child care is not finished when the son moves out from home. Both are patient advisers for me, especially in terms of mental balance and motivation. Finally, I have to thank my mother-in-law Reinhild Althans. She gave me a new home ten years ago when I married her daughter. Within this time period she has become a second mother to me. Although she would say that she has not contributed to this thesis, her final contribution is much higher than she would ever expect. Without her support during the last years, this thesis would had not been possible.

Contents

Chapter I

Introduction

Since the pioneering work of Becker (1960, 1965), an increasing number of studies analyze the impact of family decisions on economic outcomes. Originally, the focus was on family-related topics such as marriage, fertility, or time use of spouses. Nowadays, also other fields of economics take the family dimension seriously, for example macroeconomics, see Doepke and Tertilt (2016) and Greenwood et al. (2017) for overviews; development economics, see for example Doepke and Tertilt (2009) and Fernández (2014); and financial economics, see Addoum et al. (2016) and Hertzberg (2019). My Thesis contributes to this growing literature by analyzing intra-household decision-making of couple households and the resulting outcomes, in particular, labor-supply decisions of couples, intra-household risk sharing, and inefficiencies caused by limited commitment between spouses. I follow the majority of the literature on intra-household decision-making by assuming that couples behave cooperatively. This is in line with the findings of Del Boca and Flinn (2012) that the majority of all couples, about 75%, behave cooperatively.[1]

The first generation of family-economics models have been unitary models of the household. Starting with Becker (1965), unitary models abstract from intra-household decision-making, i.e., individual preferences do not matter in the sense that the couple has joint preferences which are exogenously given. This implies that, in the unitary framework, individual preferences cannot be deduced from an empirical analysis. An important implication of the unitary model is *income pooling*, implying that the distribution of incomes among spouses does not matter for the allocation of goods within the household. Several empirical studies provide evidence against the unitary-model approach.[2]

The second generation of couple-household models are collective household models with full commitment. Full commitment means that, at the time of household formation, spouses choose a contingency plan for their future lives and can make binding promises to stick to this plan. Starting with the work of Chiappori (1988), collective models assume that individual preferences of spouses determine the behavior of couples.[3] Any Pareto efficient allocation within a couple can be derived from a weighted sum of individual preferences. The corresponding weights of the spouses are called Pareto weights. In collective models, Pareto weights are usually the result of a black-box intra-household decision-making process since any decision procedure which guarantees efficiency is compatible with the collective approach. In

[1] The literature on non-cooperative couple behavior is much smaller, see, e.g., Meier and Rainer (2015), and Doepke and Tertilt (2019) for recent examples.

[2] See for example Fortin and Lacroix (1997), Browning and Chiappori (1998), Anderson and Baland (2002).

[3] Some early studies are Manser and Brown (1980) and McElroy and Horney (1981) who analyze intra-household decision-making using Nash bargaining which is a special case of the collective approach of Chiappori (1988).

general, the Pareto weights are functions of prices, incomes, and distribution factors. Since Pareto weights depend on the distribution of incomes of spouses they have an impact on the allocation which is a contradiction to the *income pooling* hypothesis. Several empirical studies have provided evidence that the collective household model with full commitment characterizes well intra-household decision-making in static environments, see Browning et al. (2014) and Chiappori and Mazzocco (2017) for overviews. However, there is strong empirical evidence that commitment between spouses is limited over time, see, e.g., Mazzocco (2007), Cesarini et al. (2017), and Lise and Yamada (2019).[4]

Starting with the seminal paper by Mazzocco (2007), the third generation of couple-household models considers frameworks where commitment between spouses is limited. In these collective household models with limited commitment, the Pareto weights of the spouses become time-varying. Full and limited-commitment models have in common that spouses join the household only if their expected lifetime utility within the couple is at least as high as their corresponding expected lifetime utility outside the couple, i.e., their participation constraints must be fulfilled. A key difference between both model types is the frequency of relevant participation constraints. In the full commitment case, participation constraints must be fulfilled only at the stage of household formation. Under limited commitment, by contrast, participation constraints must be fulfilled in every period and in any state of the world because household members are free to leave the household at any time (unilateral divorce). Whenever one spouse has an incentive to leave the household, this spouse must be compensated by the other spouse such that the incentive is removed. This is achieved by adjustments in the Pareto weights which is the reason why these weights become time-varying in the case of limited commitment. This can be interpreted in at least two ways. The first interpretation stems from the literature on implicit contracts, see Marcet and Marimon (2019). Spouses engage in an implicit contract at marriage which must be self-enforcing since such contracts are not legally enforceable. This contract has to include all possible future scenarios. Then, when one spouse has an incentive to leave the household, the self-enforcing compensation mechanism applies which was bargained over in advance at marriage and which guarantees that both spouses remain within the couple. A second interpretation stems from the family-economics literature. Whenever one spouse has an incentive to leave the household, spouses re-bargain over the Pareto weights such that both spouses remain in the couple.

All three generations of couple-household models are used in the literature depending on

[4]In Chapter III, I discuss the limited commitment literature in detail.

the particular research question. In my Thesis, I consider all three types of household models. The main part of the Thesis consists of four chapters. The chapters are connected by contributing to two major topics. A first major topic is how the estimation of Frisch labor-supply elasticities is affected by intra-household decision making. In Chapter II, insights from family economics are used to overcome an estimation bias in labor-supply elasticities when couples are borrowing-constrained. In Chapter IV, it is shown that also limited commitment between spouses introduces a bias when estimating Frisch elasticities and how family economics can be used to address the estimation problem. The second major topic connecting the chapters of this Thesis is the analysis of economic outcomes of intra-household decision-making with limited commitment among spouses. In Chapter III, I investigate the impact of limited commitment on intra-household risk sharing. Chapter V focuses on savings behavior under limited commitment. A particular contribution of these chapters is to develop fast numerical methods to solve collective models with limited commitment and savings opportunities. Chapter VI presents a summary of the Thesis.

In **Chapter II** (joint work with Christian Bredemeier and Falko Jüßen) we use a (unitary) couple-household model to contribute to the literature on the micro/macro puzzle of labor-supply elasticities.[5] Usually, quantitative models in macroeconomics have to assume large values for the Frisch labor-supply elasticity to match empirical targets well. In contrast, microeconometric studies tend to find small estimates for Frisch elasticities. The literature discusses several biases which may help to explain this puzzle, see Keane and Rogerson (2015) for an overview. Using a Bachelor household perspective (and thus no family decision making), Domeij and Floden (2006) have provided an explanation that is related to the presence of borrowing constraints. They have shown shown that estimates of labor-supply elasticities are downward biased when individuals are borrowing constrained. The intuition is that individuals cannot smooth consumption perfectly through (dis-)saving/borrowing when they have low wage rate realizations compared to their mean wage rate and when they are additionally borrowing constrained. Individuals then have to raise their labor supply to smooth consumption because dis-saving is not possible. Thus, the reaction of labor supply to wage-rate changes is not solely determined by the Frisch labor-supply elasticity, but also by the willingness to borrow. Domeij and Floden (2006) show that the correlation between the willingness to borrow and wage-rate changes causes an omitted variable bias which biases the estimate of the labor-supply elasticity downward. The unique solution of Domeij and

[5]This chapter has been published in Bredemeier et al. (2019). My Thesis contains a slightly different version of this article.

Floden (2006) to overcome this bias is to estimate labor-supply elasticities only from households where asset holdings are sufficiently large. We show in Chapter II that an alternative approach to address the estimation bias can be derived from family economics. In particular, we show that exploiting the couple structure of most households can overcome the estimation bias due to borrowing constraints. We extend the model of Domeij and Floden (2006) to a (unitary) couple-household model. Using the couple structure, we derive analytically an unbiased estimator of the Frisch elasticity using the first-order conditions of the couple-household maximization problem. In particular, we show analytically that changes in hours growth between two periods depend linearly on an individual's contribution to household earnings. The lower the contribution of a spouse to total household labor income, the lower is the incentive of the household to borrow against future wage growth of this spouse after a negative wage rate shock. Exploiting this relation, we derive a regression approach where we include an interaction term between the average contribution to total household income and expected wage growth, the latter variable being the regressor that recovers the Frisch elasticity. We show analytically that our approach yields an unbiased estimate of the Frisch labor-supply elasticity. To investigate the quantitative performance of the estimation approach, we conduct Monte-Carlo experiments using simulated data from our calibrated model. Finally, we apply our estimation approach to household panel data from the Panel Study of Income Dynamics (PSID) for the U.S.. For men, we estimate a Frisch labor-supply elasticity of about 0.7 and about 1 for women. These estimates are larger than estimates reported in the majority of microeconometric studies. A general implication of our results is that differences in estimated Frisch elasticities of population groups which tend to have different earner roles in the household tend to be overestimated. A key advantage of our approach is that it does not require data on assets but only information on wage rates and hours worked of both spouses.

In **Chapter III**, I analyze risk-sharing implications of limited commitment between spouses. A large literature finds that intra-household risk sharing is an important source of insurance, see for example Ortigueira and Siassi (2013), Blundell et al. (2016), and Blundell et al. (2018). Risk-averse spouses wish to share income risk and smooth consumption as much as possible. Usually, studies on intra-household risk sharing assume that spouses can make binding promises to future intra-household allocation rule, i.e., spouses can fully commit to (future) allocation rules, implying perfect intra-household risk sharing. Thus, the ratio of marginal utilities of individual consumption of the spouses is fixed over time and states of nature. However, several studies provide empirical evidence that commitment between spouses is limited, see for example Mazzocco (2007), Cesarini et al. (2017), and Lise and Yamada (2019). Spouses wish to share income risk also when commitment between spouses is limited, i.e., spouses have the opportunity to leave the household. While it is initially rational for either spouse to promise to share positive income shocks with the partner in return for the promise of insurance against negative shocks, incentives to keep the promises are reduced once shocks have materialized. For example, after being promoted, an individual may not be willing to share the increased salary with the partner because the partner's promise of support in case the promotion had not taken place is now essentially worthless. Then, spouses re-bargain over the consumption share such that the incentive is vanished. Hence, risk sharing is incomplete when commitment is limited which increases the volatility of consumption for individuals and induces a welfare loss. After discussing the literature on limited commitment among spouses in detail, I contribute to the intra-household risk-sharing literature by developing and calibrating an incomplete-markets model with limited commitment in dual-earner households. In particular, I calibrate the model to be in line with an empirically reasonable amount of commitment issues using recent evidence on the extent of commitment issues by Lise and Yamada (2019). This is of particular importance since I want to make quantitative predictions of the impact of limited commitment. In my calibrated model, I find that variances of individual consumption increase by about 25% compared to a reference model with full commitment. The Kaplan and Violante (2010) insurance coefficient is reduced by about 36% due to limited commitment. This reduced ability to share risks cause a significant welfare loss of about 0.6% of consumption when commitment between spouses is limited. This welfare loss is in the range of the lower bound of the welfare costs of the business cycle in recent studies.[6] After the quantitative analysis of risk-sharing, I use my model to assess which parts of the population are particularly subject to incomplete intra-household risk sharing due to limited commitment. In the cross-section, I find that risk-sharing possibilities are particularly reduced for individuals whose patience or risk aversion is below average while individuals with strong gains from marriage or high degrees of altruism are less affected. Further, when the correlation in spouses' incomes is low or when couples face little income risk, individuals are affected rather strongly by limited commitment. Quantitatively, small differences in terms of direct utility gains from marriage and the degree of altruism around their respective mean values are predicted to lead to rather large differences in the extent to which intra-household risk sharing is worsened by limited commitment between spouses. These results are helpful for empirical studies as they help to identify subgroups

[6]Krebs (2007) estimates welfare costs of business cycles of about 0.5% and the number reported by Krusell et al. (2009) is over 1%

of the population which tend to have a higher or lower extent of commitment issues. My results have important policy implications as they show that individuals face substantially higher consumption risk than is suggested by a pure household perspective that implicitly assumes full commitment. This raises demand for other forms of insurance, including insurance through public programs such as unemployment insurance or progressive taxation. Second, there are also implications for the design of such programs. Limited commitment in the household implies that it should be the individual suffering from a negative income shock who is entitled to a compensation transfer payment, and not the household in which the individual lives. My results further show that there can be individuals suffering from low consumption even in households that in principle have sufficient means. Accordingly, my analysis suggest that means testing of transfer programs should at least in part be performed at the individual level and not at the household level.

In **Chapter IV** (joint work with Christian Bredemeier and Falko Jüßen), we consider how limited commitment between spouses affects estimates of Frisch labor-supply elasticities. We show that estimates are downward biased when commitment between spouses is limited and we develop new estimation approaches that reduce or fully eliminate this bias. The Frisch elasticity of labor supply can be estimated in a regression of hours worked on the hourly wage rate when controlling for consumption. Theory suggests to control for consumption of the individual worker but most household panel surveys contain consumption information only at the household level (if at all).[7] We show that proxying individual consumption by household consumption causes substantial biases in estimated Frisch elasticities when commitment between household members is limited. The labor-supply conditions comprise the individual Pareto weights of the spouses. Since one cannot observe these Pareto weights in the data, a potential omitted variable bias occurs. This is unproblematic when spouses can fully commit since individual fixed-effects capture Pareto weights as long as they are time-invariant. In contrast, when commitment between spouses is limited, Pareto weights may change over time such that a panel regression with household consumption and individual fixed effects is subject to a bias. We develop a limited-commitment model with dual earners and idiosyncratic risk and exploit first-order conditions to show analytically that estimates for the Frisch elasticity are biased downward under limited commitment. We provide a graphical illustration of the associated bias and perform Monte-Carlo experiments to quantify the bias using synthetic panel data from our calibrated model. We find that estimates of the Frisch labor-

[7] Until recently, the Panel Study of Income Dynamics (PSID) had incomplete consumption data and only information on food consumption was provided. Starting in 1999, the PSID was redesigned and more detailed information on household consumption has become available.

supply elasticity are up to 30% below its true value when household rather than individual consumption is controlled for and hence changes in individual bargaining positions are not taken into account properly. We develop two new estimation approaches that reduce or fully eliminate this bias. Both approaches exploit insights from family economics and explicitly exploit the couple structure of households. The key idea is to use information from *relative* labor-supply conditions of spouses. In the first approach, we use an approximation of first-order conditions to show that an improved estimation approach is obtained when including the wage rate and hours worked of the individual's spouse in the labor-supply regression. We show that this approach becomes more accurate when higher-order terms are included as additional controls. The approximated estimation approach including higher-order terms yields almost unbiased estimates. Our second approach is an iterative procedure. We deduce the Pareto weights using first-order conditions by guessing starting values for the parameters, including the Frisch labor-supply elasticities. Using these deduced Pareto weights, we update the guess using a labor-supply regression where we use the Pareto weights as additional controls. This procedure is iterated until convergence. Using Monte-Carlo experiments, we show that the procedure converges to the true parameter values but the iterative approach is more difficult to apply in practical applications than our approximated approach. We present an empirical application using PSID data where we use our preferred approximated estimation approach. This approach is recommended as it offers the best combination of close to unbiased estimates and empirical applicability.[8] As an intermediate step, we show that a regression using total household consumption leads to an underestimation of the Frisch elasticity as total household consumption is an imperfect proxy for individual consumption. Our preferred estimation approach yields a Frisch elasticity of 0.71. In line with the predictions of our model, this estimate is substantially larger than those obtained from labor-supply regressions that do not control for behavior of the partner and hence fail to account for the effects of limited commitment. Our estimate for the Frisch elasticity is also very similar to our own estimate that we obtain using a different estimation approach in first-differences where one has to control for a bias due to borrowing constraints, see Chapter II of this Thesis.

In **Chapter V** (joint work with Christian Bredemeier and Falko Jüßen), we analyze quantitatively how limited commitment between spouses affects savings decisions of couples. Ábrahám and Laczó (2018) show that couples save more when commitment between spouses is limited. Within their model, higher asset holdings make divorce more costly since spouses

[8] Specifically, the estimated model is a linear labor-supply regression that includes hours worked and the wage rate of the individual's partner as well as second-order terms in these variable and the individual's own wage as additional controls.

lose all assets upon divorce. Thus, the wealthier the household, the more an individual would lose upon divorce. In consequence, higher household savings relax future participation constraints. This behavior is inefficient from an ex-ante perspective since couples save inefficiently high which would not have been necessary if spouses could fully commit. Ábrahám and Laczó (2018) derive these results using restrictive assumptions. Incomes of spouses are assumed to be perfectly negatively correlated so that total household income is assumed to be constant, spouses lose all assets upon divorce, and individuals are not allowed to save in the outside option. The assumption of constant household income is useful for their purposes because it shuts down precautionary savings under full commitment and hence isolates savings due to limited commitment. The assumptions regarding asset losses and saving possibilities upon divorce facilitate the model solution. While Ábrahám and Laczó (2018) discuss that their main results are independent of these assumptions, there is still a lack of *quantitative* results on savings and risk sharing under limited commitment. Chapter V of my Thesis fills this gap. In particular, we compare the quantitative importance of the limited-commitment saving motive to the well-studied precautionary saving motive. Additionally, we analyze quantitatively the impact of property division rules upon divorce on economic outcomes, complementing the analysis by Ábrahám and Laczó (2016).[9] First, we develop a generalized version of the Ábrahám and Laczó (2018) set-up where we introduce general stochastic income processes, relax the assumption that all household assets are lost upon divorce, and allow divorcees to save. To perform our quantitative analysis, we develop a practically feasible and fast numerical approach which constitutes another contribution to the literature.[10] We show that limited commitment between spouses has a sizable effect on savings behavior of couples. Couples increase their yearly saving by about 4% which leads to an average rise in accumulated wealth of about 11%. Despite living in wealthier households, married individuals share risk less efficiently with their spouses when commitment between them is limited. Quantitatively, limited commitment reduces the insurance value of marriage by about 20% and the wellbeing of married individuals by an amount equivalent to 0.3% of lifetime consumption. About 40% of this welfare loss is due to over-saving Regarding the quantitative effects of property division rules, we find that asset losses upon divorce and penalties for the spouse who files for divorce can have substantial effects on decision making within the marriage. Both measures have the potential to reduce intra-household inefficiencies substantially. Large asset losses or penalties of about 70% can even raise consumption insurance within the marriage to the full-commitment level. However, the largest gains are obtained for moderate losses or penal-

[9] Ábrahám and Laczó (2016) is a follow-up paper to Ábrahám and Laczó (2018).

[10] Under limited commitment, the additional participation constraints complicate the numerical solution.

ties. Losses or penalties in the range of 20% already close half of the gap in consumption insurance due to limited commitment. By contrast, uneven property division rules that do not depend on events within the marriage or on who files for divorce affect efficiency within the marriage only to a small degree.

Chapter II

Estimating Labor-Supply Elasticities with Joint Borrowing Constraints of Couples

1 Introduction

The Frisch elasticity of labor supply measures the percentage reaction of hours worked to a one percent change in the net wage rate holding the marginal utility of wealth constant.[1] Thus, the Frisch elasticity determines adjustments in labor supply to wage-rate changes that trigger pure intertemporal substitution effects but no accompanying income effects. There are various examples for wage-rate changes that have this property. First, under perfect capital markets, purely transitory wage-rate changes should have no impact on the marginal utility of wealth. Second, if agents are forward-looking and can borrow freely, wage-rate changes that can be expected by the agent in advance should also leave the marginal utility of wealth unchanged. Accordingly, the Frisch elasticity is important for reactions to transitory tax or productivity shocks and to predictable life-cycle patterns in wage rates.[2]

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

A particular estimation problem has been highlighted by Domeij and Floden (2006), who have shown that, in presence of borrowing constraints, conventional methods to estimate the Frisch elasticity are subject to a downward bias. This bias is important since borrowing constraints are a substantial restriction to many households in the U.S. (see, e.g., Díaz-Jiménez et al. 2011). In this paper, we derive a new estimation approach for the Frisch elasticity that yields unbiased estimates even in samples of potentially borrowing-constrained households. Our approach critically exploits the couple structure of households, i.e., we

[1] A slightly different version of this chapter is published in Bredemeier et al. (2019). 0734-306X/2019/3704-0006$10.00

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

exploit information from households with two potential earners.[3] Our approach is appealing to the applied researcher as it takes the form of a simple interaction-term regression with minimum data requirements. When we apply our method to household data from the Panel Study of Income Dynamics (PSID), we estimate relatively large values for the Frisch elasticity in comparison to previous studies.

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

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

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

from the estimation sample— the Frisch elasticity cannot be estimated correctly in an Altonji (1986) set-up. Yet, from a practical point of view, reliable household panel data on assets are hardly available, and even if they are, such data are often not observed in the same panel as labor earnings and working time.

We contribute by extending the analysis of Domeij and Floden (2006) to a two-person household set-up and by deriving an unbiased estimator of the Frisch elasticity. Our approach critically exploits the couple structure of the model and the data but does not require information on household assets. Intuitively, in a double-earner household, also the partner can react to one's own wage-rate shocks, i.e., also the partner's labor supply can be used to smooth consumption. This is particularly important if the partner earns relatively much. Then, a given negative wage-rate shock can be smoothed relatively easily as the partner's hours have to be raised by only relatively little. Importantly, this relation holds even when the household is borrowing constrained. And, when it is predominantly the partner's labor supply that smooths consumption, one's own hours change is again mostly (in the limit, only) determined by the Frisch elasticity. Accordingly, to derive an unbiased estimator of the Frisch elasticity in presence of borrowing constraints, we can exploit the relation that the household's desire to borrow against wage growth is the less important the less an individual contributes to total household earnings.[4]

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

We then evaluate our estimator in Monte Carlo experiments using a calibrated incomplete-markets model populated by double-earner households, and finally we use the method for es-

[4]Blundell et al. (2016) provide direct empirical evidence that household consumption reacts more strongly to husbands' wage shocks than to wives' wage shocks which is line with our model since husbands on average contribute larger shares to household earnings. In Guner, Kaygusz, and Ventura (2011, 2012), Domeij and Klein (2013), and Bick (2016), similar mechanisms to ours affect labor-supply reactions to permanent wage-rate changes. In these studies, income effects are weaker for women (who are often secondary earners) such that their reactions to permanent wage-rate changes mostly reflect substitution effects governed by the Frisch elasticity.

timations using PSID data. Importantly, our approach critically exploits the couple structure of our model and the data, the key issue being that, only in a population of double-earner households, there is variation in individuals' percentage contribution to household earnings which we use to identify the Frisch elasticity.[5] Our estimations using PSID data suggest Frisch elasticities for men of about 0.7. We also take into account modifications of our interaction-term approach to cope with challenges that arise when estimating Frisch elasticities for women. For women, we find Frisch elasticities of around one.

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

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

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

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

The model is a partial-equilibrium incomplete-markets model with two household members. Households differ from one another by asset holdings and wage rates. Members of a household are subject to joint budget and borrowing constraints and take decisions cooperatively under full commitment, so that the resulting allocations are Pareto optimal. Households are potentially borrowing constrained and use precautionary savings in a non-state contingent asset and labor supply of both household members to insure against bad wage-rate realizations. This behavior is similar as in the model of Ortigueira and Siassi (2013) and extends the model of Domeij and Floden (2006) to a two-person setup.

2.1 Decision problem

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

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

subject to the household budget constraint

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

and the borrowing constraint

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

where u_i is the instantaneous utility function, c is household consumption, n_i is hours worked by household member i, β is the rate of time preference, E is the expectation operator, a denotes the household's asset holdings, ω is the vector of wage rates of both household members, $\omega = (w_1, w_2)$, and r is the exogenous interest rate.[1] A prime ($'$) denotes next period values.

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

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

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

where σ denotes risk aversion and α_i is the taste for leisure. This utility function has the property that the true value of the Frisch elasticity is given by the curvature parameter η_i. In Appendix A.6.1, we consider an alternative model specification with non-separable preferences which yields similar results as our baseline model. The preference parameters are indexed by i as, in our quantitative evaluations, we will account for potential differences in these parameters between household members.

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

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

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

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

In our baseline model, we assume that consumption is a household public good, i.e., there is no consumption rivalry between spouses. Along with additive separability, the public good assumption allows a simple notion of Frisch elasticities in a context with two earners. Specifically, the issue of whose marginal utility of wealth is held constant (husband's, wife's, or household's) does not arise, since, if one of them is constant, the other two are constant as well, independent of bargaining. For completeness, we also considered a model version with private instead of public consumption. In this version, we obtain almost identical results, see Appendix A.6.2. Even allowing for endogenous time-varying Pareto weights in the spirit of a limited-commitment model (see, e.g., Ligon et al. 2002) would have no substantial impact on our results since the weights would mostly react to unexpected changes in wage rates while the Frisch elasticity is identified through changes in expected wage rates. In our baseline model, we further abstract from non-linear taxation. This assumption allows to recover the true Frisch elasticity consistently in absence of borrowing constraints. We also consider a model version with progressive joint taxation of spouses in Appendix A.6.3. Also in this version, we obtain similar results as in our baseline economy.

In a further model extension, we follow the literature (Guner et al. 2011, 2012, Bick 2016) and take into account the possibility that labor supply of women is also affected by fluctuations in the disutility of work originating from taste-for-work shocks, e.g., capturing shocks to home production or child care, see Section 6.2 for details. This model extension delivers important insights for our empirical investigation of labor-supply elasticities for women.

2.2 Equilibrium conditions

The first-order conditions of the household problem are

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

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

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

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

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

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

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

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

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

We now derive a procedure for obtaining an unbiased estimate of the Frisch elasticity in presence of borrowing constraints. We first derive our approach analytically and then evaluate it numerically using Monte Carlo experiments. To derive the estimator analytically, we apply a simplifying assumption on data frequency, which will be relaxed in the Monte Carlo experiments where we will consider a realistic, i.e., annual, data frequency. Specifically, to obtain closed-form solutions, we assume an arbitrarily small period length. Due to this assumption, it is sufficient to consider the group of borrowing-constrained households and the group of unconstrained households and, in first differences, one can neglect households that move from one group to the other.[1] For both groups, we can derive the relation between hours changes and expected wage growth analytically, and then we can pool both groups to derive the population estimate. In Appendix A.2, we derive analytical results which are independent of the period length. While the derivations are more cumbersome, the main results presented here under the assumption of an arbitrarily small period length carry over to the more general case. For simplicity, we assume in the analytical part that spouses' Frisch elasticities are identical, $\eta_1 = \eta_2 = \eta$. In the quantitative model analysis in Section 4, we account for potential heterogeneity in the true Frisch elasticity to capture gender differences. Further, we assume in the analytical part that the process for stochastic wage-rate components is homogeneous across the population. In the quantitative model analysis in Section 4, we take into account gender differences in these processes.

3.1 Households unaffected by borrowing constraints

For households unaffected by borrowing constraints, a regression of hours growth on expected wage growth yields an unbiased estimate of the Frisch elasticity. For bachelor households, this has been shown in the seminal paper by Altonji (1986). Our case of a double-earner household is a straightforward extension. In Appendix A.1.1, we show that, after taking logs and first differences, the Frisch elasticity can be recovered through regressions of the form

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

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

for household members $i = 1, 2$, where $\xi' = \ln \lambda' - \mathrm{E} \ln \lambda'$ is an expectation error which results from using the Euler equation to substitute marginal utility of consumption from the labor-supply conditions. The terms ω_i', $i = 1, 2$, are unexpected components of wage growth which result from a decomposition of observed wage growth in an expected and unexpected component. As shown by Altonji (1986), the combined residual $\eta \cdot (\xi' - \omega_i')$ is uncorrelated with the regressor expected wage growth, see Appendix A.1.1 for an intuitive explanation. The terms $\eta \cdot \ln(1 + r)$ and $\eta \cdot \ln \beta$, can be captured by time fixed effects and a constant, respectively. Thus, when borrowing constraints are not binding, a simple regression of hours growth on expected wage growth ("Altonji (1986) regression") identifies the Frisch elasticity.

3.2 Borrowing-constrained households

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

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

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

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

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

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

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

see Appendix A.1.2 for a derivation.[2] We measure the earnings contribution in the point of approximation $\overline{s}_i$ by the individual's average contribution to household earnings during the sample period. Put differently, the point of approximation is the situation where both spouses contribute their usual shares to household income.[3]

[2]Equation (16) is a first-order approximation of the budget constraint (in logs). In Appendix A.4, we evaluate the importance of the approximation for our results and find that the approximation has a negligible effect.

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

In Appendix A.1.2, we solve this system to obtain

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

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

Equation (17) separates the two effects of expected wage growth on hours growth in a borrowing-constrained household. First, as in an unconstrained household, expected wage growth induces the wish to substitute labor into periods where it is paid more. This intertemporal-substitution effect is governed by the Frisch elasticity, η. Second, expected wage growth induces the willingness to borrow against expected future earnings in order to smooth consumption. However, a borrowing-constrained household can only smooth consumption by supplying more labor which counteracts the intertemporal-substitution effect. The strength of this willingness-to-borrow effect depends on the individual's usual contribution to household earnings $\overline{s}_i$. Expected wage growth for individuals with low earnings contributions induces only relatively small expected changes in total household earnings and can more easily be smoothed through labor-supply adjustments of the partner. Hence, for these individuals, the willingness-to-borrow effect is weak and hours growth is mostly determined by intertemporal substitution and, thus, the Frisch elasticity.[4]

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

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

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

[4]A similar mechanism applies to permanent changes in wage rates for which borrowing constraints are less important but classical income effects play an important role. For individuals who contribute little to household earnings, changes in hourly wage rates induce a small change in household earnings which mutes the income effect of changes in wage rates. As a consequence, the labor-supply response to these changes is mostly driven by substitution effects and, hence, tends to be stronger than for individuals who contribute larger shares to household earnings or who are the sole earners in their households. This mechanism can help to understand findings reported by Guner, Kaygusz, and Ventura (2011, 2012), Domeij and Klein (2013), and Bick (2016) who all document that, in quantitative macro models with double-earner households, labor supply is particularly responsive for groups that can be expected to contribute small shares to household earnings.

3.3 Mixed population

We now consider a sample that includes individuals from both borrowing-constrained and unconstrained households. We denote the sample shares of the constrained and unconstrained households by p and $1-p$, respectively. As an intermediate step, we consider a group of individuals with usual earnings contribution $\overline{s}$ but which include both, unconstrained and constrained households. In such a sample, a standard Altonji (1986) regression of hours growth on expected wage growth yields the following estimate for the Frisch elasticity:

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

which uses that $\mathrm{E}\left(\Delta\ln n'\right)\cdot\mathrm{E}\left(\mathrm{E}\,\Delta\ln w'\right)=0$.[5] The final step is to consider a sample where individuals differ in their usual contributions to household earnings. In such a sample, the OLS estimate averages over the different $\overline{s}$ such that the coefficient on expected wage growth is

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

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

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

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

3.4 Deriving an unbiased estimator

The relation between the earnings contribution and the covariance term $\operatorname{cov}(\Delta \ln n', \mathrm{E}\,\Delta \ln w')$ in (19) holds two key insights for deriving a regression specification that yields an unbiased estimate of the Frisch elasticity, even in samples that include borrowing-constrained households. First, in a population of double-earner households, there is variation in individuals' contribution to household earnings which can be used to identify the Frisch elasticity. Second, the covariance $\operatorname{cov}(\Delta \ln n', \mathrm{E}\,\Delta \ln w')$ is a *linear* function of $\overline{s}$, see (18).

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

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

gives

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

and

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

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

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

Figure 1: Policy functions for labor supply.

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

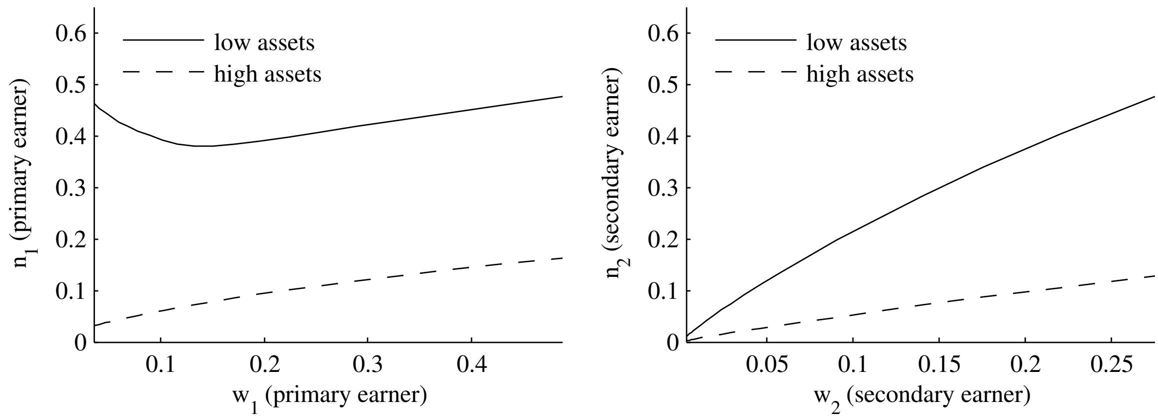

NOTE.–Policy functions refer to household type X, which is a household type with pronounced long-run intra-household wage differences. In the left panel, the wage rate of the secondary earner is at its lowest possible grid value. In the right panel, the wage rate of the primary earner is at its highest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

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

3.5 Graphical illustration

Figure 1 shows policy functions from a numerical solution of our calibrated full model.[6] For the graphs, we compare the household member who contributes, in the long run, more to household earnings (the primary earner) and the member who contributes less (the secondary earner). For illustration, we consider a household with strong wage-rate differences between household members such that the willingness-to-borrow effect is strong for the primary earner and weak for the secondary earner.

The labor-supply curve of the primary earner (left panel) is globally upward-sloping if the household is wealth-rich (dashed line), reflecting the intertemporal-substitution effect governed by the Frisch elasticity, due to the household's ability to smooth consumption through dis-saving when wage rates are low. The standard Altonji (1986) regression identifies the Frisch elasticity from this upward-sloping shape of the labor-supply curve. By contrast, for a household with low asset holdings (solid lines), the borrowing constraint is binding when

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

wage rates are low. Then, the labor-supply curve of the primary earner has a downward-sloping range where a further wage decrease triggers an increase in labor supply (rather than a decrease), because consumption cannot be smoothed through borrowing but only through an increase in labor supply. This negative relation between wage-rate changes and labor-supply changes leads to a downward estimation bias.

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

4 Estimating labor-supply elasticities from synthetic data

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

4.1 Calibration

Our baseline PSID sample used for the calibration covers the period 1972-1997, see Appendix A.3.1 for details on the sample selection. Due to our focus on double-earner households, we consider household heads and their partners for whom both partners' wage rates are observed. Further, we apply similar sample selection criteria as Altonji (1986) and Domeij and Floden (2006). In particular, we consider individuals between age 25 and 60.

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

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

We estimate the parameters of the stochastic wage processes, i.e., autocorrelations ρ_m, ρ_f and innovation variances $\sigma^2_{m,\varepsilon}$, $\sigma^2_{f,\varepsilon}$, separately for men (m) and women (f).[1] We first obtain residual wages by filtering deterministic cross-sectional variation using an OLS regression. We then identify autocorrelations and innovation variances from gender-specific Generalized Method of Moments (GMM) estimations, see Appendix A.3.2 for details.[2]

Our interaction-term approach that corrects for the bias due to borrowing constraints exploits variation in individuals' usual percentage contribution to household earnings. In order to assess our method in the model, the simulated economy has to feature sufficient

[1] Blundell et al. (2016) document that alternatively using a combination of permanent and transitory shocks leads to similar estimation results for preference parameters such as Frisch elasticities.

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

and realistic variation in individuals' contribution to household earnings. We therefore solve and simulate our model with ten household types. Household types differ in the constant (=permanent) wage components ψ_i of its members which we set to match average male and female wage rates in the ten deciles of the empirical distribution of relative wage rates of spouses in couple households in our PSID sample. We then calibrate household-type specific preference weights α_m and α_f to match average hours worked by gender and group, and, as a result, our calibrated model displays a realistic distribution of relative labor earnings within households.[3]

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

For the remaining preference parameters we use standard values from the literature. Relative risk aversion is set to $\sigma = 1.5$.[4] Following Domeij and Floden (2006), we set $\beta = 0.95$ (annual model frequency), and calibrate the interest rate so that the bottom 40% of the wealth distribution own 1.4% of total wealth. Table 7 in Appendix A.3.3 summarizes all parameter values of our baseline model.

4.2 Simulation set-up

We simulate a synthetic panel data set with similar size as our baseline PSID sample. Specifically, we simulate households for a long period of time and calculate hours growth, expected wage growth, average wage rates, and average contributions to household earnings. We then draw 10,000 samples of 15,000 household-year observations which we use for the regressions and report mean point estimates and mean standard deviations. In the estimations, we consider separate samples of men and women to take into account gender differences in both, true Frisch elasticities and usual earner roles. In the main text, we report the estimation

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

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

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

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

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

results for men while results for women are similar and can be found in Appendix A.5. To determine the regressor expected wage growth, we exploit the properties of the wage process (21), i.e., we calculate $\mathrm{E}_t \, \Delta \ln w_{ijt+1} = (\rho_i - 1) \cdot z_{ijt}$.

4.3 Monte Carlo results

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

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

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

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

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

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

We have considered several extensions of our baseline model and have investigated the performance of our interaction-term approach in these extended model environments. Specifically, we have incorporated non-separable preferences, progressive income taxes, and private instead of public consumption, see Appendices A.6.1-A.6.3 for details. In all model extensions, we find that our preferred interaction-term approach delivers an estimate of the Frisch elasticity which is close to its true value while standard Altonji (1986) regressions underestimate it considerably. Finally, we have developed modifications of our interaction-term approach to cope with challenges when estimating the Frisch elasticity for women instead of men. In Appendix A.6.4, we show that our modified approaches are robust in presence of taste-for-work shocks originating from, e.g., child care or home production, in particular our approach using predicted instead of actually observed earnings contributions.

5 Estimating labor-supply elasticities from PSID data

In this section, we present empirical results for our interaction-term approach using PSID data. Details on the data and sample selection can be found in Appendix A.3.1. As shown above, our approach corrects for the bias due to borrowing constraints and is able to deliver an almost unbiased estimate of the Frisch elasticity. The outline of the empirical analysis closely follows our Monte Carlo experiments, i.e., we first investigate two key implications of our model and then apply our preferred interaction-term approach to the data. Before presenting estimation results, we discuss some econometric aspects that are relevant in an empirical application of our approach.

5.1 Econometric aspects

As in our theoretical model, we analyze the choice of hours worked at the intensive margin in double-earner households.[1] When estimating labor-supply regressions from PSID data, we use individual characteristics to determine expected future wage changes, $\mathrm{E}_t\,\Delta w_{ijt+1}$, in gender-specific OLS regressions. Specifically, we follow, e.g., MaCurdy (1981) and Domeij and Floden (2006) and use as predictors age, age squared, years of schooling, and an interaction term between age and years of schooling. If there were no borrowing constraints, predictable wage growth would leave the marginal utility of wealth unchanged and would hence identify the Frisch elasticity. Using individual characteristics as predictors has the advantage that measurement error in these variables is uncorrelated with measurement error in wage rates.[2]

In empirical data, individual labor supply may also be affected by taste shifters.[3] Using first-differenced data is helpful in addressing this aspect. First-differencing eliminates the

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

[2] Measurement error in hours on the left-hand side of the regression reduces the R^2 of the regression but the estimate for η is consistent. As discussed by Altonji (1986), Domeij and Floden (2006) and Keane (2011), the instruments used to determine expected wage growth are potentially weak, one reason being that most wage changes may simply be unexpected. A potentially strong instrument is the lagged wage rate but this instrument should be avoided because it magnifies biases stemming from measurement error in the wage data (see Altonji, 1986, for further discussion). Using higher lags of the wage rate would mitigate this problem, but, in our sample, such instruments are barely informative for future wage changes. In our gender-specific first-stage regressions to obtain expected wage growth, the F statistics are 18.59 for men and 11.69 for women. In the IV literature (see Stock et al. 2002), instruments are regarded as reliable if, in the case of one endogenous regressor, the F statistic exceeds 10.

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

need to control for permanent taste shifters which are likely correlated with wages, such as education. In turn, this means that only transitory taste shifters may still be present in the first-differenced regression. As argued by, e.g., Keane (2011), transitory taste shifters are less likely to be correlated with expected wage changes.

Empirical estimates of labor-supply elasticities can also be affected by non-linear taxation (e.g., Aaronson and French 2009). When taxes are progressive, changes in gross wage rates overstate changes in net wage rates. Even if marginal net wages were observable, they would be endogenous as changes in hours affect marginal tax rates under progressive income taxation. In Appendix A.6.3, we present Monte-Carlo estimations for a model version with progressive income taxation which show that, in our context, the biases due to progressive taxation are small compared to the biases arising from borrowing constraints. Relatedly, it may be argued that taxes will largely drop out of a labor-supply condition in yearly differences as the household's marginal tax rate typically does not change substantially from year to year, see, e.g., Altonji (1986).

5.2 Testing two key implications of our model

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

Our theory predicts a negative coefficient for the interaction term. This is confirmed in our estimations using PSID data, see Table 2. For both, men and women, the incremental effect of being a primary earner on the estimated Frisch elasticity is significantly negative. This corroborates that labor-supply elasticities are estimated to be substantially smaller for primary than for secondary earners when borrowing constraints are ignored. Considering gender-specific regressions is important to make this point as they show that the estimated differences in labor-supply elasticities of primary and secondary earners are indeed related to differences in earner status and do not primarily pick up gender differences in the true Frisch elasticities. In Appendix A.7, we present additional evaluations corroborating this point. Specifically, we consider alternative definitions of primary and secondary earners and we compare Altonji (1986) regressions in several ranges of the relative contribution to household

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

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

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

earnings.

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

The second testable implication of our analysis is that differences in estimated Frisch elasticities from Altonji (1986) regressions should become smaller when the samples are less affected by borrowing constraints. As in our Monte Carlo experiments, we test this prediction by comparing male primary and secondary earners in samples of households with different liquid wealth. In particular, we repeat the estimations including an interaction term with the primary-earner dummy but only consider households with liquid wealth above a certain threshold which we increase step by step. For this evaluation, we build on Domeij and Floden (2006) and restrict the PSID data to three 3-year panels for which detailed asset data are available. Our theoretical model predicts the coefficient on the interaction term to decrease with an increasing wealth cut-off. We find that this pattern is confirmed in the PSID data, see Appendix A.7.3 for details.

5.3 Frisch-elasticity estimates for men

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

We estimate a substantially larger Frisch elasticity when we use our preferred interaction-term approach that exploits the couple structure. This is in line with our theoretical analysis where we have shown that our approach corrects for the negative bias due to borrowing constraints. Our interaction-term approach using the husband's average percentage earnings contribution yields an estimated Frisch elasticity (the coefficient on non-interacted expected wage growth) of 0.72, see column (2). Compared to the estimation without the interaction term in column (1), the bias-corrected estimate is hence about three quarters higher.[4] Comparing the estimation results in columns (1) and (2) suggests that the bias due to borrowing constraints in Altonji (1986) regressions amounts to more than 40% for men which is quantitatively in line with our Monte-Carlo experiments. The negative coefficient on the interaction term in column (2) reflects that men with higher earnings contributions have a weaker connection between expected wage growth and hours growth. This corroborates that their labor supply is particularly strongly exposed to the effects of borrowing constraints which induce a negative co-movement of expected wage growth and hours growth counteracting the positive co-movement induced by intertemporal substitution and governed by the Frisch elasticity.

As in our Monte Carlo experiments, we also investigate in how far our estimates are driven by male secondary earners in the sample. As discussed before, these individuals may be particular in various aspects and one may be skeptical when identification would largely depend on these individuals. In order to address this concern, we re-estimate our interaction-term specification for a restricted sample, where we only include men who are primary earners, see column (3) of Table 3. We find that the estimate for the primary-earner only sample is similar to the one obtained for the full sample, in line with our results from the Monte Carlo analysis. This corroborates that male secondary earners are not solely responsible for identification although we use the covariance between expected wage growth and hours growth if the male earnings contribution *were* small.

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

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

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

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

6 Implications for labor-supply elasticities of different population groups

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

6.1 Primary and secondary earners

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

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

6.2 Men and women

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

For example, we could measure low contributions to household earnings for women who work only few hours and have a relatively elastic labor supply because of child-care obligations, as shown by Alesina et al. (2011). Then, our derived estimator would put relatively much weight on a group of women whose labor-supply elasticity is not representative for the total population. We therefore develop modifications of our baseline interaction-term approach to address challenges when estimating the Frisch elasticity for women.[1]

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

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

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

The shock h can be understood as a home production requirement, e.g., the presence of children without the availability of informal or affordable formal child care.[2] Similar to Guner, Kaygusz, and Ventura (2011, 2012), we model h as a two-state Markov process with states h_{low} and h_{high} and transition probabilities κ_1 from h_{low} to h_{high} and κ_2 from h_{high} to h_{low}.[3]

In Appendix A.6.4, we present a detailed analysis of the model extension with preference shocks and we show that our baseline interaction-term approach tends to over-estimate the true Frisch elasticity in this setting. We therefore suggest modifications of our baseline approach and we show in Monte-Carlo experiments that, with these modifications, we obtain almost unbiased estimates also in the model with preference shocks. First, we adopt an approach where we consider a sample restriction and only consider women who contribute at least 30% to household earnings. Second, we apply an approach where we replace the wife's actual earnings contribution by the predicted earnings contribution based on a regression with observable determinants as regressors.[4] This alternative measure of the earnings contribution is less affected by idiosyncratic determinants (such as child care needs).

When we apply these two approaches to the PSID data, we obtain an estimated Frisch elasticity for women of around one, see columns (2) and (3) of Table 4. These estimates for women are only about 45% larger than the one for men. By contrast, comparing gender-specific labor-supply elasticities on the basis of Altonji (1986) regressions, i.e., comparing the results in column (1) of Table 4 and column (1) of Table 3, would suggest considerably larger gender differences of about 90%.[5] Also the calibration of our theoretical model suggests rather small gender differences in labor-supply elasticities. In fact, a difference of only about 40% is needed to rationalize the substantially larger difference in Altonji (1986) estimates. In summary, our analysis for women suggests that potential gender differences in the true elasticities are magnified by the differential importance of the estimation bias so that differences in true elasticities are in fact smaller than suggested by previous studies. This way, our analysis has implications for, e.g., the taxation of couples (Kleven et al. 2009) or gen-

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

[3] We follow Guner et al. (2011) and Bick (2016) to calibrate these additional parameters, see Appendix A.6.4 for details.

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

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

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

	(1)	(2)	(3)
expected	0.78	1.08	1.05
wage growth	(0.17)	(0.27)	(0.23)
expected wage growth		-1.49	
$\times$ earnings contribution (%)		(0.46)	
expected wage growth			-0.77
$\times$ predicted contribution			(0.43)
time effects	yes	yes	yes
sample	all	$\bar{s}_{ij} \geq 0.3$	all
observations	14,340	8,966	14,340

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

ders (Alesina et al. 2011) where arguments often rely on gender differences in labor-supply elasticities.

6.3 High and low earnings

Our analysis also implies that conventional methods overestimate the differences in labor-supply elasticities between groups with high and low *levels* of earnings.[6] When we distinguish between men in the upper 25% of the distribution of current labor earnings and those in the bottom 75%, see columns (1) and (3) of Table 5, estimates from standard Altonji (1986) regressions suggest that the labor supply of individuals with high earnings is considerably less elastic. Accordingly, one might draw the conclusion that strong tax progressivity is efficient, see, e.g., Saez (2001) who relate optimal income tax rates to labor-supply elasticities.[7] However, our analysis suggests that this difference in labor-supply elasticities is over-estimated as individuals with high earnings on average also contribute larger *shares* to household earn-

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

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

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

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

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

ings. In fact, when we apply our preferred interaction-term approach, estimated labor-supply elasticities are found to be more similar for both earnings groups, see columns (2) and (4) of Table 5. Remaining differences in estimated elasticities may reflect, among other things, that the top-earnings group is on average older and more educated than the rest of the population. We also considered a specification where we exclude individuals who switch back and forth between the two earnings groups. In this specification, we obtain very similar results as for the total sample, see Appendix A.7.4 for details.

7 Conclusion

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

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

A Appendix

A.1 Deriving an unbiased estimator of the Frisch elasticity: Analytical details

A.1.1 Households unaffected by borrowing constraints

For households unaffected by borrowing constraints, the multiplier on the borrowing constraint equals zero, $\phi = 0$. Taking logs of the first-order conditions (7), (8), and (9) gives

$$\ln \lambda + \ln w_1 = \ln \mu + \ln \alpha_1 + \frac{1}{\eta} \cdot \ln n_1, \tag{23}$$

$$\ln \lambda + \ln w_2 = \ln \mu + \ln \alpha_2 + \frac{1}{\eta} \cdot \ln n_2, \tag{24}$$

$$\ln \lambda = \ln (1 + r) + \ln \beta + \ln \mathrm{E}\, \lambda'. \tag{25}$$

Iterating forward (23) and (24) and taking first differences gives

$$\Delta \ln n_1' = \eta \cdot \Delta \ln w_1' + \eta \cdot \Delta \ln \lambda', \tag{26}$$

$$\Delta \ln n_2' = \eta \cdot \Delta \ln w_2' + \eta \cdot \Delta \ln \lambda', \tag{27}$$

where, for a generic variable y, $\Delta \ln y' = \ln y' - \ln y$. The next step is to use the Euler equation to substitute for the unobservable term $\Delta \ln \lambda'$. Rearranging the Euler equation (25), we obtain

$$\ln \mathrm{E}\, \lambda' - \ln \lambda = - \ln (1 + r) - \ln \beta.$$

While the Euler equation relates to the expected marginal utility of consumption, it is the actual change in the marginal utility of consumption that is part of the labor-supply conditions, see (26) and (27). We therefore introduce an expectation error $\xi' = \ln \lambda' - \mathrm{E} \ln \lambda'$ such that

$$\Delta \ln \lambda' = \ln \lambda' - \ln \lambda = - \ln (1 + r) - \ln \beta - \xi'. \tag{28}$$

In (28), we use that $\ln \mathrm{E}\, \lambda' = \mathrm{E} \ln \lambda'$ up to first order and we neglect higher-order terms. In Appendix A.4, we evaluate the importance of the approximation for our results.

Substituting (28) into (26) gives

$$\Delta \ln n_1' = \eta \cdot \Delta \ln w_1' - \eta \cdot \ln (1 + r) - \eta \cdot \ln \beta - \eta \cdot \xi', \tag{29}$$

$$\Delta \ln n_2' = \eta \cdot \Delta \ln w_2' - \eta \cdot \ln (1 + r) - \eta \cdot \ln \beta - \eta \cdot \xi'. \tag{30}$$

Note that the residual $\eta \cdot \xi'$ is correlated with $\Delta \ln w_1'$ as well as with $\Delta \ln w_2'$. Intuitively, wage growth, if not perfectly foreseen, leads to an increase in consumption and a reduction of marginal utility compared to their previously planned levels. Put differently, a household that enjoys unforeseen wage growth for either spouse will increase its consumption and marginal utility will be lower than it was expected one period before, $\xi' = \ln \lambda' - \mathrm{E} \ln \lambda' < 0$. To address the resulting endogeneity problem, Altonji (1986) suggests a decomposition of wage growth into an expected and an unexpected component,

$$\Delta \ln w_1' = \mathrm{E}\, \Delta \ln w_1' + \omega_1',$$
$$\Delta \ln w_2' = \mathrm{E}\, \Delta \ln w_2' + \omega_2',$$

where ω_1' and ω_2' are the unexpected components of wage growth. Using these decompositions in (29) and (30) gives (13) for $i = 1, 2$ from the main text. These equations can be estimated by OLS as the expectation errors ξ' and ω_i', $i = 1, 2$, which form the joint residual, are both uncorrelated with the regressor $\Delta\, \mathrm{E} \ln w_i'$: Under rational expectations, the expectation error ω_i' is uncorrelated with the expectation $\mathrm{E}\, \Delta \ln w_i'$ itself. Further, expected wage growth $\mathrm{E}\, \Delta \ln w_i'$ does not cause an adjustment of the marginal utility of consumption and hence a non-zero value of ξ' if the household is not affected by borrowing constraints. Intuitively, the household will use expected wage growth to finance already contemporaneous increases in consumption through either dis-saving or borrowing.

A.1.2 Borrowing-constrained households

To derive (14), (15), and (16), we first take logs of (6), (8), and (9):

$$-\sigma \ln c = \ln \lambda, \tag{31}$$
$$\ln \lambda + \ln w_1 = \ln \mu + \ln \alpha_1 + \frac{1}{\eta_1} \cdot \ln n_1, \tag{32}$$
$$\ln \lambda + \ln w_2 = \ln \mu + \ln \alpha_2 + \frac{1}{\eta_2} \cdot \ln n_2. \tag{33}$$

In the point of approximation, it holds that

$$-\sigma \ln \overline{c} = \ln \overline{\lambda}, \tag{34}$$
$$\ln \overline{\lambda} + \ln \overline{w}_1 = \ln \mu + \ln \alpha_1 + \frac{1}{\eta_1} \cdot \ln \overline{n}_1, \tag{35}$$
$$\ln \overline{\lambda} + \ln \overline{w}_2 = \ln \mu + \ln \alpha_2 + \frac{1}{\eta_2} \cdot \ln \overline{n}_2. \tag{36}$$

Subtracting (34), (35), and (36) from (31), (32), and (33), respectively, gives

$$\ln\left(\lambda/\overline{\lambda}\right) = -\sigma \ln\left(c/\overline{c}\right), \tag{37}$$

as well as the expressions (14) and (15) from the main text.

To obtain (16), we log-linearize the budget constraint. For borrowing-constrained households ($a = a' = 0$), the budget constraint reads

$$c = w_1 n_1 + w_2 n_2.$$

Applying a first-order Taylor approximation gives[1]

$$c - \overline{c} = \overline{w}_1 \cdot (n_1 - \overline{n}_1) + \overline{n}_1 \cdot (w_1 - \overline{w}_1) + \overline{w}_2 \cdot (n_2 - \overline{n}_2) + \overline{n}_2 \cdot (w_2 - \overline{w}_2).$$

Dividing by $\overline{c} = \overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2$ and expanding the terms on the right-hand side gives

$$\begin{aligned}\frac{c-\overline{c}}{\overline{c}} =& \frac{\overline{w}_1 \cdot \overline{n}_1}{\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2} \cdot \frac{n_1 - \overline{n}_1}{\overline{n}_1} + \frac{\overline{n}_1 \cdot \overline{w}_1}{\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2} \cdot \frac{w_1 - \overline{w}_1}{\overline{w}_1} \\ &+ \frac{\overline{w}_2 \cdot \overline{n}_2}{\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2} \cdot \frac{n_2 - \overline{n}_2}{\overline{n}_2} + \frac{\overline{n}_2 \cdot \overline{w}_2}{\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2} \cdot \frac{w_2 - \overline{w}_2}{\overline{w}_2}.\end{aligned}$$

Using the definitions $\overline{s}_1 = \overline{w}_1 \cdot \overline{n}_1 / \left(\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2\right)$ and $\overline{s}_2 = \overline{w}_2 \cdot \overline{n}_2 / \left(\overline{w}_1\overline{n}_1 + \overline{w}_2\overline{n}_2\right)$, and using that, for a generic variable y, $\left(y - \overline{y}\right)/\overline{y} \approx \ln\left(y/\overline{y}\right)$, we obtain

$$\ln\left(c/\overline{c}\right) = \overline{s}_1 \cdot \left(\ln\left(w_1/\overline{w}_1\right) + \ln\left(n_1/\overline{n}_1\right)\right) + \overline{s}_2 \cdot \left(\ln\left(w_2/\overline{w}_2\right) + \ln\left(n_2/\overline{n}_2\right)\right). \tag{38}$$

Substituting (38) into (37) gives (16) from the main text.

To obtain the labor-supply policy functions, we solve (14), (15), and (16) for $\ln\left(n_i/\overline{n}_i\right)$, $i = 1, 2$,

$$\begin{aligned}\ln\left(n_i/\overline{n}_i\right) =& \left(\eta - \frac{\sigma\eta\left(\eta+1\right)}{\sigma\eta+1} \cdot \overline{s}_i\right) \cdot \ln\left(w_i/\overline{w}_i\right) \\ &- \frac{\sigma\eta\left(\eta+1\right)}{\sigma\eta+1} \cdot \left(1 - \overline{s}_i\right) \cdot \ln\left(w_{-i}/\overline{w}_{-i}\right),\end{aligned} \tag{39}$$

which show that labor supply depends on one's own wage rate as well as on the wage rate of the partner. As we want to recover the own-wage Frisch elasticity, we will focus on the reaction to one's own wage-rate shocks and summarize the response to the partner's wage

[1] In Appendix A.4, we evaluate the importance of approximation errors for our results.

rate in a residual. Taking first differences yields

$$\Delta \ln n_i' = \ln\left(n_i'/\overline{n}_i\right) - \ln\left(n_i/\overline{n}_i\right) = \\ \left(\eta - \frac{\sigma\eta\left(\eta+1\right)}{\sigma\eta+1}\cdot\overline{s}_i\right)\cdot\Delta\ln w_i' + \zeta',$$

where $\zeta' = -\frac{\sigma\eta(\eta+1)}{\sigma\eta+1}\cdot(1-\overline{s}_i)\cdot\Delta\ln w_{-i}'$. As for unconstrained households, we can now decompose wage growth $\Delta\ln w_i'$ into an expected and an unexpected component, where the latter becomes part of the combined residual. This gives equation (17) from the main text.

A.2 Generalizing the relation between the borrowing-constraint bias and the relative earnings contribution

In this Appendix, we show that the estimation bias due to borrowing constraints in Altonji (1986) regressions is decreasing in an individual's contribution to household earnings for an arbitrary data frequency. For the general case, the estimation bias cannot be derived in closed form but we can express it in a comprehensive way as a function of endogenous moments as also done by Domeij and Floden (2006). First, we show that a standard Altonji (1986) regression is subject to a downward bias also in our two-earner model. This result and its derivation are very similar to the corresponding result in the bachelor model of Domeij and Floden (2006). Then, we show that this bias is the less important the less an individual contributes to household earnings. This second result critically requires the double-earner structure.

We start by deriving the true relation between hours growth and expected wage growth in our model. Taking logs and first differences of the labor-supply conditions (8) and (9), we obtain (for $i = 1, 2$)

$$\Delta \ln w_i' = \frac{1}{\eta} \cdot \Delta \ln n_i - \Delta \ln \lambda'. \tag{40}$$

A log-linear approximation of the Euler equation (7) implies

$$\Delta \ln \lambda' = -\ln \beta - \ln (1 + r) - \frac{\phi}{\lambda} + \xi', \tag{41}$$

where $\xi' = \ln \lambda' - \mathrm{E} \ln \lambda'$ denotes an expectation error. Inserting (41) into (40) and rearranging yields

$$\Delta \ln n_i' = \eta \cdot \Delta \ln w_i' - \eta \cdot \ln \beta - \eta \cdot \ln (1 + r) - \eta \cdot \frac{\phi}{\lambda} + \eta \cdot \xi'. \tag{42}$$

Note that this expression nests both relations, for unconstrained and constrained households, respectively, derived in the main text. For households unaffected by borrowing constraints, $\phi = 0$, such that (42) simplifies to (13) for $i = 1, 2$, when using $\Delta \ln w_i' = \Delta \mathrm{E} \ln w_i' + \omega_i'$. For borrowing-constrained households, the multiplier ratio ϕ/λ can be substituted to obtain (17). Further, (42) also covers households who transition between being borrowing constrained and unconstrained.

As pointed out by Altonji (1986), the residual ξ' is correlated with wage growth $\Delta \ln w_i'$ but not with expected wage growth $\mathrm{E} \Delta \ln w_i'$. As in Section 3.1, one therefore applies a decomposition of $\Delta \ln w_i'$ into $\mathrm{E} \Delta \ln w_i'$ and an unexpected component ω_i'. However, the multiplier ratio ϕ/λ which measures the household's willingness to borrow affects labor supply but is not observable in empirical data and is therefore part of the combined residual. This causes the estimation bias due to borrowing constraints because the willingness to borrow is correlated with wage growth.[2]

We will now show that the associated estimation bias differs depending on an individual's earner role in the household. In our derivations, we distinguish between two cases regarding the observability of expected wage growth. First, when using synthetic data, one can calculate the expected wage change $\mathrm{E}\,\Delta \ln w_i'$ from the particular wage process assumed. In the quantitative evaluations of our model, we assume an AR(1) process with fixed effects, $w_i' = \psi_i + z_i$, $z_i' = \rho \cdot z_i + \varepsilon_i'$, where ψ_i is a fixed effect, ρ determines persistence and ε is a wage-rate shock. We also apply this assumption here. Thus, we can calculate expected wage growth as

$$\mathrm{E}\,\Delta \ln w_i' = \mathrm{E}(z_i' - z_i) = \mathrm{E}((\rho - 1) \cdot z_i + \varepsilon_i') = (\rho - 1) \cdot z_i. \tag{43}$$

Second, 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). In the following, we cover both cases and derive important properties of the results of Altonji (1986) regressions in our model.

Proposition 1 *When expected wage growth $\mathrm{E}\,\Delta \ln w'$ is known, the percentage bias in the estimate for the Frisch elasticity in an Altonji (1986) regression is*

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

If expected wage growth is identified using an instrument y, the bias is

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

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

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

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

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

$$\widehat{\eta} = \frac{\mathrm{cov}\,(\Delta \ln n', \mathrm{E}\,\Delta \ln w')}{\mathrm{var}\,(\mathrm{E}\,\Delta \ln w')} = \frac{\mathrm{cov}\,(\Delta \ln n', (\rho-1)\,z)}{\mathrm{var}\,((\rho-1)\,z)} = \frac{\mathrm{cov}\,(\Delta \ln n', z)}{(\rho-1)\cdot \mathrm{var}\,(z)}. \tag{44}$$

[2] This problem is not resolved by using expected wage-rate 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-load consumption through borrowing.

We rearrange the true labor-supply relation (42) to

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

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

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

Inserting this into (44) gives the estimate as

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

The percentage bias then is

$$\begin{aligned}\frac{\widehat{\eta}-\eta}{\eta}&=-\frac{\operatorname{cov}\left(\frac{\phi}{\lambda},z\right)}{(\rho-1)\operatorname{var}(z)}=-\frac{\operatorname{cov}\left(\phi/\lambda,\mathrm{E}\,\Delta\ln w'/(\rho-1)\right)}{(\rho-1)\operatorname{var}(z)}\\&=-\frac{\operatorname{cov}\left(\phi/\lambda,\mathrm{E}\,\Delta\ln w'\right)}{(\rho-1)^2\cdot\operatorname{var}(z)}=-\frac{1+\rho}{(1-\rho)\cdot\sigma_\varepsilon^2}\cdot\operatorname{cov}\left(\phi/\lambda,\mathrm{E}\,\Delta\ln w'\right),\end{aligned}\tag{45}$$

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

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

$$\widehat{\gamma}=\frac{\operatorname{cov}\left(\Delta\ln w',y\right)}{\operatorname{var}(y)},\quad\widehat{\Delta\ln w'}=const_1+\frac{\operatorname{cov}\left(\Delta\ln w',y\right)}{\operatorname{var}(y)}\cdot y.$$

Then, the second-stage regression is

$$\Delta\ln n'=const+\eta\cdot\frac{\operatorname{cov}\left(\Delta\ln w',y\right)}{\operatorname{var}(y)}\cdot y+u,$$

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

$$\widehat{\eta}=\frac{\operatorname{cov}\left(\frac{\operatorname{cov}(\Delta\ln w',y)}{\operatorname{var}(y)}\cdot y,\Delta\ln n'\right)}{\operatorname{var}\left(\frac{\operatorname{cov}(\Delta\ln w',y)}{\operatorname{var}(y)}\cdot y\right)}=\frac{\frac{\operatorname{cov}(\Delta\ln w',y)}{\operatorname{var}(y)}\cdot\operatorname{cov}\left(y,\Delta\ln n'\right)}{\left(\frac{\operatorname{cov}(\Delta\ln w',y)}{\operatorname{var}(y)}\right)^2\cdot\operatorname{var}(y)}=\frac{\operatorname{cov}\left(y,\Delta\ln n'\right)}{\operatorname{cov}\left(\Delta\ln w',y\right)}.\tag{46}$$

We rearrange the true relation (42) to

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

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

$$\operatorname{cov}\left(\Delta \ln n', y\right) = \eta\gamma \cdot \operatorname{var}(y) - \eta \cdot \operatorname{cov}(y, \phi/\lambda).$$

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

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

We can use the first-stage results to determine the covariance in the denominator of the estimate $\widehat{\eta}$ in (46) as

$$\operatorname{cov}\left(\Delta \ln w', y\right) = \widehat{\gamma} \cdot \operatorname{var}(y).$$

Inserting this into (46) gives

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

Hence, the percentage bias is

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

as stated in the proposition. ■

For the following analytical results, we consider a simplified version of the model, where we assume $\eta_1 = \eta_2 = \eta$, $\alpha_1 = \alpha_2 = 2$, $\mu = 1/2$, $\rho = 0$, and $\operatorname{var}(\varepsilon_1) = \operatorname{var}(\varepsilon_2)$. Importantly, the assumptions on the preference weights α and μ imply that differences in long-run earner roles stem solely from differences in the wage fixed effects ψ. The higher a spouse's fixed effect, the more this person contributes to household earnings (see Lemma 3 below). Therefore, it is convenient to first perform comparative statistics in the fixed effects ψ and to transfer the results to relative earner roles thereafter.

As an intermediate step, we can state that the multiplier ratio in the true labor-supply relation (42) 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. 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 (7), the multiplier ratio ϕ/λ is

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

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 $\mathrm{E}\,\lambda'$ in case of being borrowing constrained as a constant, which we denote by $\overline{\lambda}$.

Taken together, we can express the multiplier ratio as

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

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

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:

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

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

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

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

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

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

Second, the covariance between the symmetric function Λ of the spouses' wage rates and an individual's stochastic wage component z_i increases in the individual's relative wage-fixed effect in the household Ψ_i/Ψ_{-i}. This implies that the relative strength of the estimation bias also increases in the relative wage fixed effect.

Lemma 2 *The relative estimation bias in the household,* $\mathrm{bias}_i\,/\,\mathrm{bias}_{-i}$ *is an increasing function of the relative wage fixed effects,* Ψ_i/Ψ_{-i}*. Further,* $\mathrm{bias}_i \to 0$ *for* $\Psi_i \to 0$ *and for* $\Psi_{-i} \to \infty$*.*

Proof. From Proposition 1, it follows that the relative estimation bias in the household is given by $\operatorname{cov}(\phi/\lambda, \mathrm{E}\ln w_i' - \ln w_i) / \operatorname{cov}\left(\phi/\lambda, \mathrm{E}\ln w_{-i}' - \ln w_{-i}\right)$. Using the specification of the wage process, this is equivalent to

$$\frac{\text{bias}_i}{\text{bias}_{-i}} = \operatorname{cov}\left(\phi/\lambda, z_i'\right) / \operatorname{cov}\left(\phi/\lambda, z_{-i}'\right). \tag{47}$$

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 $\operatorname{cov}(\phi/\lambda, z_i) \leq 0$.

Now, consider the covariance between Λ and the individual stochastic wage component z_i. For household member i, $\operatorname{cov}(\Lambda, Z_i) = \operatorname{cov}\left(\Psi_i^{1+\eta} Z_i^{1+\eta} + Z_{-i}^{1+\eta}, Z_{-i}\right) = \Psi_i^{1+\eta} \cdot \operatorname{cov}\left(Z_i^{1+\eta}, Z_i\right) + \operatorname{cov}\left(Z_i^{1+\eta}, Z_{-i}\right) = \Psi_i^{1+\eta} \cdot \operatorname{cov}\left(Z_i^{1+\eta}, Z_i\right)$. Since $\operatorname{cov}\left(Z_i^{1+\eta}, Z_i\right)$ is an exogenous constant determined by the parameters of the wage process, $\operatorname{cov}(\Lambda, Z_i)$ is an increasing function of Ψ_i with $\operatorname{cov}(\Lambda, Z_i) \to 0$ for $\Psi_i \to 0$. For the other household member $-i$, we have $\operatorname{cov}(\Lambda, Z_{-i}) = \operatorname{cov}\left(Z_{-i}^{1+\eta}, Z_{-i}\right)$. Due to $\rho = 0$, and $\operatorname{var}(\varepsilon_1) = \operatorname{var}(\varepsilon_2)$, we can also use that $\operatorname{cov}\left(Z_{-i}^{1+\eta}, Z_{-i}\right) = \operatorname{cov}\left(Z_i^{1+\eta}, Z_i\right)$ which implies that

$$\frac{\operatorname{cov}(\Lambda, Z_i)}{\operatorname{cov}(\Lambda, Z_{-i})} = \frac{\Psi_i^{1+\eta}}{\Psi_{-i}^{1+\eta}}. \tag{48}$$

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 proof continues by stating results conditional on the beginning-of-period asset holdings a and later continues by aggregating over a.

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

$$\begin{aligned}\operatorname{cov}(f(\Lambda), Z_i|a) &= \mathrm{E}(\partial(\phi/\lambda)/\partial\Lambda|a) \cdot \operatorname{cov}(\Lambda, Z_i|a)\\ &= \mathrm{E}(\partial(\phi/\lambda)/\partial\Lambda|a) \cdot \operatorname{cov}(\Lambda, Z_i).\end{aligned}$$

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.

Now, aggregating over the different a gives $\operatorname{cov}(f(\Lambda), Z_i) = \int \operatorname{cov}(f(\Lambda), Z_i|a) \;\mathrm{d}\,h(a) = \operatorname{cov}(\Lambda, Z_i) \cdot \int \mathrm{E}(\partial(\phi/\lambda)/\partial\Lambda|a) \;\mathrm{d}\,h(a)$ and hence

$$\frac{\operatorname{cov}(f(\Lambda), Z_i)}{\operatorname{cov}(f(\Lambda), Z_{-i})} = \frac{\operatorname{cov}(\Lambda, Z_i) \cdot \int \mathrm{E}(\partial(\phi/\lambda)/\partial\Lambda|a) \;\mathrm{d}\,h(a)}{\operatorname{cov}(\Lambda, Z_i) \cdot \int \mathrm{E}(\partial(\phi/\lambda)/\partial\Lambda|a) \;\mathrm{d}\,h(a)} = \frac{\operatorname{cov}(\Lambda, Z_i)}{\operatorname{cov}(\Lambda, Z_i)} = \frac{\Psi_i^{1+\eta}}{\Psi_{-i}^{1+\eta}}, \tag{49}$$

where the final step uses (48).

Finally, we use $Z_i \approx 1 + z_i$ and, hence, $\operatorname{cov}(\phi/\lambda, Z_g) \approx \operatorname{cov}(\phi/\lambda, z_g)$ and combine (47) and (49) to

$$\frac{\text{bias}_i}{\text{bias}_{-i}} \approx \frac{\Psi_i^{1+\eta}}{\Psi_{-i}^{1+\eta}}.$$

From this, it follows that the estimation bias for household member approaches zero if either $\Psi_i \to 0$ or $\Psi_{-i} \to \infty$. Further, the estimation bias for household member i is strictly increasing in his or her relative wage fixed effect Ψ_i/Ψ_{-i}. ■

Lemma 3 *The relative contribution of household member i to household earnings is monotonically increasing in his or her relative wage fixed effect,* Ψ_i/Ψ_{-i}.

Proof. Consider the first-order conditions for labor supply, (8) and (9), to see $w_i n_i / (w_{-i} n_{-i}) = (w_i/w_{-i})^{1+\eta}$ for any combination of w_i, w_{-i}, a. Hence, for any asset level, relative earnings are an increasing function of relative wage rates. Taking logs gives $\ln w_i n_i - \ln w_{-i} n_{-i} = (1+\eta) \cdot (\ln w_1 - \ln w_2)$ and taking expectations gives

$$\mathrm{E}\left(\ln w_1 n_1 - \ln w_2 n_2\right) = (1+\eta) \cdot \left(\mathrm{E} \ln w_1 - \mathrm{E} \ln w_2\right) = (1+\eta) \cdot \ln\left(\Psi_1/\Psi_2\right).$$

Hence, the mean earnings gap is an increasing function of the relative wage fixed effects Ψ_i/Ψ_{-i}. Finally, it is straightforward to show that the relation between the earnings gap and the earnings contribution is positive and monotonic. Define $v = \ln w_1 n_1 - \ln w_2 n_2$ and $x_1 = w_1 n_1 / (w_1 n_1 + w_2 n_2)$. Then $s = v/(v+1)$. ■

Proposition 4 *The estimation bias* $(\widehat{\eta} - \eta)/\eta$ *for an individual is monotonically related to the individual's percentage contribution to household earnings: The higher is the contribution to household earnings, the stronger is the estimation bias. The estimation bias converges to zero for individuals whose percentage contribution to household earnings converges to zero.*

Proof. Follows directly from Lemma 2 together with Lemma 3. ■

A.3 Sample selection and calibration

A.3.1 Sample selection

We use observations for the years 1972-1997 from the PSID. Before 1972, there is no information on wives' education. After 1997, the PSID switched from annual to biennial interviews. We consider household heads and wives for whom both partners' wage rates are available. 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. We restrict the sample to 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 (see Domeij and Floden 2006). 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.

Note that our baseline sample is relatively large in comparison to related studies. For instance, Domeij and Floden (2006) have to restrict the data to three subpanels around the years 1984, 1989 and 1994, where the PSID contained a supplement on household wealth. In our specifications where we exploit asset holdings to control for borrowing constraints, we also consider this subsample of the PSID data, see Section 5 and Appendix A.7.3.

A.3.2 Estimation of stochastic wage process

We distinguish between two components of wage rates. First, there is a component that captures (observed or unobserved) characteristics of the individual and that leads to long-run wage differences between individuals (both, within and across households). Second, there is a component that captures idiosyncratic and temporary fluctuations in wage rates which may induce borrowing constraints to bind. We use a combination of microeconometric estimation (idiosyncratic and temporary wage components) and calibration (long-run wage-rate differences within and across households) to obtain the parameters of the model.

Idiosyncratic and temporary wage components. We assume in our model that the idiosyncratic and temporary component follows an AR(1) process. The parameters of this process are important for frequency and expected duration of binding borrowing constraints and for the process of expected wage growth which is key to labor-supply regressions. We quantify the parameters of the AR(1) process through gender-specific Generalized Method of Moments (GMM) estimations. As discussed in Section 4.1, we quantify long-run wage differences through calibration.

To determine the parameters of the AR(1) process, the following steps have to be carried out that do not have direct counterparts in the theoretical model. To save on notation, we do not account for a gender index $g = m, f$ in the following. We first filter predictable influences on observed log wage rates $\ln \widetilde{w}_{it}^{*}$ using a filter regression

$$\ln \widetilde{w}_{it}^{*} = q\left(o_{it}\right) + \widehat{w}_{it},$$

where o_{it} denote characteristics (time dummies, age dummies, and education dummies interacted up to a quadratic age trend) of individual i in year t.[3] The wage process estimation is then performed for residual log wage rates $\widehat{w}_{it}$.[4]

In the empirical process for residual wages, we account for individual fixed effects and we follow the applied literature, see, e.g., Heathcote, Perri, and Violante (2010), by incorporating time-varying factor loadings that allow the permanent and transitory components to change over time in a way that is common across individuals. While the variance terms are time-invariant in our theoretical model, the empirical literature has shown the importance of allowing for such flexibility in the estimated processes to correctly identify persistence and idiosyncratic risk. Hence, residual log wages are decomposed into

$$\widehat{w}_{it} = \pi_t \cdot \chi_i + \zeta_t \cdot \widetilde{z}_{it}, \tag{50}$$

where π_t and ζ_t are factor loadings, χ_i is an individual fixed effect, and $\widetilde{z}_{it}$ is the transitory component of observed wage rates.

Next to an autoregressive component, we also incorporate a moving-average term in the process for the transitory wage component $\widetilde{z}_{it}$ to take into account measurement error in empirical wage-rate data that would otherwise bias the estimated AR coefficients[5],

$$\widetilde{z}_{it}, = \rho \cdot \widetilde{z}_{it-1} + \theta \cdot \varepsilon_{it-1} + \varepsilon_{it}, \tag{51}$$

where ρ is the persistence parameter, θ is the MA parameter, and ε_{it} is the shock to the transitory wage component with variance σ_{ε}^{2}.

We estimate the process for residual wages, (50) and (51), by a Generalized Method

[3] 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 filter regression to obtain the final vector of residual wage rates, $\widehat{w}_{it}$.

[4] A measure of dispersion after controlling for observables is the relevant concept in the incomplete markets literature where the focus is on idiosyncratic uncertainty, see, e.g., Storesletten et al. (2004), Krueger and Perri (2006), Heathcote, Storesletten, and Violante (2010), and Bayer and Juessen (2012).

[5] Accounting for a moving-average component to deal with measurement error is widely adopted in the empirical literature, see, e.g., Meghir and Pistaferri (2004).

of Moments (GMM). Specifically, we exploit as moment conditions the variance-covariance matrix of residual wages $\widehat{w}_{it}$ which has diagonal elements

$$\sigma_1^2 = \pi_1^2 \sigma_\chi^2 + \zeta_t^2 \sigma_{w1}^2$$
$$\sigma_t^2 = \pi_t^2 \sigma_\chi^2 + \left\{ \zeta_t^2 \left(\rho^{2t-2} \sigma_{w1}^2 + K \sum_{j=0}^{t-2} \rho^{2j} \right) \right\}, t > 1$$

and off-diagonal elements

$$\operatorname{cov}\left(\widehat{w}_t, \widehat{w}_{t+s}\right) = \pi_t \pi_{t+s} \sigma_\chi^2 + \zeta_t \zeta_{t+s} \left(\rho^s \sigma_{w1}^2 + \rho^{s-1} \theta \sigma_\varepsilon^2 \right), t = 1 \text{ and } s > 0$$
$$\operatorname{cov}\left(\widehat{w}_t, \widehat{w}_{t+s}\right) = \pi_t \pi_{t+s} \sigma_\chi^2 + \zeta_t \zeta_{t+s} \left(\rho^{2t+s-2} \sigma_{w1}^2 + \rho^2 K \sum_{j=0}^{t-1} \rho^{2j} + \rho^{s-1} \theta \sigma_\varepsilon^2 \right), t > 1 \text{ and } s > 0,$$

where $K = \sigma_\varepsilon^2 \left(1 + \theta^2 + 2\rho\theta\right)$. Concerning initial conditions, we follow the approach by MaCurdy (1982) and treat the variance at the start of the sample period, σ_{w1}^2, as an additional parameter to be estimated. GMM estimation is carried out by replacing population moment conditions by their sample analogues. For both genders, the parameter vector to be estimated is

$$\varphi = \left\{ \sigma_\chi^2, \rho, \sigma_{w1}^2, \sigma_\varepsilon^2, \theta, \zeta_2 ... \zeta_T, \pi_2 ... \pi_T \right\}.$$

For identification, the first-period factor loadings π_1, ζ_1 are set to one. In our unbalanced panel data, each sample moment is constructed using all available observations covering the respective time span.[6] 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 6 summarizes the parameter estimates.[7] For men, the estimated autocorrelation of idiosyncratic wages is $\rho_m = 0.82$ and the estimated standard deviation is $\sigma_{m,\varepsilon} = 0.22$. These results are well in line with the literature, see, e.g., Domeij and Floden (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.

Long-run wage-rate differences within and across households. Next to the autoregressive component, we include constant terms ψ in the wage process of our model to induce

[6] 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.

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

Table 6: Estimated processes for residual wage rates, men and women.

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

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

long-run differences in earner roles. We quantify these constant terms by calibration. Specifically, we distinguish between ten household types matching average male and female wage rates in the ten deciles of the empirical distribution of relative wage rates of spouses in our PSID data. An alternative approach would be to target the estimated variance of fixed effects, σ^2_χ from the microeconometric wage process estimation. While this would capture the gender-specific *across*-household variance of (residual) wage rates appropriately, we implement the former approach to obtain a realistic distribution of *within*-household wage differences.

A.3.3 Parameter values

Table 7 summarizes the parameter values of our baseline model.

Table 7: Parameter values, baseline model

Description	Parameter	Value(s)										
Aggregate parameters												
Interest rate	r	1.45%										cal.
Discount factor	β	0.95										set
Risk aversion	σ	1.5										set
Pareto weight	μ	0.5										set
Gender-specific parameters		male					female					
Autocorrelation z	ρ_g	0.84					0.81					est.
Standard deviation z	σ_g	0.24					0.44					est.
True Frisch elasticity	η_g	0.65					0.90					est.
HH-type specific parameters		I	II	III	IV	V	VI	VII	VIII	IX	X	
Const. wage comp. husband	ψ_{mj}	-2.8	-2.6	-2.6	-2.5	-2.4	-2.5	-2.3	-2.3	-2.2	-2.0	cal.
Const. wage comp. wife	ψ_{fj}	-2.5	-2.7	-2.8	-2.9	-2.9	-3.1	-3.1	-3.2	-3.3	-3.5	cal.
Labor-disutility husband	α_{mj}	48.4	61.2	68.3	73.7	75.8	84.1	84.5	86.4	99.6	105.7	cal.
Labor-disutility wife	α_{fj}	60.5	51.1	49.3	47.4	45.3	45.2	43.1	42.0	41.9	36.1	cal.

NOTE.–HH-type: household type. Roman numbers (I-X) indicate household types. cal.: calibrated. est.: estimated

A.4 Log-linearization and approximation bias

The derivation of our interaction-term approach (20) uses first-order approximations. In this appendix, we use second-order Taylor approximations to evaluate the importance of approximation errors.

Compared to the expressions in Appendix A.1.1, a second-order approximation of the Euler equation for unconstrained households yields the additional term, $\xi' + \frac{1}{2}(\xi')^2$, where ξ' is the ex-post percentage expectation error in marginal utility $\xi' = (\lambda' - E\lambda')/\lambda'$, see Domeij and Floden (2006) for a derivation. When substituted into the labor-supply conditions (26) and (27), this higher-order term enters additively and multiplied with a negative constant on the right-hand side of equation (13) in the main text.

For borrowing-constrained households, we have to approximate the budget constraint (2). A second-order approximation yields the additional term $\overline{s}_i \cdot (\ln(w_i/\overline{w}_i) \cdot \ln(n_i/\overline{n}_i)) + \overline{s}_{-i} \cdot (\ln(w_{-i}/\overline{w}_{-i}) \cdot \ln(n_{-i}/\overline{n}_{-i}))$ compared to the expressions in Appendix A.1.2. When substituted into the system of first-order conditions, this higher order term enters in first differences, additively, and multiplied with a negative constant on the right-hand side of equation (16) in the main text.

To evaluate the importance of approximation errors, we calculate both higher-order terms (h.o.t.) for unconstrained and constrained households, respectively, in our simulation. When determining the expectation term in the expression for constrained households, we first calculate ξ and ξ^2 as functions of state variables and then determine their expectations by applying transition probabilities.[8] We then include the h.o.t. as additional regressors in our interaction-term regression (20). Table 8 shows that this has almost no effect on the non-interacted coefficient on expected wage growth (which is the estimated Frisch elasticity). Thus, higher-order terms are almost irrelevant for our results and our preferred first-order approximation is sufficient for obtaining an almost unbiased estimate of the Frisch elasticity.

[8] A regression-based approach as applied by Domeij and Floden (2006) delivers almost identical results.

Table 8: Estimation results from synthetic household panel data, the role of approximation errors.

	(1)	(2)	(3)	(4)
expected	0.63	0.62	0.62	0.61
wage growth	(0.10)	(0.10)	(0.10)	(0.10)
expected wage growth	-0.32	-0.31	-0.32	-0.30
$\times$ earnings contribution (%)	(0.15)	(0.15)	(0.14)	(0.15)
h.o.t. (unconstrained)		0.10		-0.33
		(0.22)		(0.23)
h.o.t. (constrained)			-0.03	-0.03
			(0.00)	(0.00)

NOTE.–Results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the simulation. h.o.t. (unconstrained): $\mathrm{E}_t(\xi_{jt+1} + \xi_{jt+1}^2/2)$. h.o.t. (constrained): $\Delta\left(\overline{s}_{ij} \cdot (\ln(w_{ijt+1}/\overline{w}_{ij}) \cdot \ln(n_{ijt+1}/\overline{n}_{ij})) + \overline{s}_{-ij} \cdot (\ln(w_{-ijt+1}/\overline{w}_{-ij}) \cdot \ln(n_{-ijt+1}/\overline{n}_{-ij}))\right)$. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses.

A.5 Monte Carlo results for women

In this appendix, we repeat the baseline Monte Carlo experiments for women. Column (1) in Table 9 shows the results from a standard Altonji (1986) regression. As women's average earnings contribution is smaller than men's, the downward bias is less pronounced for women than it is for men. Column (2) shows the results of our preferred interaction-term approach, see equation (20). The remaining bias is only about 5%, comparable to the one for men.

Table 9: Estimation results from synthetic household panel data, for women, baseline model.

	(1)	(2)
expected wage growth	0.78	0.85
	(0.02)	(0.07)
expected wage growth × earnings contribution (%)		-0.18
		(0.19)
observations	15,000	15,000

NOTE.–Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the simulation. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses.

A.6 Model extensions

A.6.1 Model version with 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. 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

$$v = \frac{\left(\alpha \cdot c^{\gamma} + (1-\alpha) \cdot \left(\zeta \cdot n_i^{\theta} + (1-\zeta) \cdot n_2^{\theta}\right)^{-\gamma/\theta}\right)^{(1-\sigma)/\gamma} - 1}{1-\sigma}, \tag{52}$$

where c, n_1, and n_2 denote consumption, and hours of household members 1 and 2, respectively.[9]

Other than in the baseline model, we do not calibrate the elasticities of the utility function to match the empirical estimates of Altonji (1986) regressions. Since it is the purpose of this model extension to check our estimation approach in an environment with realistic degrees of complementarities, we take the substitution elasticities γ and θ directly 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. Similar to the calibration of our baseline model, we calibrate the preference weights α and ζ for each of the 10 household types to match average hours worked by husband and wife. Again, we set the economy-wide interest rate r to match the wealth share of the bottom 40% of the distribution. All other parameters are unchanged. In our calibrated model, the average true Frisch elasticity of men is 0.80.

Column (1) in Table 10 shows estimation results from a standard Altonji (1986) regression using simulated data from the model with non-separable preferences. We find that the downward bias is quantitatively similar. The estimated Frisch elasticity using our preferred interaction-term approach is much closer to the true value, see column (2).

[9] 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. Accordingly, the household directly maximizes $V = v + \beta \cdot \mathrm{E}\, V'$ in this model version.

Table 10: Estimation results from synthetic household panel data, model version with non-separable preferences.

	(1)	(2)
expected	0.42	0.62
wage growth	(0.02)	(0.10)
expected wage growth		-0.26
$\times$ earnings contribution (%)		(0.14)
observations	15,000	15,000

NOTE.–Results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the simulation. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses.

A.6.2 Model version with private consumption

For the case where consumption is a private good within the household, the household budget constraint (2) 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 both spouses and are choice variables. Further, instead of the first-order condition (6), it holds that $\frac{\partial v}{\partial c_1} = \mu \cdot c_1^{-\sigma} = \frac{\partial V(a,\omega)}{\partial a} = \lambda$ and $\frac{\partial v}{\partial c_2} = (1-\mu) \cdot c_2^{-\sigma} = \frac{\partial V(a,\omega)}{\partial a} = \lambda$. For the calibration, we use the same targets as in the baseline model. The true Frisch elasticity of men in this model version is also 0.65. Table 11 shows that we obtain very similar estimation results as in our baseline model with private consumption.

A.6.3 Model version with progressive taxation

In this appendix, we augment our baseline model by progressive taxation. We apply the parametric tax function used by Blundell et al. (2016) and Heathcote et al. (2017). Specifically, the average tax rate is

$$1 - (1-\xi) \cdot (w_1 n_1 + w_2 n_2)^{-\tau}\,, \tag{53}$$

such that the household budget constraint is $c + \frac{a'}{1+r} \leq (1-\xi) \cdot (w_1 n_1 + w_2 n_2)^{1-\tau} + a$, instead of (2). Guner et al. (2014) have estimated this tax function and we use their estimates for the "all married couples" sample to quantify the parameters in (53), $\xi = 0.1260$ and $\tau = 0.060$.[10] For the calibration, we use the same targets as in our baseline model. The true Frisch

[10] In Guner et al. (2014), the counterpart to ξ is $0.087 \cdot \overline{y}^{\tau}$, where 0.087 is an estimate and $\overline{y}$ is mean household income. In our model, this gives $\xi = 0.1260$. Guner et al. (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.

Table 11: Estimation results from synthetic household panel data, model version with private consumption.

	(1)	(2)
expected wage growth	0.41 (0.01)	0.64 (0.10)
expected wage growth × earnings contribution (%)		-0.34 (0.14)
observations	15,000	15,000

NOTE.–Results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the simulation. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses.

elasticity of men in this model version is 0.70.

Table 12 shows the results of labor-supply regressions for the model with progressive income taxation. As the estimation 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 estimation bias in an Altonji (1986) regression, see column (1), is somewhat stronger but overall similar to the baseline model. Put differently, the additional bias due to progressive taxation is small compared to the bias stemming from borrowing constraints. Most importantly, our proposed interaction-term approach delivers a good estimate of the Frisch elasticity also in this model version, see column (2) of Table 12.

A.6.4 Model version with preference shocks

With preferences given by (22), the first-order condition for labor supply of the wife is

$$(n_f + h)^{1/\eta} = \lambda \cdot w_f,$$

from which we can derive the Frisch elasticity as a function of h. Taking logs and the derivative with respect to w_f gives

$$\left.\frac{\partial n_f}{\partial w_f} \cdot \frac{w_f}{n_f}\right|_{\lambda} = \frac{n_f + h}{n_f} \cdot \eta.$$

Hence, a large realization of h increases the Frisch elasticity conditional on hours worked. By contrast, for $h = 0$, the Frisch elasticity is η as before.

Next to raising the Frisch elasticity, large values of h also increase the marginal disutility

Table 12: Estimation results from synthetic household panel data, model version with progressive taxation.

	(1)	(2)
expected	0.41	0.65
wage growth	(0.02)	(0.10)
expected wage growth		-0.36
$\times$ earnings contribution (%)		(0.15)
observations	15,000	15,000

NOTE.–Results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the simulation. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses.

of labor supply inducing women to work fewer hours.[11] Hence, for a given realization of the wife's wage rate, a large realization of h leads to low but elastic labor supply. As a consequence, one observes women with low contributions to household earnings (due to low labor supply) to have a strong connection between hours growth and expected wage growth.

We rely on Guner et al. (2011) and Bick (2016) to calibrate the new parameters in this model extension. We have to quantify (i) how many households are in the h_{high} state in the long-run distribution, (ii) for how long households stay in this state, and (iii) the numbers for h_{high} and h_{low}. For the first statistic, we refer to Bick (2016) who reports that about 90% of parents have grandparents living within one driving hour. Considering grandparents as an available form of informal child care, we target the population share with $h = h_{high}$ in the ergodic distribution, $\kappa_1 / (\kappa_1 + \kappa_2)$, to be 10%. We set the expected duration of $h = h_{high}$ to ten years (to mimic the time from child birth to start of middle school or junior high school). Hence, we set the transition probability κ_2 to 0.1. We set h_{low} to zero as Guner et al. (2011) do for women without (young) children, which implies that there is a large fraction of the population with a Frisch elasticity given by η as in our baseline model, and h_{high} to 0.1, which corresponds to a third of average male working time as in Guner et al. (2011).[12]

We calibrate the remaining parameters to match the same targets as in the baseline model. The decisive parameter is of course the curvature parameter in female labor disutility, η_f.

[11] In this model version, also non-participation of the wife is possible when $h > 0$. As in the empirical analysis, we drop households where one spouse does not work from the sample and consider only double-earner households in the estimations.

[12] While Guner et al. (2011) assign this value to all mothers but only while children are young (ages 0-4), we consider the case that a small fraction of household does not have access to child care at later ages of the children and, accordingly, has to bear these time costs for longer.

Table 13: Estimation results for women, from synthetic household panel data, extended model version with preference shocks.

	(1)	(2)	(3)	(4)
expected	0.78	0.97	0.91	0.87
wage growth	(0.02)	(0.04)	(0.04)	(0.08)
expected wage growth		-0.51	-0.35	
$\times$ earnings contribution (%)		(0.09)	(0.09)	
expected wage growth				-0.26
$\times$ predicted contribution				(0.24)
sample	all	all	$\bar{s}_{ij} \geq 0.3$	all
observations	15,000	15,000	15,000	15,000

NOTE.–Results for women. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution $\bar{s}_{ij}$ is the average percentage contribution of individual i to labor earnings of household j in a seven-period window around the observation. Average estimates from 10,000 Monte-Carlo draws, average standard errors in parentheses.

We calibrate $\eta_f = 0.84$ which directly translates into the true Frisch elasticity of the 90% of women with $h = 0$. From the remaining 10% of the female population with $h > 0$, about 9% supply zero hours and are hence excluded from the sample. The 9.2% of women in the remaining sample who have $h > 0$ supply 0.18 hours on average and hence have an average Frisch elasticity of about 1.3. Accordingly, the average Frisch elasticity in the total population of working women is about 0.88.

As in the baseline model, we use simulated data of double-earner households for labor-supply regressions. Table 13 shows the results for women while results for men are similar to those of the baseline model and are hence not shown. Column (1) shows that we have calibrated also the extended model to yield a coefficient on expected wage growth of 0.78 in a standard Altonji (1986) regression. Column (2) shows the results for our baseline interaction-term specification that we used for men in the main text. Here, it is important that we calculate the average earnings contribution of wives over a seven-year span mimicking the average length of observing a given household in our PSID sample. We find that, first, the coefficient on the interaction term is negative reflecting that individuals with low contributions to household earnings have a more elastic labor supply. Second, the coefficient on non-interacted expected wage growth (0.97) *exceeds* the true average Frisch elasticity (0.88). This shows that our baseline interaction-term approach puts relatively much weight on the small share of women with high Frisch elasticities.

Columns (3)-(4) address this point and provide guidance how to obtain a more representative estimate of the Frisch elasticity for the majority of the population. Instead of conditioning directly on h, which is difficult in empirical data, we consider approaches that can more easily be implemented in empirical applications. First, in column (3), we eliminate women with contributions to household earnings below 30%. In this sample, our baseline interaction-term approach yields an average Frisch elasticity of 0.91 which is closer to the economy-wide true female Frisch elasticity of 0.88 than the estimate from the full sample. While the sample restriction yields an improved estimate and is easy to implement empirically, one can improve further by taking out the idiosyncratic choices of hours worked when determining the earnings contribution. Specifically, in column (4), we consider an approach where we use predicted earnings contributions instead of the actually observed ones but estimate from the total sample of double-earner households. In the Monte Carlo experiment, we use the prediction from a regression of earnings contributions on household-type dummies which capture the only source of deterministic heterogeneity in our model.[13] We obtain an estimated Frisch elasticity of 0.87 which is close to the true value of 0.88. Thus, both of our modified approaches are robust in the presence of taste shocks and in particular our approach using predicted earnings contributions delivers almost unbiased estimates of the Frisch elasticity. Table 4 in the main text shows the results of both specifications estimated from PSID data.

[13] In an empirical application, this can be captured by characteristics such as education and age.

A.7 Additional regression results from PSID data

A.7.1 Alternative definitions of primary and secondary earners

In this appendix, we consider alternative definitions of primary and secondary earners to corroborate that standard Altonji (1986) regressions assign smaller estimates of the Frisch elasticity to primary earners than to secondary earners. To make sure that our results are not driven by gender differences together with the fact that most secondary earners are women, we rely on a within-gender perspective and report results for samples of men.

Column (1) of Table 14 replicates the results for men using our baseline definition where we defined primary earners as having the higher average wage rate over the observation period compared to the spouse (see Table 2). Column (2) of Table 14 refers to a specification where we compare average earnings rather than average wage rates, i.e., a husband is identified as primary earner if his average earnings exceed those of his wife. Compared to our baseline definition, this criterion has little within-gender variation as 90% of men in our sample have higher average earnings than their wives while only 81% have higher average wage rates.[14] Nevertheless, we obtain similar results as for our baseline definition.

In column (3) of Table 14, we define the earner status using a comparison of contemporaneous wage rates instead of average wage rates, i.e., the primary-earner dummy d_{jt} equals one if the husband's average wage rate in years t and $t+1$ exceeds the wife's average wage rate in these two years (also in this estimation, we consider a sample of males only). Compared to the baseline definition, this definition has the advantage that it does not include past or future earnings potentials. However, the disadvantage is that it is heavily affected by contemporaneous wage-rate shocks while our theoretical results relate to the usual contribution to household earnings. Again, we obtain similar results as for our baseline definition.

A.7.2 Altonji (1986) estimates in different ranges of the contribution to household earnings

In this appendix, we corroborate the negative relation between the relative contribution to household earnings and the estimated Frisch elasticity in standard Altonji (1986) regressions. Specifically, we run regressions where we compare several groups of workers defined through ranges of their percentage contribution to household earnings rather than using a binary distinction between primary and secondary earners. We include interaction terms between expected wage growth and three dummy variables that indicate whether the individual's average earnings contribution exceeds thresholds of one third, one half, and 90%, respectively.

[14] Note that, in our preferred interaction-term approach, the usual contribution to household earnings is a *continuous* variable and displays substantial within-gender variation.

Table 14: Empirical labor-supply regressions for men, PSID data, alternative definitions of primary and secondary earners.

	(1)	(2)	(3)
	average wage rates	average earnings	contemp. wage rates
expected wage growth	0.52	0.59	0.57
	(0.12)	(0.14)	(0.12)
expected wage growth × primary earner	-0.15	-0.21	-0.25
	(0.09)	(0.12)	(0.09)
time effects	yes	yes	yes
observations	14,340	14,340	14,340

NOTE.–Results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Individuals identified as primary earners if the mean realized wage rate in the sample $\overline{w}_{ij}$ exceeds the mean realized wage rate of the spouse $\overline{w}_{-ij}$ (column 1), the mean realized earnings in the sample $\overline{y}_i$ exceed the mean realized earnings of the spouse $\overline{y}_{-i}$ (column 2), and the mean realized wage rate in years t and $t+1$, $(w_{ijt}+w_{ijt+1})/2$, exceeds the mean realized wage rate of the spouse, $(w_{-ijt}+w_{-ijt+1})/2$ (column 3), respectively. Standard errors in parentheses.

The implied estimated elasticities for the different groups of individuals can be obtained by summing up the coefficients appropriately. Our model predicts all coefficients on these interaction terms to be negative as implied estimated elasticities should be smaller for groups who contribute larger shares to household earnings. Table 15 shows that this is confirmed in the PSID data. The group of husbands contributing more than 90% to household earnings is particularly interesting as these men are almost single earners in their respective households. In line with our model, we find that a standard Altonji (1986) regression yields a particularly small point estimate (0.13) for this group.

A.7.3 Conditioning on asset holdings in the Domeij and Flodén (2006) sample

In this appendix, we investigate the implication of our model that differences in estimated Frisch elasticities should become smaller when the samples are less affected by borrowing constraints. For this, we repeat the estimations including an interaction term with the primary-earner dummy but only consider households with liquid wealth above a certain threshold which we increase step by step. While information on wealth components is not available in most waves of the PSID, we can use the subsample constructed by Domeij and Floden (2006) for this purpose. Domeij and Floden (2006) pool three subsamples of the PSID around the

Table 15: Empirical labor-supply regressions for men, PSID data, distinction by ranges of the percentage contribution to household earnings.

	(1)	(2)	(3)	(4)
expected wage growth	0.70	0.59	0.41	0.69
	(0.41)	(0.14)	(0.10)	(0.41)
expected wage growth	-0.30			-0.12
$\times\ I(\bar{s} > 1/3)$	(0.41)			(0.42)
expected wage growth		-0.21		-0.20
$\times\ I(\bar{s} > 0.5)$		(0.12)		(0.12)
expected wage growth			-0.27	-0.25
$\times\ I(\bar{s} > 0.9)$			(0.23)	(0.23)

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

years where data on household wealth is available: 1983-1985, 1988-1990, and 1993-1995. Liquid assets are measured 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. This subsample is substantially smaller than our baseline sample and, accordingly, estimates turn out to be imprecise. Nevertheless, we use the subsample to test the hypothesis whether the coefficient on the interaction with the primary-earner dummy decreases with an increasing wealth cut-off.

From Table 16, one can see that this pattern is confirmed in the PSID data. As expected, the estimated coefficient on the interaction term, measuring the difference between estimated elasticities between primary and secondary earners, is the smaller (in absolute terms), the wealthier is the considered group of households. 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, in line with our model.

A.7.4 Labor-supply elasticities of individuals with high and low earnings

Some individuals move up and down the income distribution and hence switch back and forth between the two earnings groups considered in Section 6.3. These individuals do not necessarily change their own true Frisch elasticity but still may change the *average* true Frisch elasticity in the groups. In order to limit the influence of individuals who move up and down the earnings distribution, we consider a specification where we exclude these individuals.

Table 16: Empirical labor-supply regressions for men, PSID data, distinction by earner status and household liquid wealth.

	(1)	(2)	(3)	(4)
	all men	assets > 0	assets $> 0.5 \cdot y^m$	assets $> y^m$
expected wage growth	0.60 (0.20)	0.63 (0.23)	0.57 (0.23)	0.52 (0.24)
expected wage growth $\times$ primary earner	-0.35 (0.19)	-0.31 (0.22)	-0.25 (0.22)	-0.19 (0.23)
time effects	yes	yes	yes	yes
observations	9,780	6,886	6,023	5,511

NOTE.–Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Individuals identified as primary earners if the mean realized wage rate in the sample $\overline{w}_{ij}$ exceeds the mean realized wage rate of the spouse $\overline{w}_{-ij}$. y^m is an average full-time male monthly income, y^m = \$2014 (1983 dollars). Assets: 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. Standard errors in parentheses.

While this reduces the sample size by more than 20%, we find that the results remain very similar.

Table 17: Empirical labor-supply regressions for men, PSID data, by earnings group, excluding individuals moving up and down the earnings distribution.

	(1)	(2)	(3)	(4)
	top 25% earnings		bottom 75% earnings	
expected	0.17	0.60	0.54	0.87
wage growth	(0.19)	(0.47)	(0.14)	(0.27)
expected wage growth		-0.63		-0.55
× earnings contribution (%)		(0.64)		(0.39)
time effects	yes	yes	yes	yes
observations	2,514	2,514	8,784	8,784

NOTE.–Estimation results for men. Dependent variable is hours growth $\Delta \ln n_{ijt+1}$ of individual i in household j in period $t+1$. Constant included but not shown. Usual earnings contribution is the average percentage contribution of individual i to labor earnings of household j in the sample. Earnings groups defined using the position of current labor earnings in the distribution of individual labor earnings in the year of observation. Here, we exclude individuals who change between earnings groups more than once. Standard errors in parentheses.

Chapter III

Risk-sharing Implications of Limited Commitment in Dual-Earner Couples - A Quantitative Analysis

1 Introduction

A large literature quantifies insurance possibilities of individuals under market imperfections (see for example Blundell et al. 2008, Kaplan and Violante 2010, and Heathcote et al. 2014). An important channel of insurance against idiosyncratic income risk is intra-household risk sharing, i.e., to pool income risk with other members of the family (see Ortigueira and Siassi 2013, Blundell et al. 2016, and Blundell et al. 2018). For example, in the model of Ortigueira and Siassi (2013), family insurance can cushion an individual's consumption drop upon unemployment by up to 75%. The existing literature typically uses models of the family that assume full commitment between family members. Full commitment means that, at the time of household formation, spouses choose a contingency plan for their future lives and can make binding promises to stick to this plan. By contrast, in the family-economics literature a number of recent papers have taken a new direction in modeling decision-making by considering limited commitment in household decision making, i.e., that individuals are free to leave the household whenever this is preferable to them independent of promises made in the past. Chiappori and Mazzocco (2017) provide a first survey of the limited-commitment literature focusing on the general theoretical modeling strategy and applications.

Limited commitment between household members reduces the possibility to share risk within the family. This point has been discussed by Voena (2015) and Ábrahám and Laczó (2018). In general, risk-averse spouses would like to share their income risk and smooth consumption as much as possible. When spouses are able to commit to plans, they can achieve perfect intra-household risk sharing. A negative shock to one spouse's income is then equally shared by both spouses, i.e., consumption drops for both spouses similarly. In turn, after a positive shock to one spouse's income, this spouse has to share the additional income with the other spouse. While it is initially rational for either spouse to promise to share positive income shocks with the partner in return for the promise of insurance against negative shocks, incentives to keep the promises are reduced once shocks have materialized. For example, after being promoted, an individual may not be willing to share the increased salary with the partner because the partner's promise of support in case the promotion had not taken place is now essentially worthless. Hence, risk sharing is incomplete when commitment is limited which increases the volatility of consumption for individuals and induces a welfare loss.

In this paper, we quantify the impact of limited commitment in household decision making on the degree of intra-household risk sharing. We develop a quantitative incomplete-markets model with dual-earner households whose individual members cannot fully commit to promises made to each other. A particular contribution of this chapter is that we calibrate the model to obtain a realistic degree of limited commitment. This is important as our aim is to make quantitative predictions about the risk-sharing implications of limited commitment. Specifically, we use the evidence provided recently by Lise and Yamada (2019) who use Japanese data on individual consumption shares in a household and document that a fifth of the variance of these consumption shares occurs *within* households over time. This observation is informative about the degree of limited commitment within households as perfect risk sharing under full commitment would imply this within-household variance to be zero. Matching this target, our model implies that limited commitment increases individual consumption variances on average by about 25% compared to a reference model with full commitment. The aggregate Kaplan and Violante (2010) risk-sharing coefficient is reduced by about one third through spouses' inability to fully commit.

We use our model to assess the welfare consequences of foregone intra-household insurance. The aggregate welfare loss from incomplete intra-household risk sharing due to limited commitment amounts to an equivalent of 0.6% of life-time consumption. This number can be compared to the estimated welfare costs of other frictions or the costs of business cycles. Guner et al. (2011) estimate a consumption equivalent of 0.2% for the welfare effect of joint taxation of married couples. In their incomplete-markets set-ups, Krebs (2007) estimates welfare costs of business cycles of about 0.5% and the number reported by Krusell et al. (2009) is about 1%.

We also use our model to assess which parts of the population are particularly subject to incomplete intra-household risk sharing due to limited commitment. To do so, we incorporate further exogenous heterogeneity into the model. In particular, we consider differences in time preferences, risk aversion, direct utility gains from marriage, altruism for the partner, as well as in the variance and intra-household correlation of income shocks. We find that the effect of limited commitment on the degree of intra-household risk sharing varies substantially across the population. It is particularly strong for impatient and rather risk-tolerant individuals. By contrast, individuals with high direct utility gains from marriage and those who care strongly about their partners' well-being can achieve high degrees of intra-household risk sharing even in presence of limited commitment. Further, when the correlation in spouses' incomes is low or when couples face little income risk, individuals are affected rather strongly by limited commitment. Quantitatively, small differences in terms of direct utility gains from marriage and the degree of altruism around their respective mean values are predicted to lead to rather large differences in the extent to which intra-household risk sharing is worsened by limited commitment between spouses.

Our results have important policy implications. Given that imperfections in intra-household risk sharing due to limited commitment are substantial, policy makers should take them into account when designing policies to protect people against income risk. First, our results show that individuals face substantially higher consumption risk than is suggested by a pure household perspective that implicitly assumes full commitment. This raises demand for other forms of insurance, including insurance through public programs such as unemployment insurance or progressive taxation. Second, our results also have implications for the design of such programs. Limited commitment in the household implies that it should be the *individual* suffering from a negative income shock who is entitled to a compensation transfer payment, and not the *household* in which the individual lives. This ensures that the bulk of the transfer payment actually reaches the targeted individual and is not consumed by the household member with the largest bargaining power, i.e., in such situations most often the other spouse. Our results further show that there can be individuals suffering from low consumption even in households that in principle have sufficient means. Accordingly, our analysis suggest that means testing of transfer programs should at least in part be performed at the individual level and not at the household level.

Our cross-sectional results show that households most likely to become reliant on social transfer programs are particularly affected by risk-sharing inefficiencies due to limited commitment. It is well documented that the risk of poverty is higher for the impatient (e.g., Lawrance 1991) and for those more prone to divorce (e.g., McManus and DiPrete 2001, Peterson 1996, Holden and Smock 1991). Our results show that these characteristics are also associated with higher risk-sharing inefficiencies due to limited commitment.

The remainder of this chapter is organized as follows. Section 2 reviews the literature on limited commitment in the household. Section 3 presents our model. Section 4 describes the numerical solution technique. Section 5 provides results regarding the aggregate risk-sharing effects of limited commitment. Section 6 studies cross-sectional heterogeneity in these effects. Section 7 concludes.

2 Literature

The majority of empirical studies on household decisions have considered collective models under full commitment, see Browning et al. (2014) and Chiappori and Mazzocco (2017) for surveys of the empirical literature. In this section, we survey the literature on limited commitment. The theoretical basis of limited-commitment models was established in the field of development economics, see Ligon et al. (2002). We give an overview of these studies in Section 2.1. In Section 2.2, we review studies that have provided empirical evidence for commitment problems between spouses. In Section 2.3, we give an overview of applications of limited-commitment models in family economics.

While we focus on tests and applications of limited commitment in family economics, the theoretical basis of limited commitment has proven useful in other fields of research. For instance, similar theoretical frameworks have been used in the context of labor contracts, see Thomas and Worrall (1988) and Oikonomou (2018); for the analysis of informal insurance in village economies, see Kimball (1988), Coate and Ravallion (1993), Kocherlakota (1996), Ligon et al. (2000,2002), Foster and Rosenzweig (2001), Dubois et al. (2008), Laczó (2014, 2015), and Morten (2019); for analyzing risk-sharing agreements between rather than within households, see Attanasio and Rıos-Rull (2000) and Krueger and Perri (2006); in international economics, see Kehoe and Perri (2002); in political economy, see Dixit et al. (2000), Aguiar and Amador (2011), and Acemoglu et al. (2011); and for optimal taxation, see Park (2014).

2.1 Theoretical basis of limited commitment models

The theoretical basis of limited-commitment models stems from the field of development economics, in particular, Kimball (1988), Coate and Ravallion (1993), Kocherlakota (1996) and Ligon et al. (2002).[1] These early limited-commitment models consider risk sharing between inhabitants of village economies.

Kocherlakota (1996) provides first theoretical implications of limited commitment, e.g.,

[1] Kimball (1988) and Coate and Ravallion (1993) were the first to consider commitment issues between households in village economies. These early studies analyze the self-enforcing of insurance agreements where households transfer resources to other households such that marginal utilities of consumption are equalized across households. When this allocation rule of resources violates at least one participation constraint, the transfers are adjusted such that all participation constraints are fulfilled. However, in contrast to subsequent studies, transfers in the current period do not depend on previous period transfers. Thus, whenever possible, marginal utilities of consumption of households are equal independent of previously made transfers.

that commitment problems tend to decrease with the patience of agents.[2] He also connects limited-commitment models to non-cooperative dynamic games by showing that the prevalent solutions of limited-commitment models are subgame-perfect.

A full characterization of the limited-commitment solution was provided by Ligon et al. (2002). This paper has become the theoretical basis for almost all applications employing models of limited commitment and in Sections 3 and 4, we will provide a more detailed discussion how their results can be used to solve quantitative limited-commitment models. Ligon et al. (2002) analyze risk-sharing in a limited-commitment environment where agents have the opportunity to self-insure through private savings. They highlight that the opportunity to save can increase or decrease welfare depending on the particular savings technology. Additionally, they provide a first discussion about a numerical analysis of a limited-commitment model. Ligon et al. (2002) test implications of their model by investigating risk-sharing behavior in three Indian villages. While their model can explain the reaction of individual consumption to individual income, it cannot explain the consumption distribution within the villages.

In a different paper, Ligon (2002) connects the limited-commitment model of Ligon et al. (2002) to more traditional bargaining procedures like the Nash solution. He shows that the limited-commitment model fulfills two classical axioms of the Nash (1950) bargaining solution: invariance (INV) and independence of irrelevant alternatives (IIA). However, as discussed in the introduction to this chapter, the resulting allocation in a limited-commitment model is *ex-ante* inefficient. Thus, the Pareto efficiency axiom of Nash is not fulfilled. Ligon (2002) introduces two additional axioms: individual rationality (IR) and constrained Pareto optimality (CPO). The latter axiom states that an allocation must be *ex-post* efficient. Ligon (2002) shows that the limited-commitment model fulfills the four axioms, INV, IIA, IR, and CPO. Thus, the solution to the limited-commitment model can be interpreted as the solution to a dynamic bargaining problem. This result establishes the interpretation of re-bargaining between spouses when the Pareto weight changes, see Mazzocco (2007). Furthermore, Ligon (2002) shows that the solution of a limited-commitment model is Pareto superior to a sequence of static Nash bargaining solutions.

Ábrahám and Laczó (2018) use a limited-commitment model to analyze the incentives of spouses to save hiddenly.[3] In the model, agents face perfectly negatively correlated income shocks, aggregate household income is fixed, and all joint asset holdings are lost upon divorce.

[2] We replicate this result in the analysis in Section 6.1.

[3] They analyze a general model with two agents who share idiosyncratic risks so that their results can be applied to the case of couple households.

Ábrahám and Laczó (2018) show that agents have an incentive to save more under limited commitment than under full commitment which is caused by the assumption that all joint assets are lost upon divorce. Thus, welfare losses from divorce increase in asset holdings and outside options become ceteris paribus less attractive the higher are joint asset holdings. In the second part of their paper, Ábrahám and Laczó (2018) analyze a model where spouses can additionally save in hidden private asset holdings. While the household-public asset holdings get lost upon divorce, hidden private savings remain even upon divorce. Ábrahám and Laczó (2018) show that hidden private savings do not occur in equilibrium. The reason is that the household planner internalizes the incentives to save hiddenly. Thus, by saving in joint asset holdings, any incentives to save in hidden private asset holdings are eliminated. However, due to the opportunity to save in private asset holdings also in autarky, consumption is ceteris paribus more dispersed since commitment issues increase due to ceteris paribus improved outside options. Whether the opportunity to save is welfare increasing or not depends on the return on savings. If the interest rate is sufficiently high, welfare is higher than in a model version without any saving opportunities.

2.2 Empirical evidence on commitment problems

Browning et al. (2014) provide a detailed discussion of issues in empirical analyses of collective household models. In general, the empirical analysis of commitment issues requires panel data on individual income and, ideally, consumption of spouses in a household. With such data, one can test whether intra-household risk sharing is perfect as predicted by the full commitment model. Perfect intra-household risk sharing implies that the origin of a stochastic income shock does not matter for the intra-household allocation of consumption. In a world of full commitment, perfect intra-household risk sharing must hold due to the once and for all decision on a particular intra-household allocation rule. By contrast, intra-household risk sharing is imperfect in a world of limited commitment. Empirical studies use expenditures on different consumption goods to test whether the use of household income depends on individual income shocks of the spouses. However, most datasets that provide consumption data at all only report consumption at the household level and not at the individual level. To handle these data issues, econometricians frequently use a priori assumptions which goods are consumed relatively more by men and women, respectively. To test for imperfect risk sharing, one then tests hypotheses such as, for example, an increase of the income of the male may increase expenditures on football tickets while an increase of the income of the female may increase expenditures on theater tickets. If this is the case, one concludes that spouses cannot fully commit to a particular allocation rule.

Dercon and Krishnan (2000) investigate risk-sharing behavior across and within households in villages in Ethiopia with respect to health risk. It is assumed that households share risk to smooth the nutritional status across and within households over time and state of natures of the world. They use data on food and nutrients consumption to proxy for the nutritional status and find that risk sharing is rejected for poor households within the villages. Furthermore, their empirical analysis shows that idiosyncratic health shocks are not shared perfectly within households in some regions of Ethiopia. Especially health shocks of women are less perfectly shared in couple households than mens' ones.

Duflo and Udry (2004) use panel data from Côte d'Ivoire to test the perfect risk-sharing prediction of the full commitment model. They reject perfect intra-household risk sharing by showing that expenditures for different consumption goods depend on individual income shocks.

Mazzocco (2007) uses U.S. data to estimate Euler equations of couple households, one for goods that can be considered private consumption of the individual spouses such as food or clothes and one for consumption of household-public goods such as housing or housekeeping services. To identify commitment issues, spouses' incomes are included in the Euler equation as distribution factors. In both model types, full and limited commitment, incomes should enter the Euler equations as a proxy for the probability distribution of incomes at marriage formation. In the limited-commitment model, incomes should enter the Euler equation additionally as a proxy for a change in the Pareto weights, which captures time-variation in relative intra-household decision power. Mazzocco (2007) develops a test of commitment by using that, if spouses can fully commit, the distribution factors (the incomes of spouses) should enter the Euler equation only as interaction terms with consumption. In contrast, if spouses cannot commit, the distribution factors should additionally enter the Euler equation directly. Mazzocco (2007) strongly rejects the full commitment model while the no-commitment model cannot be rejected.

In an experimental study using data from Kenya, Robinson (2012) shows that an individual spouse increases her consumption by more if she is hit by a positive shock to her own income in comparison to a situation where the other spouse receives the positive income shock. Thus, perfect intra-household risk sharing is rejected, which Robinson (2012) also confirms using non-experimental data.

Blau and Goodstein (2016) investigate potential commitment issues of older couples using U.S. data on individual inheritances. If spouses can fully commit, the labor supply of a particular spouse should not depend on the distribution of inheritances, i.e., on who obtains an inheritance. The empirical analysis of Blau and Goodstein (2016) rejects perfect intra-household risk sharing and thereby the ability of spouses to fully commit to a particular allocation.

Cesarini et al. (2017) use panel data from Sweden to analyze the impact of lottery wins on labor supply, both in single and couple households. They show that, in couple households, the labor supply of the person who wins the lottery is reduced more strongly than the labor supply of the spouse which contradicts full commitment. Cesarini et al. (2017) suggest possible explanations for their result. For example, spouses may decide cooperatively but cannot commit to an allocation rule once and for all. A second possible explanation is that the behavior of couples can be characterized by the "separate spheres" model of Lundberg and Pollak (1993) where the outside option of the bargaining process is non-cooperative behavior within the couple. Cesarini et al. (2017) favor this "separate spheres" approach. However, even if the bargaining procedure was of this type, the resulting allocation would be determined once and for all at the household formation step if spouses can fully commit. Thus, a separate spheres" model could only explain the empirical results of Cesarini et al. (2017) if bargaining is repeated in every period, and, repeated bargaining effectively corresponds to an (alternative) limited-commitment model.

Lise and Yamada (2019) estimate a full household-decision model allowing for potential time-variation in the bargaining weights of spouses. They use a panel data set for Japan which includes individual consumption data as well as individual data on time spent on leisure, hours worked, and homework. In contrast to the majority of limited-commitment models, which solve for the endogenous evolution of the Pareto weights, Lise and Yamada (2019) assume a particular functional form for the Pareto weight, where full commitment is a nested case.[4] The parameters of this functional form are estimated among other parameters. The results show that the reactions of hours and consumption are not in line with a full commitment model. In particular, the reaction of hours to unpredictable changes in the relative wage rate is too small, while the reaction of consumption is too strong for a full-commitment model to rationalize them. By contrast, these results are well in line with a limited-commitment model since an unpredictable change in the relative wage rate may increase the Pareto weight of this spouse. In consequence, this spouse consumes relatively more market goods and leisure time. Importantly, the results of Lise and Yamada (2019) can help calibrating a limited-commitment model with respect to a target value for the severity of commitment problems

[4]Strictly speaking, Lise and Yamada (2019) consider a no-commitment model where the Pareto weight changes whenever the state of nature of the world changes.

between spouses. Lise and Yamada (2019) show that initial differences in Pareto weights account for roughly 80% of the total variance of Pareto weights. The remaining 20% are due to changes in the Pareto weight over time. In a full-commitment model, these latter changes do not occur and can hence be seen as a measure for the strength of problems due to limited commitment. In our quantitative model evaluations in Section 5 and Section 6, we will compare this statistic to the Lise and Yamada (2019) evidence.

While the studies discussed so far test for perfect intra-household risk sharing, Chiappori et al. (2019) use the labor-supply behavior of couples to test whether couples can fully commit or not. This study is among the first to test all three collective models, full commitment, no-commitment, and limited commitment, against each other. In a full commitment model, the Pareto weights do not depend on current or past wage shocks of the spouses. The Pareto weights are fixed once and for all implying that the allocation rule is fixed once and for all. In contrast, under limited commitment, Pareto weights do depend on both current and past wage rate shocks of the spouses, reflecting the Markovian structure of a limited-commitment model. Pareto weights change only when a spouse faces a binding participation constraint, i.e., has an incentive to leave the couple given a particular allocation rule. Thus, the allocation rule may change depending on previous and current wage shocks. In the no-commitment model, spouses re-bargain in every period. In consequence, only current period wage shocks matter since past wage shocks have no impact on the current bargaining. Using the panel structure of the PSID, Chiappori et al. (2019) find that current period labor-supply behavior of couples strongly depends on current and past wage shocks. Thus, their estimates reject both, the full commitment and the no-commitment model, and provide support for the limited-commitment model.

2.3 Applications of limited-commitment models in family economics

In general, there are two theoretical approaches which can explain the empirical evidence discussed in the last section, a collective model with limited commitment or a non-cooperative couple model. Most papers on non-cooperative household models use deterministic frameworks, e.g., Lundberg and Pollak (1994), Konrad and Lommerud (1995), and Doepke and Tertilt (2019) among others. The incomplete markets framework considered in this chapter is a dynamic stochastic framework though. There is only a small literature on non-cooperative couple models in the context of dynamic stochastic models, see Hertzberg (2019) for an example. Del Boca and Flinn (2012) test different models of couple behavior, such as efficient cooperative behavior, constrained-efficient cooperative behavior, and non-cooperative behavior against each other, using the different implications of these behaviors derived from theoretical models. Constrained-efficient cooperative behavior captures that none of the spouses should have lower well-being with cooperative behavior than with non-cooperative behavior. This situation is not ruled out in their efficient cooperative behavior model. Using PSID data, Del Boca and Flinn (2012) find that three-fourth of all couple households behave cooperatively. Furthermore, they find that ex-ante heterogeneity in Pareto weights is important for a substantial part of couples in the data in line with limited-commitment models of intra-household decision-making. Given the incomplete-markets framework of this chapter and the results of Del Boca and Flinn (2012), our literature survey focuses on collective household models with limited commitment.[5]

In the literature, there are both positive and normative papers taking into account limited commitment. Some papers study specific policy reforms while other papers study the consequence of limited commitment for the behavior of spouses without referring to a specific policy.

Divorce laws. Voena (2015), Bayot and Voena (2015), Ábrahám and Laczó (2016), and Reynoso (2019) consider divorce laws. Voena (2015) analyzes the impact of divorce laws on intra-household decision-making and the resulting allocations, including spousal welfare. She analyzes a transition from a mutual consent divorce regime to an unilateral divorce regime. A direct implication of the introduction of unilateral divorce is that it makes commitment more difficult since both spouses can file for divorce if their outside option is more attractive than their inside option. In contrast, in a regime of mutual consent, spouses can only file for divorce if both spouses agree, so that the marriage continues as long as one spouse has an incentive to maintain the marriage which makes commitment easier. Using an estimated version of her model, Voena (2015) shows that the introduction of unilateral divorce increases female decision power within the household. As a consequence, the female participation rate decreases.[6] After the introduction of unilateral divorce, couples save more on average. The estimated model implies that risk sharing, or equivalently consumption insurance, is better under a mutual consent regime than under an unilateral divorce regime. Additionally, Voena (2015) analyzes different rules of property division upon divorce: A title-based regime, a community property regime, and an equitable distribution regime. In a title-based regime,

[5]Up to our knowledge, there is only one paper in the context of family economics which uses a limited commitment approach without considering couples. Mommaerts (2016) uses a limited-commitment model to analyze intergenerational imperfect risk sharing where commitment between parents and their children is limited.

[6]Voena (2015) only considers the extensive margin of labor supply and further assumes that men always work so that a change in the divorce regime has no impact on the labor supply of men.

both spouses receive their individual asset holdings upon divorce. This requires that all assets are assigned individually to one of the spouses. In the community property regime, all assets are divided equally upon divorce. In the equitable distribution regime, courts decide over the distribution upon divorce to achieve equality in after-divorce environments. Voena (2015) obtains similar results for the different rules of property division upon divorce. In a mutual consent divorce regime, risk sharing is not affected by different property division rules. However, consumption of spouses upon divorce depends strongly on the particular rule of property division. For example, upon divorce, consumption of women decreases less in a regime of community property division.

Bayot and Voena (2015) in their study for Italy analyze the impact of prenuptial contracts on the behavior of married couples. Couples can choose between two regimes of property rights over household resources, either community property, which is the default regime, or separation of property. While all resources of the couple are joint property in the former case, all household resources are assigned to one of the spouses in the latter case. This free choice at the household formation implements a public prenuptial contract at low or no costs, respectively. Using data from Italy, they show that most couples in Italy choose the separation of property regime. Additionally, Bayot and Voena (2015) find that the community property regime is predominantly used by couples with a higher number of children and a wife who does not participate in the labor market. Using a limited-commitment model similar to Voena (2015), Bayot and Voena (2015) highlight that this type of model cannot explain the choice of community property of some couples. They find that the separation of property regime is the constrained efficient choice under limited commitment. This is caused by the fact that couples restrict their future choices to allocate assets between spouses in the case of the community property regime. Bayot and Voena (2015) extend their model by a potential non-cooperative period before divorce. Since spouses anticipate that they may not cooperate preceeding divorce, they may have an incentive to restrict asset decisions in this case to guarantee better consumption smoothing in the case of divorce. Bayot and Voena (2015) find that the modified model is line with Italian data. Furthermore, they find that a separation of property regime implies higher female labor-market participation, lower divorce rates, and higher savings.

Ábrahám and Laczó (2016) consider a limited-commitment model with the opportunity to divorce. Depending on the realization of a love shock, spouses have an incentive to divorce. Upon divorce, a constant fraction of joint asset holdings gets lost. In their numerical analysis, Ábrahám and Laczó (2016) investigate the size of this fraction and the size of the share of assets that spouses obtain who file for divorce. They highlight that optimal parameter values exist where the total loss of assets is between zero and one-hundred percent and where the asset share of the spouse, who files for divorce, is smaller than fifty percent.

Reynoso (2019) uses a limited-commitment model to analyze the impact of the introduction of unilateral divorce on marriage market outcomes and the corresponding welfare consequences. She shows that positive assortative mating in education occurs, that human capital increases on average, that less people marry and that gains from marriage decrease for most people. Spouses have an incentive to invest more in human capital since it increases their outside options. This is beneficial for the particular spouse both in the case of divorce and also within marriage when commitment is limited. However, this higher human capital stock may be inefficient from an *ex-ante* perspective. While unilateral divorce is advantageous for people who want to divorce, for example due to bad match quality shocks, it implies limited commitment between spouses as shown by Voena (2015). Unilateral divorce and with it limited commitment is disadvantageous for all (potential) couples which do not want to divorce. Which effect dominates is in general not clear. Reynoso's (2019) results suggest that for most people the second effect dominates, resulting in a smaller advantage of being married and lower marriage rates when there is unilateral divorce.

Family policies. Several papers study the effects of family policies under limited commitment. Bronson (2015) analyzes gender gaps in education, in particular different choices in college majors and the decision to attend college. Furthermore, she investigates how the evolution of skill premia impact on commitment issues within households and on the marriage market for potential couples. She then uses the estimated model to study the impact of several policies such as maternity leave, part-time work, etc. Bronson (2015) finds that part-time work entitlement policies are most efficient to narrow gender gaps in college majors since women are encouraged to choose more frequently college majors which typically have high wage penalties due to part-time work.

Foerster (2019) analyzes changes in existing alimony and child support laws in a limited-commitment model. He finds that these policies can increase ex-ante household welfare in the limited-commitment case. Alimonies and child support laws are examples for policies that have an impact on the outside options of married couples to reduce commitment issues.

Family taxation. Bronson and Mazzocco (2018) consider three tax reforms in an estimated limited-commitment model. First, a change from a joint to an individual tax basis. Second, the introduction of a deduction for the secondary earner in a system of joint taxation. Third, the introduction of child care subsidies in both a joint and an individual tax basis system.

They find that secondary earners are affected most in all three policy reforms. All three policies increase incentives for secondary earners to supply more labor. This implies that secondary earners accumulate more human capital and improve their outside options which increases their intra-household decision-power and thereby their welfare. However, Bronson and Mazzocco (2018) do not discuss the consequences for total household welfare.

Oikonomou and Siegel (2015) use a general equilibrium limited-commitment model to analyze a tax reform where a linear capital income tax is replaced by a linear labor income tax. There are no match quality shocks and hence all commitment issues are driven by idiosyncratic wage uncertainty. In the Oikonomou and Siegel (2015) framework, wealth reduces commitment problems because spouses want to avoid loosing their wealth upon divorce.[7] Hence, incentivizing households to save can help them commit to more efficient intra-household allocations and thereby raise welfare. Oikonomou and Siegel (2015) show in their quantitative analysis that the tax reform replacing capital income taxes by labor income taxes increases aggregate savings, aggregate hours worked, and aggregate welfare. In welfare terms, single females profit most from the tax reform while single males even suffer from the tax reform. Using partial equilibrium analyses, Oikonomou and Siegel (2015) show that the severity of commitment issues is reduced by the reduction of capital taxes. Higher labor income taxes strengthen commitment issues since lower household income increases the incentives to leave the household. However, all of these effects are quantitatively relatively small.[8]

Obermeier (2019) investigates the impacts of progressive labor income taxation on the intra-household allocation and intra-household inequality under limited commitment. He considers a progressive labor income tax system where spouses are taxed separately. He shows that intra-household consumption inequality captures about one fourth of total consumption variance in his calibrated economy. Then, he performs a tax experiment by increasing the progressivity of the tax system. In line with previous literature, Obermeier (2019) finds that the increased progressivity decreases the total consumption variance in the economy. He shows that also intra-household consumption inequality decreases substantially after the tax reform. A full commitment model could not capture this effect. Furthermore, Obermeier (2019) shows that changes in marriages and divorces do not have a significant impact on intra-household inequality.

[7] In a more general framework, this is a quantitative question, see Section 5.

[8] Already before the reform, commitment issues are relatively unimportant, with the welfare loss from limited commitment amounting to only about 0.06% in consumption-equivalence terms.

Labor supply. Mazzocco et al. (2013, 2014), Low et al. (2018), Theloudis (2018), Potoms (2019), and Turon (2019) study the impact of limited commitment on (female) labor supply. Mazzocco et al. (2013) use PSID data to show that couple household behavior such as labor supply, home production, savings and marital decisions are linked. Their estimated limited-commitment life-cycle model highlights the importance of a joint consideration when analyzing policies. They use the estimated model to analyze the impacts of the Earned Income Tax Credit (EITC) and other subsidy programs on household behavior and welfare. In particular, Mazzocco et al. (2013) show that an elimination of EITC would increase labor supply of singles, both men and women, while the impact on labor supply of spouses is negligible. The elimination of EITC would reduce individual welfare. Mazzocco et al. (2013) consider a replacement of EITC with another subsidy system as a second policy experiment. The alternative subsidy system is build such that labor-supply disincentives should be smaller than in the EITC. They find that labor-force participation increases sharply after the replacement while average hours worked conditional on working decrease. However, Mazzocco et al. (2013) do not discuss how much of the results of these policy experiments depend on commitment issues between spouses and through which channels.

Mazzocco et al. (2014) use a limited-commitment model to analyze the labor-supply behavior of spouses. While many changes in labor supply of married people in recent decades can be exlained by changes in the intra-household distribution of human capital and the presence of children, there remains a significant unexplained part. Mazzocco et al. (2014) show that a limited-commitment model with uncertainty about match quality can explain parts of this unexplained part of the labor-supply transition.

Low et al. (2018) use their estimated limited-commitment model to analyze the impacts of the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) reform on intra-household allocations. PRWORA limits the time for which a person can receive welfare payments during her lifetime. Thus, the reform of the welfare system increases the insurance advantage of being married. Low et al. (2018) show the reform has the strongest impact at the time of marriage while the reform has only a small impact on commitment issues during marriage in line with an increased gain from marriage. Further, they show that the reform reduces divorces and increases marriages due to higher gains from marriage in the form of increased insurance gains.

Theloudis (2018) uses an estimated limited-commitment model and shows that female bargaining power has increased over time due to a narrowing of the gender wage gap.[9] The

[9] Strictly speaking, Theloudis (2018) considers a no-commitment model where the Pareto weights change whenever the state of the world changes.

model can explain a substitution of female time through male time in homework observed by several studies, see Aguiar and Hurst (2016) for empirical evidence. Although Theloudis (2018) includes limited commitment between spouses in his model, the impact of commitment issues on the results is not discussed.

Potoms (2019) uses an estimated limited-commitment model to analyze the impact of endogenous borrowing constraints on labor supply of secondary earners. Borrowing constraints become endogenous since housing is used as collateral for debt. He uses the primary loan-to-income (p-LTI) ratio to discriminate between couple households. Couples with a high p-LTI have a higher probability of binding borrowing constraints, especially if the primary earner faces a negative income shock. In a counterfactual analysis, Potoms (2019) investigates the impact of a loosening of the endogenous borrowing constraint. Home ownership then increases and the new home owners use labor supply, and especially the labor supply of secondary earners, to finance their house. Thus, average labor supply of secondary earners increases. Potoms (2019) shows that especially households with a low p-LTI ratio are affected by the policy. The bargaining power of the household head is relatively high in these households implying a larger increase of labor supply of the secondary earner under limited commitment than under full commitment.

Turon (2019) analyzes childcare in a limited-commitment model. Using an estimated version of her model, she shows that standard childcare policies, such as subsidies or free public childcare, have almost no impact on a substantial fraction of households with young children. She shows that almost the same impact could be accomplished by changing the parental-childcare-taste distribution due to higher quality of market childcare. Additionally, she analyzes the impact of part-time work and thereby depreciating human capital on future commitment issues. While the decision to work part time has no impact on future allocations, Turon (2019) finds that these additional costs of home-produced childcare cannot be internalized by the households in the presence of limited commitment. Thus, labor-supply decisions of couples become even more distorted and inefficient.

Marriage market. Gemici and Laufer (2014), Devereux and Turner (2016), Gierlinger and Laczó (2017), and Shephard (2019) study the impact of limited commitment between spouses on the marriage market. Gemici and Laufer (2014) analyze the behavior of cohabitating couples. Since cohabitators do not have to file for divorce if they want to break up, commitment issues are expected to be stronger than in a married couple. Using PSID data, Gemici and Laufer (2014) find empirical evidence that typical couple advantages such as risk sharing and specialization are observed less for cohabitating couples. They use their estimated limited-commitment model to analyze a counterfactual increase of divorce costs. Higher divorce costs cause lower marriage rates and a higher number of cohabitating couples. Since divorce becomes less attractive, only the couples with the highest gains marry resulting in less severe commitment issues of married couples and thereby a higher degree of specialization of married couples. The welfare consequence are highest for couples who would be married under the baseline divorce costs while they cohabitate under the counterfactual divorce costs.

Devereux and Turner (2016) investigate the impact of the opportunity to search for a new partner during marriage (on-the-marriage partner search) on commitment issues and welfare of spouses. In general, commitment issues become more severe due to on-the-marriage search since outside options become more attractive. This induces welfare losses. However, Devereux and Turner (2016) highlight an additional welfare improving channel. Spouses within a bad marriage, a marriage with a bad match quality shock, have better opportunities in welfare terms to leave the couple. Thus, the aggregate welfare effects of on-the-marriage search are ambiguous. Using an estimated version of their model, Devereux and Turner (2016) find that a limited-commitment model with on-the-marriage search explains empirical patterns of re-marriage better than a limited-commitment model without on-the-marriage search.

Gierlinger and Laczó (2017) analyze the effects of limited commitment on assortative matching in a frictionless matching market. In their model, the gain from marriage is risk sharing. Spouses are heterogeneous with respect to risk aversion. The model replicates the result known from the literature that, if spouses can fully commit in such an environment, the optimal result is perfect negative assortative matching. Gierlinger and Laczó (2017) begin with considering static risk-sharing agreements. In contrast to the full commitment case, the correlation of incomes matters in the limited-commitment model. They show that, if incomes are negatively correlated, all matchings are perfectly positive assortative when commitment issues are sufficiently severe. If incomes are non-negatively correlated, the assortativeness of matchings is always negative. These results can be explained by differential risk-aversion among spouses. If incomes are negatively correlated, both spouses can reduce their consumption risk. Spouses with higher risk aversion commit to higher transfers in exchange for more insurance. If commitment issues are sufficiently severe, caused by a discount factor which is sufficiently low, this positive effect of equal risk-aversion dominates the effect of more efficient risk sharing among spouses with unequal risk-aversion. If incomes are positively correlated, the consumption risk of one spouse can only be reduced by increasing the consumption risk of the other spouse. Thus, stronger risk aversion increases commitment issues in this case.

In consequence, the amount of insurance increases in the difference of risk aversion among spouses, making a negative assortative matching more attractive. In a second step of the analysis, Gierlinger and Laczó (2017) confirm that positive assortative matching also occurs in a dynamic environment.

Shephard (2019) considers a limited-commitment model with search frictions in the marriage market allowing for age differences within the couple. The model can replicate the bivariate age distribution of married couples and several life-cycle patterns of married couples. The model is used to quantify implications of the narrowing of the gender wage gap: an increase in female employment, a decrease in male employment, delayed marriages, and a decrease in the marital age gap. Shephard (2019) finds that the narrowed gender wage gap accounts for roughly one-third of the decreased marital age gap. While Shephard (2019) includes limited commitment between spouses in his model, the impact of commitment issues on results is not discussed.

Household finance. Addoum et al. (2016) analyze the impact of limited commitment between spouses on the structure of household portfolios. They assume that women are more risk averse than men. In their model, changes in Pareto weights imply changes in the risk aversion of the couple, implying adjustments in the risk composition of the household asset portfolio. Since women are more risk averse, they prefer less risky assets than men. Thus, if a positive income shock increases the bargaining power of the wife, the portfolio of the couple becomes less risky.

3 A stochastic endowment economy

In this section, we present a simple limited-commitment model in the context of family economics. It is an incomplete markets model where spouses face idiosyncratic income risks and have no opportunity to save. The only way to insure against bad income realizations is to share risk within the couple. But this is done in an imperfect way in presence of limited commitment. In the full commitment case, there is no non-trivial decision within the household after the household formation. In the limited-commitment model, rebargaining upon binding participation constraints is the sole "decision".

3.1 Model set-up

The economy is populated by couple households consisting of two spouses, a male and a female, m and f, respectively, who are assumed to decide cooperatively. Couples discount future utility by the discount factor β. Spouses receive utility from private individual consumption, c_{it}, with $i = m, f$. The individual concave period felicity function is denoted by $u_i(c_{it})$. Each spouse receives a constant period utility gain from marriage, denoted by Ψ. This constant utility gain captures all gains from marriage not explicitly modeled such as home production, children, etc. Thus, expected lifetime utility of spouse i at household formation is

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \cdot [u_i(c_{it}) + \Psi], \text{ for } i = m, f.$$

Spouses face idiosyncratic income, y_{it}, which follows a Markov process. $\mathbf{y_t}$ is the vector of joint income realizations, $\mathbf{y_t} = (y_{mt}, y_{ft})$. At the household formation step, potential spouses decide whether to enter the household. They do so only when their expected lifetime utility within the couple is at least as high as their expected lifetime utility as a single. Thus, to state the maximization problem of the couple, we have to define the expected lifetime utility of a single (=bachelor) household.

3.2 Bachelor households

The maximization problem of a bachelor household with gender i in recursive formulation is

$$S_i(y_i) = \max_{c_i} \; u_i(c_i) + \beta \cdot \mathbb{E}\left[S_i(y_i')\right], \tag{54}$$

where a prime ($'$) denotes next period values, subject to the period budget constraint

$$c_i \leq y_i, \tag{55}$$

where $\mathbb{E}$ denotes the expectation operator which includes all information in the particular period. In particular, expectations about next period incomes are conditioned on current period income realizations. As singles can only consume their period income, the optimal consumption policy function is $\mathcal{C}_i^S(y_i) = y_i$.

The lifetime utility of the bachelor household results from the Bellman equation (54) using the optimal consumption policy:

$$S_i(y_i) = u_i\left(\mathcal{C}_i^S(y_i)\right) + \beta \cdot \mathbb{E}\left[S_i\left(y_i'\right)\right]. \tag{56}$$

These lifetime utilities are used as outside options in the participation constraints of couple households.

3.3 Couple households

Household preferences are characterized following the collective approach of Chiappori (1988). In this approach, household preferences can be stated as a weighted sum of individual preferences

$$v(c_m, c_f, \mu_m, \mu_f) = \mu_m \cdot u_m(c_m) + \mu_f \cdot u_f(c_f) + (\mu_m + \mu_f) \cdot \Psi,$$

where the weights, μ_m and μ_f, depend on prices, incomes, and distribution factors.

The collective approach is only feasible if intra-household decisions are efficient. Limited commitment implies that the intra-household allocation is *ex-ante* inefficient (see Ligon et al. 2002). However, Ligon et al. 2002 show that the resulting allocations are *ex-post* efficient. Thus, a weighted sum approach is feasible to characterize preferences of couple households even under limited commitment. Using the results of Marcet and Marimon (2019), one can show that the weighted sum approach can be used in a limited-commitment model when the Pareto weights become additional state variables that are allowed to change over time.[1] Before we consider the maximization problem of couples under full and limited commitment in detail, we discuss commitment issues between spouses in this environment.

Commitment issues. Full and limited-commitment models differ with respect to the occurrence of the participation constraint. If spouses can fully commit, the participation con-

[1] The derivation can be found in Appendix A.1.

straints of spouses need only be fulfilled at the household formation step. Thus, if both spouses have an incentive to enter the couple household at the household formation step, they decide how to allocate consumption *once and for all*. In deciding over the consumption allocation, spouses determine their particular weights in the preferences of the couple household. In terms of the Pareto weights, full commitment implies that the weights do not change over time.

In contrast, if spouses cannot fully commit, the participation constraints must be fulfilled in every period and after any history of income shocks. Thus, there must be an incentive for both spouses to remain in the couple household in any situation. Ligon et al. (2002) characterize the behavior in such an environment.[2] If one spouse has an incentive to leave the couple household, i.e. a binding participation constraint, the couple changes the intra-couple allocation such that both spouses again have an incentive to remain within the couple.[3] In particular, Ligon et al. (2002) show that the allocation is changed as little as possible such that the participation constraint of the spouse with an incentive to leave holds with equality.

This behavior of couples can be re-interpreted in terms of the Pareto weights. In the limited-commitment model, the Pareto weights become additional time-varying state variables. If the participation constraint of one spouse is binding, the couple changes the Pareto weights so that the participation constraint of the particular spouse holds with equality. It is optimal for the household to change the Pareto weight μ as little as possible to fulfill the participation constraints. Marcet and Marimon (2019) show that the current period Pareto weight of a spouse is her previous period Pareto weight plus the current period Kuhn-Tucker multiplier on her participation constraint.[4] Thus, if her participation constraint is not binding in a particular period, her Pareto weight does not change. In general, however, Pareto weights may grow over time and are unbounded in an infinite-horizon model. Thus, the Pareto weights do not add up to one in the general case. The implication is that not all periods have the same weight for the households. Whenever a participation constraint is binding, this period is more important for households and first-order gains could be achieved by transferring resources in this period such that the participation would not be binding. Intuitively, if spouses knew that a participation constraint was binding next year, they have an incentive to transfer resources from this or future periods to the next period to avoid this binding participation constraint. Therefore, it is important for *intertemporal* decisions to consider the

[2]Ligon et al. (2002) consider risk-sharing agreements in village economies. Their approach has been applied to couple households by several studies such as Mazzocco (2007) and Voena (2015).

[3]If such rearrangement of the intra-couple allocation is not possible, spouses divorce. In the calibrated models studies in Section 5, divorces do not occur in equilibrium.

[4]A derivation of this result can be found in Appendix A.1.

more general case of Pareto weights which do not add up to one. However, when there are no intertemporal decisions, couples have no opportunity to transfer resources in these periods. Thus, when there are no intertemporal decisions, the Marcet and Marimon (2019) approach yields the same allocations compared to an approach where time-varying Pareto weights of spouses are assumed to sum up to one. When only the *relative* Pareto weights matter for the final allocations, one can assume without biasing the results that Pareto weights sum up to one. This approach is not feasible in the case of intertemporal decisions, see Chapter V for a detailed discussion. In models without intertemporal decisions, we denote the male Pareto weight by μ and the female Pareto weight by $(1-\mu)$ so that the household target function is given by

$$v\left(c_m, c_f, \mu\right) = \mu \cdot u_m\left(c_m\right) + (1-\mu) \cdot u_f\left(c_f\right) + \Psi.$$

Under limited commitment, the Pareto weight is a function of the current period stochastic state, here the income realizations, and the previous period Pareto weight. In this set-up, the influence of commitment issues on the Pareto weight implies the following updating rule for the Pareto weight, see Ligon et al. (2002):

$$\begin{aligned} \mu' &= \operatorname{argmin}_{\tilde{\mu}} \left|\tilde{\mu} - \mu\right| \text{such that} \\ u_m\left(c_m\right) + \Psi + \beta \cdot \mathbb{E}\left[V_m\left(\mathbf{y}', \tilde{\mu}\right)\right] &\geq S_m\left(y_m\right) \\ u_f\left(c_f\right) + \Psi + \beta \cdot \mathbb{E}\left[V_f\left(\mathbf{y}', \tilde{\mu}\right)\right] &\geq S_f\left(y_f\right), \end{aligned} \tag{57}$$

where V_m and V_f denote the lifetime utilities of married men and women and S_m and S_f denote the outside options.

Ligon et al. (2002) show that the solution of a limited-commitment problem can be characterized in an equivalent way by a set of optimal intervals for the Pareto weight. These intervals depend on the realizations of state variables, here, income shocks. Recall that μ denotes the Pareto weight of the *husband.* The upper bounds of the intervals at a particular income combination $\mathbf{y}$, $\overline{\mu}'(\mathbf{y})$, are defined such that the wife's participation constraint holds with equality

$$u_f\left(c_f\right) + \Psi + \beta \cdot \mathbb{E}\left[V_f\left(\mathbf{y}', \overline{\mu}'(\mathbf{y})\right)\right] = S_f\left(y_f\right).$$

The male Pareto weight cannot become larger than $\overline{\mu}'(\mathbf{y})$ as the wife's share would then become too small to make her remain in the household. The lower bound, $\underline{\mu}'(\mathbf{y})$, is defined analogously for the husband

$$u_m(c_m) + \Psi + \beta \cdot \mathbb{E}\left[V_m\left(\mathbf{y}', \underline{\mu}'(\mathbf{y})\right)\right] = S_m(y_m).$$

In consequence, the updating rule (57) can be re-written as

$$\mu' = \begin{cases} \underline{\mu}'(\mathbf{y}), & \text{if } \mu < \underline{\mu}'(\mathbf{y}), \\ \mu, & \text{if } \mu \in \left[\underline{\mu}'(\mathbf{y}), \overline{\mu}'(\mathbf{y})\right], \\ \overline{\mu}'(\mathbf{y}), & \text{if } \mu > \overline{\mu}'(\mathbf{y}). \end{cases} \tag{58}$$

In the numerical solution of limited-commitment models, we use these intervals to update the Pareto weights.

Full-commitment couple. Chiappori (1988) highlights that any allocation resulting from a maximization problem of a couple under full commitment is Pareto efficient. Thus, all deviations in the limited-commitment model from the full-commitment benchmark result in welfare losses caused by commitment issues.

When spouses can fully commit, the Pareto weight is fixed once and for all in the household formation step, μ_0^{fc}.[5] Then, following the approach of Chiappori (1988), household preferences can be characterized as the weighted sum of individual preferences

$$v\left(c_m, c_f, \mu_0^{fc}\right) = \mu_0^{fc} \cdot u_m(c_m) + \left(1 - \mu_0^{fc}\right) \cdot u_f(c_f) + \Psi. \tag{59}$$

Spouses can share idiosyncratic income risks and consume their joint income. The period budget constraint is

$$c_m + c_f \leq y_m + y_f. \tag{60}$$

The maximization problem of the full-commitment couple household in recursive formulation is

$$V^{fc}\left(\mathbf{y}, \mu_0^{fc}\right) = \max_{c_m, c_f} v\left(c_m, c_f, \mu_0^{fc}\right) + \beta \cdot \mathbb{E}\left[V^{fc}\left(\mathbf{y}', \mu_0^{fc}\right)\right], \tag{61}$$

subject to the period budget constraint (60).

[5] For now, we assume that the Pareto weight is exogenously given. Strictly speaking, a model version with an exogenously given Pareto weight is a unitary model since the Pareto weight is not a function of prices, incomes and distribution factors. In Appendix A.3, we consider the endogenous determination of the Pareto weight in the household formation step. The exact determination method has no impact on the general risk-sharing implications under full commitment.

The resulting first-order conditions are

$$\mu_0^{fc} \cdot \frac{\partial u_m(c_m)}{\partial c_m} = \left(1 - \mu_0^{fc}\right) \cdot \frac{\partial u_f(c_f)}{\partial c_f} = \lambda, \tag{62}$$

together with the budget constraint (60), given the exogenous income vectors $\mathbf{y}$, and the constant Pareto weight μ_0^{fc}.[6] λ denotes the Lagrange multiplier on the budget constraint. Condition (62) implies that spouses achieve perfect risk sharing within the household, as the ratio of marginal utilities of consumption is constant over time and states of the world, and fulfills

$$\frac{\partial u_m(c_m)/\partial c_m}{\partial u_f(c_f)/\partial c_f} = \frac{1 - \mu_0^{fc}}{\mu_0^{fc}}. \tag{63}$$

Aggregate income of spouses, $y_m + y_f$, may change over time depending on income shocks and couples cannot insure against aggregate fluctuations. However, to reduce individual consumption volatility as much as possible, the couple fixes the ratio of marginal utilities of consumption of the spouses.

Limited-commitment couple. In the limited-commitment model, the Pareto weight may change over time depending on the particular values of the state variables. As before, household preferences can be characterized as the weighted sum of individual preferences

$$v\left(c_m, c_f, \mu'\right) = \mu' \cdot u_m(c_m) + \left(1 - \mu'\right) \cdot u_f(c_f) + \Psi, \tag{64}$$

but in this case μ' follows the updating rule (57). The initial value of the Pareto weight, μ_0, is determined at the household formation step.[7] When forming expectations, spouses take potential changes in the Pareto weights into account.

The maximization problem of a couple household under limited commitment is

$$V(\mathbf{y}, \mu) = \max_{c_m, c_f} \; v\left(c_m, c_f, \mu'\right) + \beta \cdot \mathbb{E}\left[V\left(\mathbf{y}', \mu'\right)\right], \tag{65}$$

subject to the period budget constraints (60), and the updating rule for the Pareto weight (57). The timing convention in the maximization problem is as follows. Spouses enter the period with the previous period Pareto weight μ. Then, income shocks realize. If one of the participation constraints is violated, spouses re-bargain and the Pareto weight is updated following condition (57). If no participation constraint is binding, the Pareto weight is unchanged. The particular allocation in a period depends on the updated Pareto weight, μ'.

[6] The derivation of the first-order conditions can be found in Appendix A.2.

[7] Again, we assume for now that the initial value of the Pareto weight is exogenously given. In Appendix A.3, we discuss how the initial Pareto weight can be determined.

The resulting first-order conditions are[8]

$$\mu' \cdot \frac{\partial u_m (c_m)}{\partial c_m} = \left(1 - \mu'\right) \cdot \frac{\partial u_f (c_f)}{\partial c_f} = \lambda, \tag{66}$$

together with the budget constraint (60), given the exogenous income vectors $\mathbf{y}$, and the updating rule for the Pareto weight (57). λ denotes the Lagrange multiplier on the budget constraint. Importantly, condition (66) implies that intra-household risk-sharing is imperfect when commitment is limited:

$$\frac{\partial u_m (c_m) / \partial c_m}{\partial u_f (c_f) / \partial c_f} = \frac{1 - \mu'}{\mu'}. \tag{67}$$

The right-hand side of (67) may vary over time and over states. In consequence, individual consumption volatilities may differ between spouses. In contrast, the right-hand side of the corresponding full-commitment condition (63) is fixed. Thus, the limited-commitment consumption allocation differs from the Pareto efficient full-commitment counterpart, implying a welfare loss. In this simple model, the sole source of inefficiency is imperfect risk-sharing due to potential time-variation and state-variation in the ratio of marginal utilities of consumption.

Individual value functions of the spouses are given by

$$V_i (\mathbf{y}, \mu) = u_i (\mathcal{C}_i (\mathbf{y}, \mu)) + \Psi + \beta \cdot \mathbb{E} \left[V_i \left(\mathbf{y}', \mu' \right) \right], \tag{68}$$

using the optimal consumption policy functions $\mathcal{C}_i (\mathbf{y}, \mu)$. These individual lifetime utilities are used when evaluating the participation constraints in the updating rule (57).

[8] The derivation of the first-order conditions can be found in Appendix A.2.

4 Numerical solution

To solve the model, we use standard computational techniques from dynamic economics, see Heer and Maussner (2009) for an overview.[1]

The general solution method follows a recursive approach. Since there are no savings opportunities, there is no intertemporal decision and the maximization problem simplifies to a sequence of static problems. Given the state variables, the solution can be found by solving the set of potentially non-linear first-order conditions.

In the stochastic endowment economy, individual incomes are the individual stochastic state variables. Incomes are assumed to follow first-order autoregressive processes and are discretized using the algorithm of Tauchen (1986).

We combine the exogenous stochastic state variables to one variable, $\mathbf{y}$, which captures the combination of husbands' and wives' incomes. The joint transition matrix of the husband-wife income combinations, P, follows from the transition matrices of the individual stochastic processes, P_m and P_f, by $P = P_f \otimes P_m$. The number of grid points of the joint income grid is $N_y = N_{y_m} \times N_{y_f}$, where N_{y_g} denotes the number of grid points of the individual income process of a spouse with gender g. Throughout, we assume the same Thus, in a couple model, the number of grid points of the stochastic state variable is squared compared to a model with bachelor households which increases the computation time substantially.

Both in the collective model with full commitment and in the collective model with limited commitment, the Pareto weight is endogenous. Therefore, an additional grid for the Pareto weight is needed in the computation. While the Pareto weight is fixed over time and states in a full-commitment model, and hence strictly speaking not a state variable, a grid is needed to find the optimal Pareto weight in the household formation step.[2] In the case of limited commitment, a grid for the Pareto weight is needed since the weight is state-dependent. As discussed before, in the limited-commitment model, the participation constraints must be fulfilled in every period, in any stochastic state, and after any history of stochastic realizations. In the numerical computation, this can be implemented by including the Pareto weight as an additional state variable. The Markovian structure of the model makes this approach feasible since the history is summarized in the current state variables.

In the numerical implementation, we build three-dimensional meshgrids. The first dimen-

[1] We use MATLAB for the numerical implementation.

[2] See Appendix A.3 for details.

sion is the stochastic income of the husband, the second dimension is the stochastic income of the wife, and the third dimension is the Pareto weight. The state space of the couple maximization problem is $\mathbf{N} = N_{y_m} \times N_{y_f} \times N_\mu$, where N_μ denotes the number of grid points for the Pareto weight. We solve for continuous solutions for the choice variables using the first-order conditions of the maximization problem.

Algorithm 1 summarizes the numerical solution of the full-commitment maximization problem.[3]

Algorithm 1 (Full commitment) *The algorithm consists of the following three steps:*

1: Set a convergence criterion ζ. Discretize the exogenous income processes using the approach of Tauchen (1986). Build an equi-distant grid for the Pareto weight $\{\mu_1, ..., \mu_{N_\mu}\}$.

2: Solve the potentially non-linear first-order conditions (60) and (62) given the grids from step 1. The results are the optimal policy function $\mathcal{C}_i^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)$, $i \in \{m, f\}$.

3: Choose initial guesses for the value functions $V_{m,0}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)$, and $V_{f,0}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)$.

3a: Given that you are in iteration step j, use the Bellman equation (61) to get individual value functions:

$$\begin{aligned} V_{m,j+1}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right) &= u_m\left(\mathcal{C}_m^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)\right) + \beta \cdot \mathbb{E}\left[V_{m,j}^{fc}\left(\mathbf{y}', \mu_0^{fc}\right)\right], \text{ and} \\ V_{f,j+1}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right) &= u_f\left(\mathcal{C}_f^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)\right) + \beta \cdot \mathbb{E}\left[V_{f,j}^{fc}\left(\mathbf{y}', \mu_0^{fc}\right)\right]. \end{aligned}$$

3b: If $\sup\left|V_{m,j+1}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right) - V_{m,j}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)\right| < \zeta$, and $\sup\left|V_{f,j+1}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right) - V_{f,j}^{fc}\left(\mathbf{y}, \mu_0^{fc}\right)\right| < \zeta$ stop. Otherwise set $j = j + 1$ and go back to Step 3a. Iterate until convergence.

The static nature of Step 2 of Algorithm 1 implies that the potentially time-consuming root-finding to obtain the consumption policy functions has only to be done once which speeds up the numerical solution substantially (this however is not possible in more general models). To determine the policy functions in Step 2, we use a fast vectorized bisection root-finding routine. Lifetime utilities can then be determined by iterating the Bellman equation (61) until convergence, see Step 3 of the algorithm, starting with an initial guess.

A special case. If spouses have identical risk aversion, $\sigma_m = \sigma_f$, the consumption policy in Step 2 of the algorithm can be determined analytically. When period utility is of the CRRA-form,

$$u_i\left(c_i\right) = \frac{c_i^{1-\sigma_i} - 1}{1 - \sigma_i}, \text{ for } i = \{m, f\},$$

[3] The corresponding algorithm for the Bachelor household can be found in Appendix A.4.

the first-order conditions of the full-commitment couple maximization problem are

$$c_m + c_f \leq y_m + y_f,$$
$$\mu_0^{fc} \cdot c_m^{-\sigma_m} = \left(1 - \mu_0^{fc}\right) \cdot c_f^{-\sigma_f}.$$

If spouses are heterogeneous with respect to σ_i, no closed-form solution of the first-order conditions exist. However, using $\sigma_m = \sigma_f = \sigma$, the consumption policy functions are

$$\mathcal{C}_m\left(\mathbf{y}, \mu_0^{fc}\right) = \left(1 + \left(\frac{\mu_0^{fc}}{1 - \mu_0^{fc}}\right)^{-1/\sigma}\right)^{-1} \cdot (y_m + y_f),$$
$$\mathcal{C}_f\left(\mathbf{y}, \mu_0^{fc}\right) = \left(\frac{\mu_0^{fc}}{1 - \mu_0^{fc}}\right)^{-1/\sigma} \cdot \left(1 + \left(\frac{\mu_0^{fc}}{1 - \mu_0^{fc}}\right)^{-1/\sigma}\right)^{-1} \cdot (y_m + y_f).$$

Limited commitment. As in the full-commitment model, the consumption policy functions have to be determined only once in this model (for the grid of possible Pareto weights) since there are no intertemporal decisions. However, under limited commitment, the participation constraints have to be checked in each iteration of the algorithm. The converged state-dependent grid of the Pareto weight, $\Theta(\mathbf{y}, \mu)$, is the set of optimal intervals of the Pareto weight from the updating rule (58). Algorithm 2 summarizes the numerical solution of the limited-commitment maximization problem.

Algorithm 2 (Limited commitment) *The algorithm consists of the following eight steps:*

1: Set convergence criteria ζ_1 and ζ_2. Discretize the exogenous income processes using the approach of Tauchen (1986). Build an equi-distant grid for the Pareto weight $\{\mu_1, ..., \mu_{N_\mu}\}$.

2: Bachelor household solution (see Appendix A.4)

3: Full-commitment solution.

3a: Perform Step 2 of Algorithm 1.

3b: Perform Step 3 of Algorithm 1.

4: Initialization limited-commitment solution.

4a: Guess initial value functions $V_{m,0}(\mathbf{y}, \mu)$, and $V_{f,0}(\mathbf{y}, \mu)$.

4b: Guess initial values for a state-dependent Pareto weight grid $\Theta_0(\mathbf{y}, \mu)$.

4c: Construct a state-independent Pareto weight grid $\tilde{\Theta}$ for the unconstrained case.

5: *Unconstrained solution (ignoring participation constraints). Given that you are in iteration step j, use the Bellman equation to get:*

$$V_m^{uncon}(\mathbf{y},\mu) = u_m\left(\mathcal{C}_m^{fc}(\mathbf{y},\mu)\right) + \beta\cdot\mathbb{E}\left[V_{m,j}(\mathbf{y}',\mu)\right], \textit{ and}$$
$$V_f^{uncon}(\mathbf{y},\mu) = u_f\left(\mathcal{C}_f^{fc}(\mathbf{y},\mu)\right) + \beta\cdot\mathbb{E}\left[V_{f,j}(\mathbf{y}',\mu)\right].$$

6: *Check participation constraints.*

6a: *Find values of the Pareto weight μ such that the participation constraints hold with equality, using the value functions $V_m^{uncon}(\mathbf{y},\mu)$, and $V_f^{uncon}(\mathbf{y},\mu)$ from Step 5. To do so, solve for the roots of*

$$V_m^{uncon}(\mathbf{y},\mu) - S_m(y_m) = 0,$$
$$V_f^{uncon}(\mathbf{y},\mu) - S_f(y_f) = 0,$$

where $S_i(y_i)$ denotes the outside option of spouse i. $\underline{\nu}(\mathbf{y})$ denotes the Pareto weight such that the participation constraint of the husband holds with equality. $\overline{\nu}(\mathbf{y})$ denotes the Pareto weight such that the participation constraint of the wife holds with equality.

6b: *Build a new grid $\Theta_{j+1}(\mathbf{y},\mu) = \tilde{\Theta}$. Replace all values for a particular state combination $\mathbf{y}$ which are smaller than $\underline{\nu}(\mathbf{y})$ by $\underline{\nu}(\mathbf{y})$. Replace all values for a particular $\mathbf{y}$ which are larger than $\overline{\nu}(\mathbf{y})$ by $\overline{\nu}(\mathbf{y})$.*

7: *Constrained case (respecting participation constraints). Set $V_{m,j+1}(\mathbf{y},\mu) = V_m^{uncon}(\mathbf{y},\mu)$ and $V_{f,j+1}(\mathbf{y},\mu) = V_f^{uncon}(\mathbf{y},\mu)$. Find all values of $V_{m,j+1}(\mathbf{y},\mu)$ smaller than $S_m(y_m)$. Replace these values by $S_m(y_m)$. Adjust the corresponding values of $V_{f,j+1}(\mathbf{y},\mu)$ by interpolating $V_f^{uncon}(\mathbf{y},\mu)$ on the corresponding relative Pareto weight values.*
Find all values of $V_{f,j+1}(\mathbf{y},\mu)$ smaller than $S_f(y_f)$. Replace these values by $S_f(y_f)$. Adjust the corresponding values of $V_{f,j+1}(\mathbf{y},\mu)$ by interpolating $V_f^{uncon}(\mathbf{y},\mu)$ on the corresponding relative Pareto weight values.

8: *If sup $|V_{m,j+1}(\mathbf{y},\mu) - V_{m,j}(\mathbf{y},\mu)| < \zeta_1$, sup $|V_{f,j+1}(\mathbf{y},\mu) - V_{f,j}(\mathbf{y},\mu)| < \zeta_1$, and sup $|\Theta_{j+1}(\mathbf{y},\mu) - \Theta_j(\mathbf{y},$ ζ_2 stop. Otherwise set $j = j+1$ and go back to Step 5. Iterate until convergence.*

The first three steps of Algorithm 2 follow directly from Algorithm 1 and Appendix A.4. The solution of the bachelor problem in Step 2 is used to check the participation constraints in Step 6. We use the lifetime utilities of the full-commitment couple as initial guesses for the value functions in the limited-commitment solution.[4] For the initialization of the set of optimal intervals, $\Theta(\mathbf{y},\mu)$, we use the Pareto weight grid from Step 1, $\{\mu_1,...,\mu_{N_\mu}\}$. A natural choice for the state-independent grid of the Pareto weight, $\tilde{\Theta}$, in Step 4c is the Pareto weight grid $\{\mu_1,...,\mu_{N_\mu}\}$.

[4]Other initial guesses are possible. However, one should not use initial values which lead to situations where both participation constraints are binding. Alternative good initial guesses are values for the value functions which are much larger than the corresponding lifetime utilities of bachelor households.

The solution algorithm for the limited-commitment model is based on the idea that, beginning with the next period, commitment between spouses is limited. In Step 5, where the unconstrained solution is determined, the lifetime utilities on the right-hand side of the Bellman equations are lifetime utilities of households/spouses who cannot fully commit. However, for the solution of the current period (given current state variables $(\mathbf{y}, \mu)$) it is assumed that spouses can fully commit. Thus, in Step 5, the participation constraints for the current period are ignored. The resulting solutions are the optimal consumption policy functions of the full-commitment problem, $\mathcal{C}_m^{fc}(\mathbf{y}, \mu)$ and $\mathcal{C}_f^{fc}(\mathbf{y}, \mu)$, since there are no intertemporal decisions in this model. Given these solutions to the first-order conditions, the lifetime utilities calculated in Step 5 are the lifetime utilities of a couple household where spouses can fully commit until the end of the current period and thereafter commitment between spouses is limited.

These lifetime utilities are then used in Step 6a to check the participation constraints. In particular, one searches for the values of the Pareto weight such that, given a particular state combination $\mathbf{y}$, the participation constraints hold with equality. To do so, we use a fast vectorized bisection root-finding routine. The idea is to determine the value of the Pareto weight the couple has to start with in the period to fulfill both participation constraints. These values are denoted by $\underline{\nu}(\mathbf{y})$ and $\overline{\nu}(\mathbf{y})$ and build an interval of the Pareto weight where no participation constraint is violated. In Step 6b, a new state-dependent grid for the Pareto weight is built (initialized by the grid $\tilde{\Theta}$) where all values which are smaller than $\underline{\nu}(\mathbf{y})$ are replaced by $\underline{\nu}(\mathbf{y})$. All values which are larger than $\overline{\nu}(\mathbf{y})$ are replaced by $\overline{\nu}(\mathbf{y})$.

In Step 7, all values of the value functions $V_{m,j+1}(\mathbf{y}, \mu)$ and $V_{f,j+1}(\mathbf{y}, \mu)$ which are smaller than the outside options are replaced by the values of the outside option. One also has to adjust the particular values of the value function of the spouse whose participation constraint is not binding. We interpolate both value functions on the new Pareto weight grid $\Theta_{j+1}(\mathbf{y}, \mu)$. This guarantees that both participation constraints are always fulfilled and that the lifetime utilities of spouses with non-binding participations constraints correspond to the correct Pareto weight value. These interpolated value functions become the initial guesses in a potential new iteration of the algorithm.

Finally, in Step 8, the convergence of all lifetime utilities and the state-dependent Pareto weight grid is checked. Until convergence, one uses the newly calculated values from Step 6b and Step 7 as new initial guesses and goes back to Step 5. The algorithm guarantees that eventually one has found lifetime utilities of both spouses and the couple household which respect the participation constraints in all periods, i.e., lifetime utilities of spouses under limited commitment.

The converged state-dependent grids of the Pareto weight are the optimal intervals of the Pareto weight from the updating rule (58). These intervals and (58) are used in the simulation of the model. At the beginning of a period, each simulated couple household has state variables $\mathbf{y}$ and the previous period Pareto weight μ. Then, the updating rule (58) is used to check whether an update of the Pareto weight is required or not.

5 Aggregate risk-sharing effects of limited commitment

We now describe the results of our quantitative analysis. We start with the aggregate effects of limited commitment on the risk-sharing behavior of married couples. To this end, we first consider a model version that has little heterogeneity across households. In Section 6, we take into account heterogeneity in a number of dimensions to study cross-sectional implications.

5.1 Functional forms and parameterization

Table 18 summarizes the parameter values for the baseline parameterization of the model. Individual utility functions are given by

$$u\left(c_i\right) = \frac{c_i^{1-\sigma} - 1}{1 - \sigma}, \tag{69}$$

where $1/\sigma$ measures the elasticity of intertemporal substitution of consumption. To highlight endogenous commitment issues in a transparent way, we abstract from gender heterogeneity in parameter values in the baseline parameterization. We assume $\sigma \rightarrow 1$ in the baseline parameterization, such that (69) becomes

$$u\left(c_i\right) = \log\left(c_i\right). \tag{70}$$

The discount factor β is set to 0.96, which is a standard value for a yearly calibration of an incomplete-markets model.

We assume that individual log income consists of two elements, a fixed effect, ψ, and a stochastic component, z:

$$\log y_i = \psi_i + z_i. \tag{71}$$

We account for a fixed-effect component to induce heterogeneity across households which is needed to calibrate the amount of commitment issues in the economy, see below. The fixed-effect component is determined once and for all before spouses marry. We follow Conesa and Krueger (2006) and assume that ψ can take two values for each spouse, so that there are four types of households with different combinations of the fixed effects for husband and wife, respectively. ψ is assumed to be log-normally distributed with mean zero and variance σ_ψ^2 and, following Conesa and Krueger (2006), we set the two values to $\psi_{i,1} = -\sigma_\psi$ and $\psi_{i,2} = \sigma_\psi$ and assume equal population shares for both types. The value for σ_ψ^2 is taken from

Table 18: Parameter values, stochastic endowment economy

Description	Parameter	Value	Source
CRRA parameter	σ	1	set
Discount factor	β	0.96	set
Marriage utility gain	Ψ	0.002	calibrated
Variance fixed effect	σ_ψ^2	0.4472	Storesletten et al. (2004)
Autocorrelation income shock	ρ	0.9	Aiyagari (1994)
Standard deviation income shock	σ_ε	0.0872	Aiyagari (1994)

NOTE.–There are four household types. In the simulated economy, income fixed effects are positively correlated between spouses with a correlation coefficient of 0.5, see main text.

Storesletten et al. (2004), see Table 18. Hyslop (2001) finds that the correlation between fixed effects of spouses is 0.5. To achieve this, we set the share of the two couple types where both partners have the same fixed effect to 37.5% each.[1] The share of couples where partners differ with respect to their fixed effects is set to 12.5% each.[2]

In line with large parts of the infinite-horizon incomplete-markets literature, we assume that the stochastic income component follows a first-order autoregressive process with autocorrelation ρ and variance σ_ε^2:

$$z_i' = \rho \cdot z_i + \varepsilon_i, \tag{72}$$

with $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$. In the baseline model, we assume that incomes are uncorrelated between spouses but we consider the case of correlated income shocks in Section 6.5. The values for ρ and σ_ε^2 are taken from Aiyagari (1994), see Table 18.

A main target of the calibration is to generate a realistic extent of commitment problems. Technically, commitment problems can be strengthened or weakened by adjusting the parameter for the utility gains from marriage, Ψ. This parameter can be interpreted as all gains from marriage not directly modeled such as home production, children, etc. we calibrate Ψ using an empirical target provided by Lise and Yamada (2019).

Lise and Yamada (2019) decompose the total variance of the (logged) male Pareto weight as the sum of the variance of the (logged) male Pareto weight across households and the variance of the (logged) male Pareto weight within households over time. Thus, the Lise-

[1]When both spouses have the same fixed effect, the correlation is 1, and when spouses have different fixed effects, the correlation is -1. The average correlation of fixed effects in the whole economy then is $2 \cdot 0.375 - 2 \cdot 0.125 = 0.5$.

[2]The fixed-effect combinations are $\psi_m = \psi_f = \sigma_\psi$ in household type I, $\psi_m = \sigma_\psi$ and $\psi_f = -\sigma_\psi$ in household type II, $\psi_m = -\sigma_\psi$ and $\psi_f = \sigma_\psi$ in household type III, and $\psi_m = \psi_f = -\sigma_\psi$ in household type IV.

Yamada-decomposition is given by:

$$\mathrm{Var}\left(\log \mu_{jt}\right) = \mathrm{Var}\left(\mathbb{E}\left(\log \mu_{jt}\right)\middle|\, j\right) + \mathbb{E}\left(\mathrm{Var}\left(\log \mu_{jt}\right)\middle|\, j\right), \tag{73}$$

where j denotes the household and t time. The first term in the Lise-Yamada-decomposition stems from heterogeneity *between* couples, $\mathrm{Var}\left(\mathbb{E}\, \mu_{jt}\middle|\, j\right)$, and the second term from *within-household* rebargaining over time due to binding participation constraints, $\mathbb{E}\left(\mathrm{Var}\, \mu_{jt}\middle|\, j\right)$. If spouses could fully commit, the within-household variance would be zero. Any positive within-household variance of the Pareto weights over time is caused by commitment issues between spouses. Any positive across-household variance of the Pareto weight is caused by heterogeneity across households, in this model by heterogeneous long-run incomes of spouses due to income fixed effects. To measure the strength of commitment issues, we use the ratio

$$\Upsilon = \frac{\mathbb{E}\left(\mathrm{Var}\left(\log \mu_{jt}\right)\middle|\, j\right)}{\mathrm{Var}\left(\mathbb{E}\left(\log \mu_{jt}\right)\middle|\, j\right)}. \tag{74}$$

Lise and Yamada (2019) use panel data from the Japanese Panel Survey of Consumers (JPSC) to analyze the extent of limited commitment in intra-household decision-making and the resulting impact on intra-household allocation rules. They document that the within-household variance is about one-fourth of the across-household variance, implying $\Upsilon = 0.25$. Thus, about 25% of the heterogeneity in Pareto weights is due to rebargaining over time. We calibrate Ψ in our model to match this empirical target.[3]

In the numerical solution, we use 21 income grid points per spouse and 501 grid points for the Pareto weight grid, resulting in a total of $220,941$ grid points. After having solved the model, we simulate a panel of $5,000$ couple households, both under full and limited commitment, using the same stochastic shocks. We also simulate $5,000$ singles of each gender.[4] We simulate $1,500$ periods and discard the first 500 periods to avoid dependence on initial conditions. Initial Pareto weights are determined using a Nash bargaining approach.[5] The resulting initial male Pareto weights are 0.5 (household types I and IV), 0.68 (household type II), and 0.32 (household type III).

5.2 Policy functions

The left panel in Figure 2 shows the Pareto weight of the wife, $1-\mu$, as a function of her income. The right panel of Figure 2 shows consumption of the wife as a function of her

[3]In the U.S., divorce rates are even higher than in Japan, see OECD (2019). Thus, the calibration target seems a lower bound for the extent of commitment issues in the United States.

[4]We compare single households to couple households in our analysis below.

[5]See Appendix A.3 for details.

Figure 2: Policy functions for the female Pareto weight and female consumption.

(a) Female Pareto weight (b) Female consumption

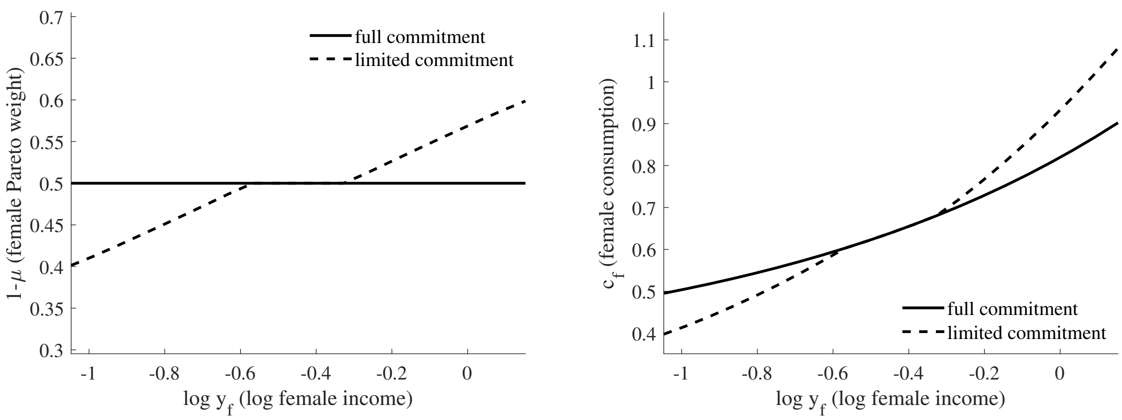

NOTE.– Policy functions for the household type where both spouses have the low fixed effect. Thus, there are no long-run income differences. In both panels, male income is at its mean value. Solid lines refer to a full-commitment couple, dashed lines refer to a limited-commitment couple.

income. In both panels, the policy functions refer to a situation where male income is at its mean level. Solid lines refer to the full-commitment model version and dashed lines refer to the limited-commitment model. For the figures, we consider household type IV which is the household type where both spouses have the low income fixed effect implying that there are no long-run income differences between spouses.[6] For plotting the policy functions of the limited-commitment model, we assume that the previous period female Pareto weight, $1 - \mu$, is given by the long-run mean which is 0.5 for household type IV.[7]

In the full-commitment model, the Pareto weight is fixed once and for all and does not depend on income shocks, see the solid line in the left panel of Figure 2. In the limited-commitment model, however, the current period Pareto weight depends on the intra-household distribution of incomes. The dashed line shows that the female Pareto weight increases when the relative income of the wife increases (recall that the income of the husband is held constant). These changes in the Pareto weights translate into changes in individual consumption shown in the right panel of Figure 2. A higher female income implies higher household income since male income is fixed at its mean value. Thus, independent of whether spouses can fully commit or commitment is limited, female consumption is increasing in her

[6] Policy functions for the other household types show the same general pattern.

[7] After shocks have realized, the new value of the Pareto weight is determined through the updating rule (58).

Table 19: Rebargaining frequency

		household type			
	Average	I	II	III	IV
Rebargaining frequency in %	37.23	37.32	36.95	36.95	37.32

NOTE.–The column 'Average' shows type-weighted averages. The other columns show household-type-specific averages. I: Both spouses have high income fixed effects. II: Male high income fixed effect, female low income fixed effect. III: Male low income fixed effect, female high income fixed effect. IV: Both spouses have low income fixed effects.

income. However, as shown by the dashed line, the slope of female consumption as a function of her income is larger in the limited commitment case compared to the full-commitment case. When relative female income is low, she receives less consumption in the limited-commitment case than under full commitment. When relative female income is high, however, she receives even higher consumption under limited commitment. These effects are driven by the change in the female bargaining position in response to her income shocks. Thus, the policy functions show that individual consumption is more volatile under limited commitment than under full commitment (analogous results hold for male individual consumption). This means that full commitment provides better consumption insurance than limited commitment and limited commitment is hence associated with a welfare loss.

5.3 Simulation results

Recall that the model is calibrated to match the statistic reported by Lise and Yamada (2019) for the variance in Pareto weights within households relative to the variance of Pareto weights across households, $\Upsilon = 0.25$. As a second measure of the strength of commitment issues, we calculate the frequency of changes in the Pareto weight in the simulated data. Any change in the Pareto weight can be interpreted as a rebargaining of spouses. Thus, higher rebargaining frequencies tend to represent more commitment issues. Table 19 summarizes the results, overall and separately for all four households types. The overall rebargaining frequency is about 37% and relatively similar across household types. Thus, on average, a couple renegotiates over the Pareto weights every 2.7 periods.

We now analyze the risk-sharing impact of limited commitment in the simulated data. In the full-commitment case, the within-household variances of individual consumption growth are identical for spouses. Further, these variances are equal to the within variance of household-aggregate consumption growth since individual consumption is proportional to

Table 20: Risk sharing

	(1)	(2)	(3)
Measure	Bachelor	Full commitment	Limited commitment
Variance of consumption growth			
household, Ξ	–	0.44	0.44
men, Ξ_m	0.83	0.44	0.53
women, Ξ_f	0.83	0.44	0.53
Kaplan-Violante coefficients			
household, ϕ	–	0	0
men, ϕ_m	0	0.50	0.32
women, ϕ_f	0	0.50	0.32

NOTE.–Ξ denotes the variance of consumption growth at the couple level and Ξ_i denotes the variance of individual consumption growth for spouse i, $\Xi_i = \text{var}(\log c_i' - \log c_i)$. ϕ_i is the Kaplan and Violante (2010) insurance coefficient of spouse i, $\phi_i = 1 - \text{cov}\,(\log c_i' - \log c_i, \varepsilon_i)/\text{var}\,(\varepsilon_i)$.

household consumption,

$$c_f = \left(\frac{\mu_0^{fc}}{1-\mu_0^{fc}}\right)^{-\frac{1}{\sigma}} \cdot c_m.$$

The limited commitment intra-household risk-sharing condition is

$$c_f = \left(\frac{\mu'}{1-\mu'}\right)^{-\frac{1}{\sigma}} \cdot c_m.$$

The Pareto weights on the right-hand side of the limited-commitment risk-sharing condition change in response to sufficiently strong income shocks, implying that individual consumption growth is in general not equal to household-aggregate consumption growth.

To measure the extent of risk sharing, we define the consumption-risk measures (aggregate, men, women) as

$$\Xi = \text{var}\left(\Delta \log\left(c_m' + c_f'\right)\right) \cdot 100, \tag{75}$$

$$\Xi_m = \text{var}\left(\Delta \log c_m'\right) \cdot 100, \tag{76}$$

$$\Xi_f = \text{var}\left(\Delta \log c_f'\right) \cdot 100, \tag{77}$$

where Δ is the first-difference operator, i.e., for a generic variable b, $\Delta b' = b' - b$. The advantage of considering consumption growth rates is that they take out variation in mean consumption across households and thus isolate individual consumption risk, i.e., volatility of individuals' consumption over time.

The upper part of Table 20 summarizes the results. In the simulated economy, the variances of individual consumption growth, Ξ_m and Ξ_f, are highest in bachelor households and lowest in full-commitment couple households. Bachelor households cannot share risk and have to consume their period income so that the variance of consumption growth (see column (1) is equal to the variance of income growth. Being a member of a couple household provides better insurance against income shocks than being a single even when commitment in the couple is limited (see column (3)). However, under limited commitment, there is less intra-household risk-sharing compared to a situation of full commitment (see column (2)). The full-commitment situation achieves perfect risk sharing and the variance of consumption growth is reduced by 39 percentage points compared to the bachelor situation. Limited commitment between spouses reduces the insurance value of marriage by about one fourth as entering a limited-commitment couple reduces the variance of consumption growth by only 30 percentage points.

As a second measure of risk sharing, we calculate the insurance coefficients suggested by Kaplan and Violante (2010),

$$\phi = 1 - \frac{\text{cov}\left(\Delta \log c_i', \varepsilon_i\right)}{\text{var}\left(\varepsilon_i\right)}, \tag{78}$$

The insurance coefficient measures the share of the variance of the shock that does not translate into consumption growth, see Kaplan and Violante (2010, page 57). We consider three insurance coefficients, one for household consumption and two four spouses' individual consumption levels,

$$\phi = 1 - \frac{\text{cov}\left(\Delta\left(\log\left(c_m' + c_f'\right)\right), \varepsilon_m + \varepsilon_f\right)}{\text{var}\left(\varepsilon_m + \varepsilon_f\right)}, \tag{79}$$

$$\phi_m = 1 - \frac{\text{cov}\left(\Delta \log c_m', \varepsilon_m\right)}{\text{var}\left(\varepsilon_m\right)}, \tag{80}$$

$$\phi_f = 1 - \frac{\text{cov}\left(\Delta \log c_f', \varepsilon_f\right)}{\text{var}\left(\varepsilon_f\right)}. \tag{81}$$

The insurance coefficient for the household, (79), measures the share of the *aggregate household* income variance which does not translate into *aggregate household* consumption growth. In models where the only source of insurance is intra-household risk-sharing, this coefficient is zero. The second and third coefficient (80) and (81) measure the share of *individual* income variance that does not translate into *individual* consumption growth. In couple households, these coefficients will differ from zero as some insurance is achieved through intra-household risk sharing.

The lower part of Table 20 summarizes the results for the Kaplan and Violante (2010)

Table 21: Kaplan and Violante (2010) insurance coefficients, household types

	household type			
	I	II	III	IV
Full commitment				
ϕ	0	0	0	0
ϕ_m	0.50	0.29	0.71	0.50
ϕ_f	0.50	0.71	0.29	0.50
Limited commitment				
ϕ	0	0	0	0
ϕ_m	0.32	0.19	0.47	0.32
ϕ_f	0.32	0.47	0.19	0.32
Reduction due to limited commitment (%)				
Men	36	34	34	36
Women	36	34	34	36

NOTE.–ϕ denotes the couple Kaplan and Violante (2010) insurance coefficient and ϕ_i denotes the insurance coefficient of spouse i. I: Both spouses have a high income fixed effect. II: Male high income fixed effect, female low income fixed effect. III: Male low income fixed effect, female high income fixed effect. IV: Both spouses have a low income fixed effect.

insurance coefficients. Bachelor households have no insurance opportunities in this model and consume their period income. Thus, any income shock translates directly into consumption growth and the insurance coefficients are zero, see the first column of Table 20. Also the couple household as a whole cannot insure against aggregate income fluctuations and hence all changes in aggregate household income directly translate into household consumption growth, see the fourth row of Table 20. However, couples do achieve (imperfect) risk sharing with respect to individual income growth, as the transmission of individual income shocks into individual consumption growth is muted. Since, under full commitment, individual consumption is proportional to the sum of the uncorrelated incomes of both spouses, only 50% of the individual income variance translates into individual consumption growth in the full-commitment household, see column (2) of Table 20. Under limited commitment, individual consumption is less than perfectly proportional to household income and depends more strongly on individual income due to changes in Pareto weights. Quantitatively, the insurance coefficient under full commitment is 0.32. Hence, $100\% - 32\% = 68\%$ of the individual income variance translates into individual consumption growth. Under limited commitment, marriage increases the insurance coefficient by only about two thirds as much as under full commitment, implying a similar reduction in the insurance value of marriage as suggested by the consumption-variance measure.

Table 21 reports insurance coefficients separately by household type. Insurance coefficients of spouses differ in household types II and III where spouses have different long-run incomes (different income fixed effects). Recall that spouses are homogeneous in all parameters except income fixed effects. Thus, the long-run variance of income shocks is smaller for the spouse with the lower fixed effect implying that it is easier for the couple to insure against income shocks of this spouse because they are less important for total household income. Since the insurance coefficient of spouse i, ϕ_i, measures her insurance with respect to her income shock, ϕ_m is smaller than ϕ_f in household type II where the husband has a high fixed effect and the wife has a low fixed effect. In contrast, in household type III $\phi_m > \phi_f$ holds.

Most important for our purposes is to compare insurance coefficients under limited commitment to those under full commitment. For all types of households and both genders, limited commitment reduces insurance coefficients by between 34 and 36%. Hence, as already suggested by the rebargaining frequencies reported in Table 19, the effect of limited commitment on intra-household risk sharing does not differ substantially by mean income of households or individuals.

We now quantify welfare losses that result from imperfect risk sharing under limited commitment, using consumption equivalents. For this, we solve the indifference conditions

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log\left((1+\zeta_m)\cdot c_{mt}^{lc}\right) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log\left(c_{mt}^{fc}\right),$$
$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log\left((1+\zeta_f)\cdot c_{ft}^{lc}\right) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log\left(c_{ft}^{fc}\right).$$

ζ_i denotes the percentage increase in consumption under limited commitment that is needed to make spouse i indifferent to the full-commitment case.

We find that, on average, one needs to give households about 0.6% more consumption under limited commitment to achieve the lifetime utility under full commitment. To put this number into perspective, recall that the welfare costs of business cycles in an incomplete-markets environment are typically estimated to be between 0.08% (Imrohoroğlu 1989) and 1% (Krusell et al. 2009) of consumption.[8] Thus, limited commitment causes a quantitatively important welfare loss.

[8]Storesletten et al. (2001) and Krebs (2007) are further studies determining the welfare costs of business cycles. Although both studies find higher welfare costs from business cycles than 1% in some parameterizations, the results in environments most similar to our baseline parameterization are below 1%. In particular, Storesletten et al. (2001) find welfare costs of 0.82% and Krebs (2007) finds welfare costs of 0.5% in environments most similar to our baseline model.

6 Cross-sectional implications

We now analyze how the risk-sharing effects of limited commitment differ across the population. To this end, we account for heterogeneity in several dimensions. One after another, we vary the parameters governing time preferences, risk aversion, and utility gains from marriage, as well as the variance and correlation of income shocks in the household. In each model extension, we refrain from heterogeneity in all other parameters except for mean wage levels (which we have seen in Section 5 to have very limited impact on the effects of limited commitment.) In an additional experiment, we introduce altruism between spouses in a marriage and allow for heterogeneity in altruism across couples.

The motivation for this analysis is twofold. First, a quantitative analysis of the cross-sectional distribution of the effects of limited commitment can inform economic policies, e.g., when limited commitment impacts particularly on already otherwise vulnerable households. When policy makers want to support specific households particularly, understanding how they are affected by limited commitment allows a better design of the respective policies. Second, in empirical tests of limited-commitment models, it can be helpful to identify cross-sectional variation in the extent to which households are subject to limited commitment. Our analysis delivers important insights in this respect as the preference parameters considered in our analysis can in principle be recovered, e.g., from survey data or experiments. For example, the Health and Retirement Study of the University of Michigan includes questions about behavior in hypothetical situations constructed to deduce respondents' time preference and risk aversion (see, e.g., Barsky et al. 1997, for an application. Further, economists have developed methods to deduce preference parameters such as the rates of time preference or risk aversion from observable behavior such as exercising, drug consumption, gambling, or risky asset holding (e.g., Brown et al. (2006), Brunnermeier and Nagel (2008); Chiappori and Paiella (2011); Finke and Huston (2013)).

6.1 Discount factor

We first incorporate heterogeneity in time preferences into the model. We include some agents whose discount factor β is below the baseline value of 0.96 and some agents whose discount factor is larger.

Kocherlakota (1996) has shown in the context of development economics that commitment

Figure 3: Heterogeneity in household discount factors.

(a) Strength of commitment issues

(b) Individual insurance coefficient

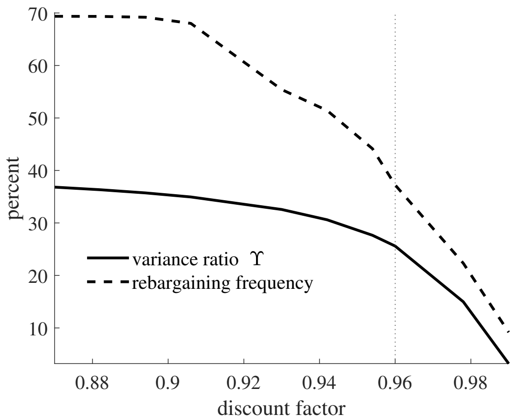

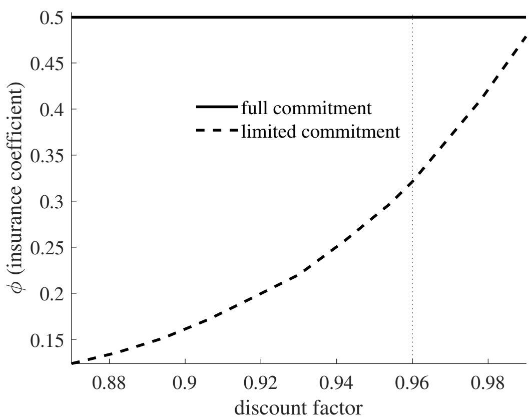

NOTE.–The solid line in the left panel shows the variance ratio (see main text for details). The dashed line in the left panel shows the rebargaining frequency. The right panel shows the individual insurance coefficient under full commitment (solid line) and under limited commitment (dashed line). Average insurance coefficients are identical for men and women since there is no gender heterogeneity in the model. The vertical dotted line indicates the parameter value in the baseline parameterization.

problems between risk-averse agents decrease with their patience. Spouses who are more patient (i.e., who have a higher discount factor β) weigh future gains from marriage more strongly. To understand how this affects commitment problems, consider the case where both spouses have the same Pareto weight. Assume that current income of the wife is above average and current income of the husband is below average. In this situation, the wife expects that her income will decrease in the future while the husband expects an increase in income. If both spouses are impatient, relatively much weight is given to short-run effects and the wife has a relatively strong incentive to leave the household so that the likelihood of rebargaining is relatively high. By contrast, if both spouses are relatively patient, they take into account more strongly future advantages from a risk-sharing agreement. This makes marriage more attractive, resulting in less severe commitment problems.

Figure 3 shows measures of the strength of commitment problems and the Kaplan-Violante insurance coefficient for individual consumption as a function of the household discount factor.[1] The left panel shows that both measures of the importance of commitment problems are smaller for households with a higher discount factor, which corresponds to Kocherlakota (1996)'s result that commitment issues decrease with the patience of spouses. In line with

[1] Average insurance coefficients are identical for men and women since there is no gender heterogeneity in the model.

Figure 4: Impact of higher utility gains from marriage.

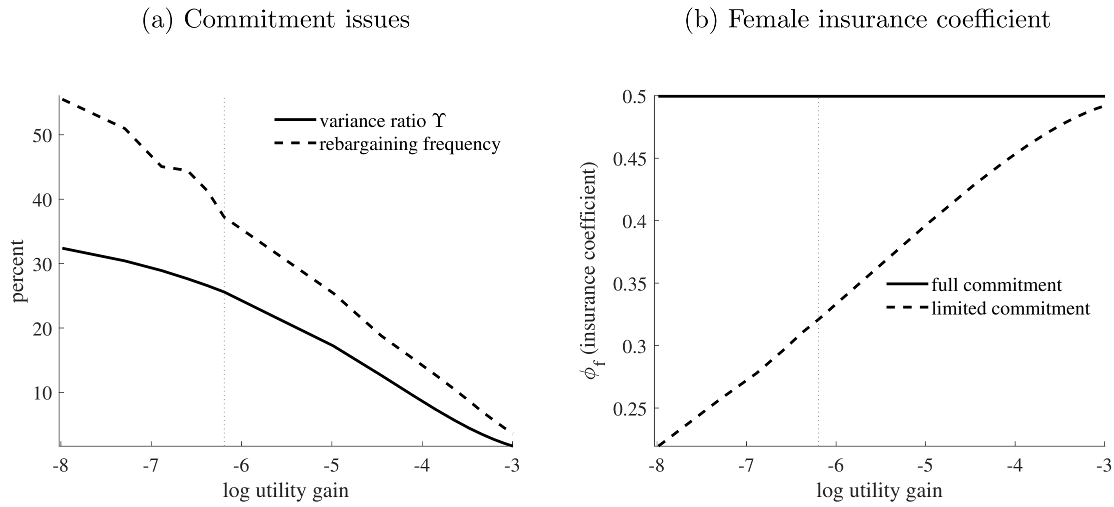

NOTE.–The left panel shows commitment issues measured by the variance ratio and the rebargaining frequency, depending on the direct utility gains from marriage. The right panel shows the female insurance coefficient under full commitment (solid line) and under limited commitment (dashed line). The vertical dotted line indicates the parameter value in the baseline parameterization.

this, the right panel of Figure 3 shows that the insurance coefficient under limited commitment increases is larger in more patient households (dashed line). When agents are very patient, the insurance coefficient under limited commitment gets closer to the full commitment counterpart (solid line).

Note that the lines in Figure 3 are steep, especially around the baseline value of $\beta = 0.96$. Households whose discount factor is only one percentage point larger than those of the baseline household have a rebargaining frequency which is seven percentage points (or 19%) lower and an insurance coefficient which is five percentage points (or 16%) larger.[2]

6.2 Utility gain from marriage

We now consider heterogeneity in terms of the direct utility gains from marriage, Ψ. The direct utility gains from marriage can be expected to vary substantially across married cou-

[2]Falk et al. (2018) find that women are on average less patient than men. We do not analyze heterogeneous discount factors in our infinite horizon model since the maximization problem would become non-stationary when discount factors differ between household members, which makes a numerical solution infeasible in this case. In particular, it would be optimal to have a decreasing consumption path for the impatient spouse and an increasing consumption path for the patient spouse. If spouses could fully commit, this would imply that the consumption share of the impatient spouse converges to zero in the long-run and the consumption share of the patient spouse would converge to one. If spouses cannot fully commit, the impatient spouse would charge for the consumption in every period such that she is indifferent between staying within and leaving the couple. This would induce a further source of inefficiency due to time-inconsistent behavior of spouses under limited commitment.

ples since they capture all gains which are not directly modeled such as home production, love, children, etc. In our baseline parameterization, we used Ψ to calibrate the strength of commitment issues to an empirical target. We now incorporate heterogeneity around the value that results from this calibration.

The left panel in Figure 4 shows that there are smaller commitment problems in households where direct utility gains from marriage are larger. The right panel shows that, in couples where Ψ is sufficiently large, the insurance coefficients under limited commitment converge to the corresponding values under full commitment. Quantitatively, risk sharing is close to perfect only in couples where direct utility gains from marriage are about 20 times larger than in the baseline couple (indicated by the dotted vertical line in the figure). However, the lines in Figure 4 are steep in the neighborhood of the baseline which implies that, in couples, where direct utility gains are moderately larger than in the baseline couple, commitment problems are substantially smaller.

6.3 Risk aversion

We now highlight the impact of different degrees of risk aversion on the severity of commitment issues. The incentives to share risk within the household are directly related to the concept of risk aversion. The more risk averse an individual is, the higher is her gain from risk sharing. Thus, commitment issues should be less severe where the risk aversion of spouses is high since total gains from marriage increase in risk aversion.[3]

Figure 5 shows measures of commitment issues and the insurance coefficient as a function of the degree of risk aversion in the household for a model where the latter varies across households. In this figure, risk aversion is assumed to be homogenous within household, i.e., every husband has the same degree of risk aversion than his wife but a different degree of risk aversion than other husbands and wives. Both, the rebargaining frequency (dashed line) and the variance ratio (solid line) are smaller in more risk averse households, see the left panel of Figure 5. In line with smaller commitment problems, the right panel in Figure 5 shows that intra-household risk sharing is better where risk aversion is higher.

Differences in risk aversion within the household. We now consider the case of heterogeneity in risk aversion *within* a household. This is motivated by the evidence that women are on average more risk averse than men, see Croson and Gneezy (2009) for a survey of the literature and Falk et al. (2018) for a recent study.

[3]If spouses are not risk averse, there are no additional utility gains from marriage next to the direct utility gains Ψ. In our model, risk-neutral spouses would have no incentive to enter and to remain in the household if $\Psi = 0$.

Figure 5: Heterogeneity in risk aversion (homogeneous risk aversion in the couple).

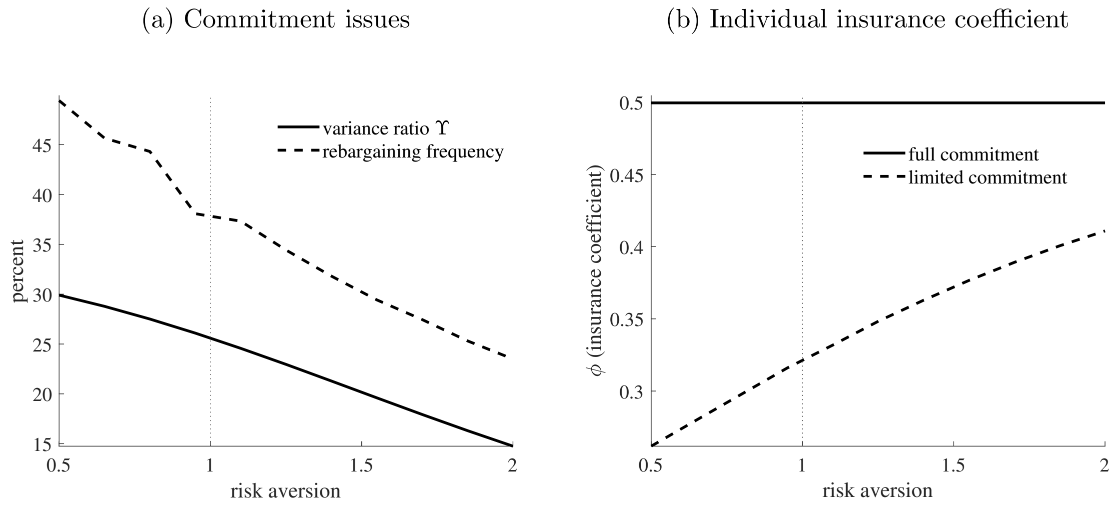

NOTE.–The left panel shows commitment issues measured by the the variance ratio (solid line), Υ, and the rebargaining frequency (dashed line). The right panel shows the female insurance coefficient under full commitment (solid line) and under limited commitment (dashed line). Both spouses have the same constant relative risk aversion. Average insurance coefficients are identical for men and women since there is no gender heterogeneity in the model. The vertical dotted line indicates the parameter value in the baseline parameterization.

Independent of whether commitment is full or limited, within-household differences in risk aversion lead to differences in the insurance coefficients of members of a household. With gender-specific parameters σ_m and σ_f, the risk-sharing condition becomes $\mu' \cdot c_m^{-\sigma_m} = (1-\mu') \cdot c_f^{-\sigma_f}$ which implies that consumption of the less risk-averse spouse will vary more strongly with *household* income. Hence, the insurance coefficient will be larger for the more risk-averse spouse under both full commitment and limited commitment. Under limited commitment, risk aversion exerts an additional effect on insurance coefficients. When one spouse has high risk aversion, he or she values the insurance provided through the marriage strongly which reduces the attractiveness of the outside option and makes it relatively unlikely that this spouse's participation constraint binds. Hence, commitment problems are smaller where risk aversion is higher even when this concerns even only spouse's risk aversion. In Appendix A.5, we provide a formal analysis of the effects of within-household heterogeneity on insurance coefficients under both full commitment and limited commitment.

Figure 6 compares insurance coefficients between households that differ in the degree of risk aversion of the wife (while risk aversion of the husband is held constant at its baseline level). The left panel refers to the case of full commitment. The dashed line shows that the female insurance coefficient is increasing in her risk aversion, holding constant risk aversion

Figure 6: Heterogeneity in risk aversion across and within households.

(a) Full commitment (b) Limited commitment

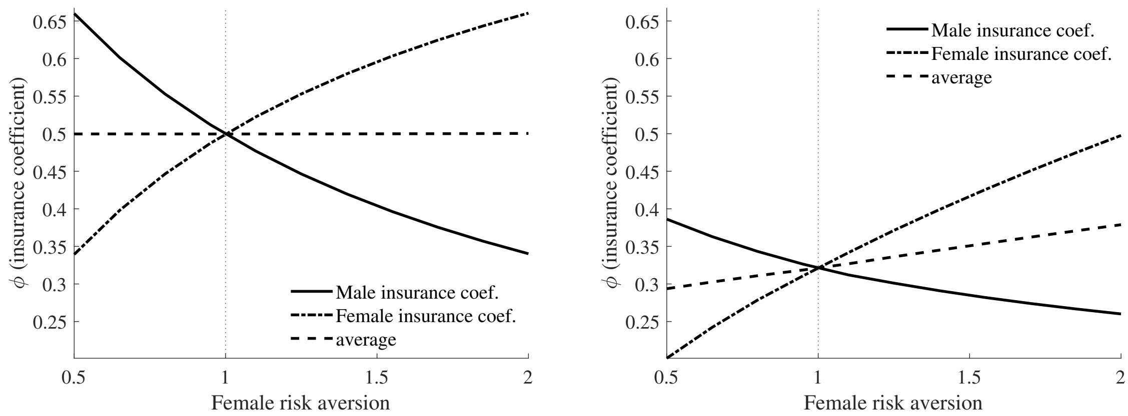

NOTE.–The figure shows male and female insurance coefficients under full commitment (left panel) and limited commitment (right). Female risk aversion is varied and male risk aversion is held constant. The vertical dotted line indicates the baseline parameter value.

of the husband. The solid line shows that the male insurance coefficient is decreasing in female risk aversion. Insurance coefficients are equal when risk aversion is identical for both household members. The average insurance coefficient (represented by the dashed line) under full commitment is constant, indicating that efficient risk sharing is achieved but less risk averse spouses efficiently accept a larger share of the household's income risk.

The right panel of Figure 6 refers to the case of limited commitment. Here, risk aversion affects the strength of commitment problems such that the average degree of consumption insurance is larger in households with more risk averse wives. As under full commitment, households share income risk unevenly with the less risk averse spouse accepting the larger share of the risk. As a consequence, one's own risk aversion has two reinforcing effects on one's insurance coefficient: the household's commitment problems are ameliorated and one has to carry a smaller share of the household's income risk. Quantitatively, the slope of the female insurance coefficient plotted against female risk aversion is larger than in Figure 5 where both spouses' risk aversion was varied. Hence, individuals achieve particularly weak consumption smoothing against income shocks when they live in households with low risk aversion and themselves are even less risk averse than their partners.

6.4 Altruism

We now take into account altruism between spouses and allow for heterogeneity in altruism across couples. Altruism means that individuals receive utility not only from consumption and leisure but also from the well-being of others, in this case, their spouses. Becker (1974a, 1974b) was the first to consider caring about spouses and several other studies have used similar approaches since, see Becker (1991, Chapter 8) and Browning et al. (2014, Chapter 3) for a detailed discussion. Foster and Rosenzweig (2001) consider a limited-commitment model in the context of development economics where households care about other households in the same village. They find empirically that, while altruism alleviates the strength of commitment issues, altruism is not strong enough to remove all commitment issues.

Including altruism requires an extension of the baseline model. Specifically, the household target function (64) in the limited-commitment model becomes

$$\begin{aligned} v\left(c_m, c_f, \mu'\right) &= \mu' \cdot \left(u_m\left(c_m\right) + \varrho_m \cdot u_f\left(c_f\right)\right) + \left(1-\mu'\right) \cdot \left(u_f\left(c_f\right) + \varrho_f \cdot u_m\left(c_m\right)\right) + \Psi, \\ &= \left(\mu' + \varrho_f \cdot \left(1-\mu'\right)\right) \cdot u_m\left(c_m\right) + \left(1 - \mu' + \varrho_m \cdot \mu'\right) \cdot u_f\left(c_f\right) + \Psi, \end{aligned} \tag{82}$$

where ϱ_i denotes the degree of altruism of spouse i. In the corresponding full-commitment model, the Pareto weights in (82) are replaced by the constants μ_0^{fc}. The baseline stochastic endowment economy discussed before assumed the special case $\varrho_m = \varrho_f = 0$, thus no altruism towards the partner. In both cases, full and limited commitment, we assume that spouses do not care about each other in the outside option. A derivation of the first-order conditions of the model with altruism can be found in Appendix A.6.

One can show analytically that commitment issues become less severe with higher degrees of altruism. Consider, for example, the case where the wife faces low income in this period while the husband has high income. Commitment becomes more credible in this period if the husband knows that his wife is an altruist and therefore will compensate him in the future for his support in this period. To see this formally, consider the case $\varrho_m = \varrho_f = 1$. In this case, the household target function (82) becomes:

$$\begin{aligned} v\left(c_m, c_f, \mu'\right) &= \left(\mu' + 1 - \mu'\right) \cdot u_m\left(c_m\right) + \left(1 - \mu' + \mu'\right) \cdot u_f\left(c_f\right) + \Psi, \\ &= u_m\left(c_m\right) + u_f\left(c_f\right) + \Psi. \end{aligned}$$

and is independent of the Pareto weight. If spouses are ex-ante identical, they will receive the same consumption and the same utility in every period and all potential commitment issues are solved by altruism.

Figure 7: Heterogeneity in the degree of altruism in the couple.

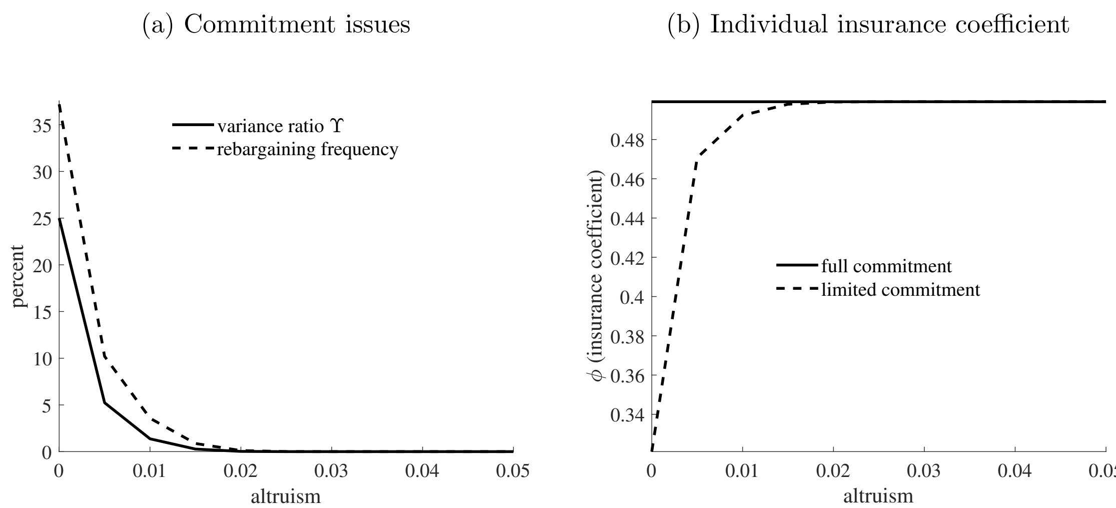

NOTE.–The left panel shows commitment issues measured by the variance ratio (solid line) and the rebargaining frequency (dashed line), depending on the degree of altruism. The right panel shows the individual insurance coefficient under full commitment (solid line) and under limited commitment (dashed line).

We now discuss our quantitative results for commitment problems and risk sharing in a cross-section of couples that differ from one another in terms of altruism. In Figure 7, we vary the altruism parameter ϱ around its baseline value of zero. Recall that the extent of commitment issues is calibrated in the baseline model. Hence, the value of zero should not be understood literally as a couple where spouses do not care about each other but as a normalization indicating the average degree of altruism between spouses. Put differently, the values for ϱ in the figure have no absolute interpretation but measure the deviation from the average degree of altruism. For example, one can interpret positive values of the altruism parameter as referring to a couple with above-average altruism.

The left panel of Figure 7 shows that both measures of commitment issues, the rebargaining frequency and the variance ratio, are smaller in households with more altruistic spouses. Even small increases in the degree of altruism compared to the baseline couple reduce commitment issues significantly. Further, couples with sufficiently high altruism can achieve close to perfect intra-household risk sharing.

6.5 Correlation of income shocks

In the baseline parameterization in Section 5.1, it is assumed that income shocks are uncorrelated among spouses. In reality, this within-household correlation differs across households- Some individuals work in the same industry or even the same firm as their respective spouses

Figure 8: Variation of the correlation of income shocks.

(a) Variance ratio Υ

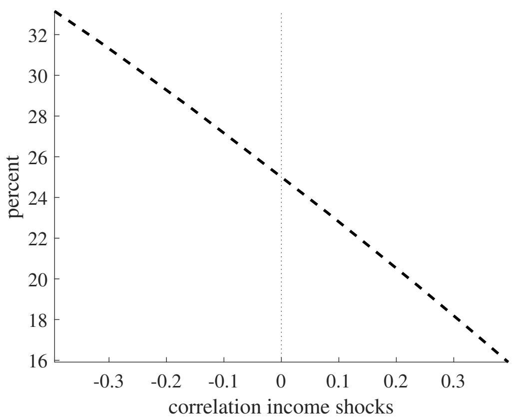

(b) Individual insurance coefficient

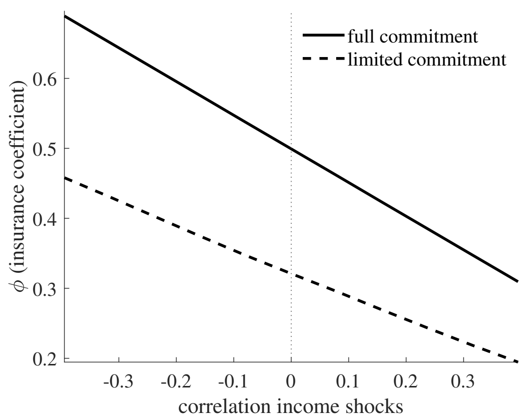

NOTE.–Panel (a) shows the variance ratio Υ as a measure of commitment problems. Panel (b) shows the individual insurance coefficient under full (solid line) and limited commitment (dashed line). Average insurance coefficients are identical for men and women since there is no gender heterogeneity in the model. The vertical dotted line indicates the parameter value in the baseline parameterization.

such that their income shocks will be positively correlated. In other couples, spouses' jobs are more different implying a lower correlation of idiosyncratic income shocks. Taking this into account, we now consider a model populated by couples who differ in terms of the within-household correlation of income shocks.[4]

A positive correlation of income shocks of spouses in a household implies that it becomes more difficult for spouses to insure each other since the probability that *both* spouses have similar income realizations increases. This tends to induce stronger commitment problems because the value of marriage is smaller and the outside option relatively more attractive. However, there is a second effect. A positive correlation of income shocks implies that the probability of large income *differences* between spouses is relatively small. Thus, this effect makes participation constraints bind less often which implies less severe commitment problems when shocks are positively correlated. Which effect dominates is a quantitative question.[5]

Figure 8 plots the measures of commitment problems and the insurance coefficient against the within-household correlation of income shocks. Recall that the model is calibrated to

[4]Hyslop (2001) estimates couple income processes taking into account within-household correlation.

[5]There are analogous effects when income shocks are negatively correlated. One the one hand, negatively correlated shocks can be insured more easily, which tends to reduce commitment issues. On the other hand, the probability of large income differences between spouses increases, which tends to strengthen commitment problems.

generate commitment issues at a correlation of zero. Hence, the numbers on the horizontal axis are to be interpreted as a deviation from the average within-household correlation of income shocks in a household.

Panel (a) in Figure 8 shows that the strength of commitment issues as measured by the variance ratio is smaller where the correlation of income shocks is larger. This implies that the effect of a higher likelihood of binding participation constraints dominates the effect of better insurance opportunities within the couple when income shocks are negatively correlated. This can be seen from panel (b) which shows a comparison of the insurance coefficients under full (solid line) and limited commitment (dashed line). In both cases, insurance is higher when income shocks correlate negatively because a higher correlation implies that one's spouse's income is a worse hedge to one's own income. However, the difference between full and limited commitment is smaller where the correlation of income shocks is low. Hence, couples with a rather high correlation of spouses' incomes are less affected by limited commitment.

6.6 Variance of income shocks

We now consider differences in income volatility across the population, i.e., we distinguish between households that face substantial income risk and household where income risk is moderate.

Higher income volatility has two opposing effects on the strength of commitment issues. Since spouses are risk averse, higher income risks increase the insurance gains from marriage such that intra-household risk-sharing becomes more attractive. Thus, this channel tends to imply less severe commitment issues. A second effect is that, when incomes are more volatile, there are more often situations where spouses have large *differences* in income. In consequence, participation constraints tend to bind more often resulting in stronger commitment issues. Which of these two channels dominates is a quantitative question.

The left panel in Figure 9 shows that the two measures of commitment problems indicate opposite effects of income risk. The variance-ratio measure of commitment problems is larger where income risk is larger. However, the rebargaining frequency is found to be small where income volatility is large. The Kaplan and Violante (2010) insurance coefficient is larger where income risk is higher, see the right panel of Figure 9. Hence, limited commitment affects intra-household risk sharing most strongly in couples that face rather little income risk. However, the differences are relatively small quantitatively.

Figure 9: Impact of higher shock volatility.

(a) Commitment issues (b) Female insurance coefficient

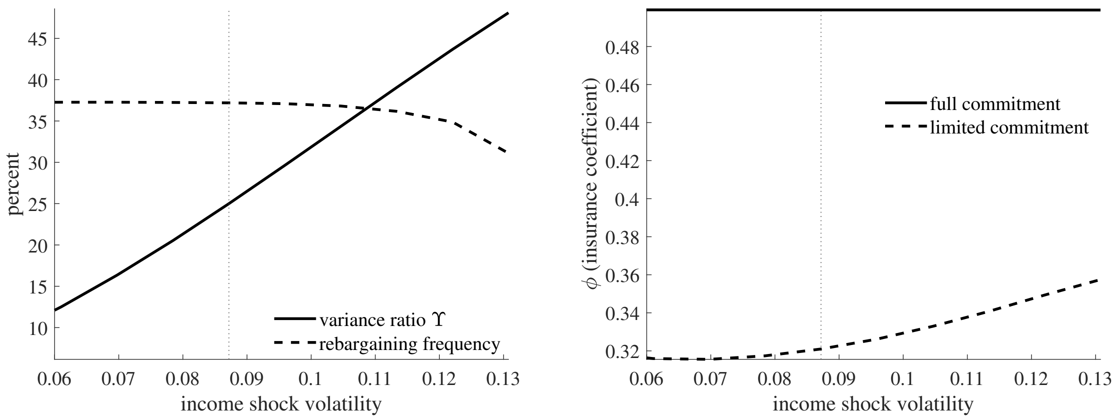

NOTE.–The left panel shows commitment issues, measured by the variance ratio Υ (solid line) and the rebargaining frequency (dashed line), for different values of the standard deviation of income shocks. The right panel shows the female insurance coefficient under full commitment (solid line) and under limited commitment (dashed line). The vertical dotted line indicates the parameter value in the baseline parameterization.

7 Conclusion

Intra-household risk sharing is an important means of insurance against idiosyncratic income shocks when there is full commitment between household members. Limited commitment between household members reduces their ability to insure each other. In this paper, we have quantified how strongly empirically reasonable degrees of limited commitment reduce individuals' possibilities to share risk within the household. To this end, we have developed and calibrated an incomplete-markets model with limited commitment in dual-earner households. Our results show that limited commitment reduces the insurance value of marriage by one fourth to one third. This is associated with a welfare loss equivalent to 0.6% of lifetime consumption.

We have also identified for which types of agents intra-household risk sharing is particularly aggravated by limited commitment. In the cross-section, risk-sharing possibilities are particularly reduced for individuals whose patience or risk aversion is below average while individuals with strong gains from marriage or high degrees of altruism are less affected.

As policy implications, our results suggest an increased need for other forms of insurance such as social safety programs and a rationale for means-testing that targets individuals rather than households.

A Appendix

A.1 The Marcet and Marimon (2019) approach

Marcet and Marimon (2019) provide an approach to rewrite general models with forward-looking constraints in recursive formulation. In general, the occurrence of forward-looking constraints prohibits a recursive formulation since the Bellman equation would become non-stationary. Marcet and Marimon (2019) show that a recursive formulation is nevertheless feasible under fairly standard conditions by introducing additional co-state variables, one for each forward-looking constraint. These co-state variables include information about previous binding forward-looking constraints. A similar approach has been used in the literature to solve models with limited commitment, including the literature on family economics with limited commitment discussed in Section 2. In this Appendix, we show that the recursive formulation of the couple household maximization problem under limited commitment presented in Section 3 can be derived from a more standard formulation of the same problem using the approach of Marcet and Marimon (2019).

To derive the recursive formulation, we start with a more standard notation of the problem. Index t denotes time. The current state is denoted by s_t and implies an income pair $\mathbf{y}(s_t) = (y_m(s_t), y_f(s_t))$. The unconditional probability of a particular state is given by $\pi(s_t)$. The history of events until time t is $s^t = (s_0, ..., s_t)$. The unconditional probability of a particular history of states is $\pi\left(s^t\right)$. The conditional probability of a particular history of states in period r, given the history of states in period t is $\pi\left(s^r \middle| s^t\right)$ with $\pi\left(s^t \middle| s^t\right) = 1$.

The expected lifetime utility of spouse i at household formation is

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t \cdot \pi\left(s^t\right) \cdot \left[u_i\left(c_i\left(s^t\right)\right) + \Psi\right], \text{ for } i = m, f.$$

Given initial Pareto weights[1] $\mu_{m0} = \mu_0$ and $\mu_{f0} = 1 - \mu_0$ for husband and wife, respectively, household lifetime utility is given by

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t \cdot \pi\left(s^t\right) \cdot \left[\mu_0 \cdot u_m\left(c_m\left(s^t\right)\right) + (1 - \mu_0) \cdot u_f\left(c_f\left(s^t\right)\right) + \Psi\right].$$

In both model versions, full commitment and limited commitment, Pareto weights must be

[1] We fix initial Pareto weights for now since the exact determination does not matter for the analysis. However, a derivation of initial Pareto weights can be found in Appendix A.3.

compatible with initial participation constraints

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot\left[u_m\left(c_m\left(s^t\right)\right)+\Psi\right]\geq\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot u_m\left(c_m^{bachelor}\left(s^t\right)\right),$$
$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot\left[u_f\left(c_f\left(s^t\right)\right)+\Psi\right]\geq\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot u_f\left(c_f^{bachelor}\left(s^t\right)\right),$$

with $\pi\left(s^0\middle|\,s^0\right)=1$. Thus, *initial* expected lifetime utility for both spouses must be at least as high within the couple than as a single.

Under limited commitment, there are additional participation constraints, as expected lifetime utility of both spouses must be at least as high within the couple as in the outside option *in any state of nature, after any history of states and in every period of time*:

$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_m\left(c_m\left(s^r\right)\right)+\Psi\right]\geq\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_m\left(c_m^{bachelor}\left(s^r\right)\right),\ \forall s^t,\ \forall t,$$
$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_f\left(c_f\left(s^r\right)\right)+\Psi\right]\geq\sum_{t=0}^{\infty}\sum_{s^r}\beta^{t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_f\left(c_f^{bachelor}\left(s^r\right)\right),\ \forall s^t,\ \forall t,$$

with $\pi\left(s^t\middle|\,s^t\right)=1$.

Using the budget constraint

$$c_m\left(s^t\right)+c_f\left(s^t\right)\leq y_m\left(s_t\right)+y_f\left(s_t\right),$$

the Lagrangian of the limited-commitment maximization problem is

$$\begin{aligned}\mathcal{L}=&\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\left\{\mu_0\cdot u_m\left(c_m\left(s^t\right)\right)+\left(1-\mu_0\right)\cdot u_f\left(c_f\left(s^t\right)\right)+\Psi\right.\\&+\nu_m\left(s^t\right)\cdot\begin{pmatrix}\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_m\left(c_m\left(s^r\right)\right)+\Psi\right]-\\\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_m\left(c_m^{bachelor}\left(s^r\right)\right)\end{pmatrix}\\&+\nu_f\left(s^t\right)\cdot\begin{pmatrix}\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_f\left(c_f\left(s^r\right)\right)+\Psi\right]-\\\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_f\left(c_f^{bachelor}\left(s^r\right)\right)\end{pmatrix}\\&\left.+\lambda\left(s^t\right)\cdot\left(y_m\left(s_t\right)+y_f\left(s_t\right)-c_m\left(s^t\right)-c_f\left(s^t\right)\right)\right\},\end{aligned}$$

where $\nu_i\left(s^t\right)$ denotes the Kuhn-Tucker multiplier on person i's participation constraint and $\lambda\left(s^t\right)$ denotes the multiplier on the budget constraint. The Kuhn-Tucker multipliers ν_m and ν_m are zero whenever the corresponding participation constraint does not bind. In all other cases, the Kuhn-Tucker multipliers are growing over time since even higher multipliers are needed such that participation constraints are fulfilled. In consequence, the Lagrangian is non-stationary.

The main idea of Marcet and Marimon (2019) is to combine elements of this Lagrangian approach in a way such that the maximization problem becomes stationary. The left-hand sides of the participation constraints (the inside options), which are the first term in the brackets in the second and third line of the Lagrangian, can be combined with the weighted utility of the particular spouse in the household target function. Marcet and Marimon (2019) show that the introduction of additional state variables transforms the problem to a stationary one. In particular, Marcet and Marimon (2019) show that a "new" Pareto weight defined according to $\mu_i\left(s^t\right)=\mu_i\left(s^{t-1}\right)+\nu_i\left(s^t\right)$ with $\mu_m\left(s^{-1}\right)=\mu_0$ and $\mu_f\left(s^{-1}\right)=1-\mu_0$, stationarizes the problem. Thus, whenever the participation constraint of spouse i is binding, i.e., $\nu_i\left(s^t\right)>0$, his or her weight in the household target function is increased as little as possible such that the participation constraint holds with equality. Marcet and Marimon (2019) show that the transformed maximization problem has the same solution as the original problem. Using this Marcet and Marimon (2019) transformation, the Lagrangian becomes

$$
\begin{aligned}
\mathcal{L}=&\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\left\{\mu_m\left(s^t\right)\cdot u_m\left(c_m\left(s^t\right)\right)+\mu_f\left(s^t\right)\cdot u_f\left(c_f\left(s^t\right)\right)\right.\\
&+\left(\mu_m\left(s^t\right)+\mu_f\left(s^t\right)\right)\cdot\Psi\\
&-\nu_m\left(s^t\right)\cdot\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_m\left(c_m^{bachelor}\left(s^r\right)\right)\\
&-\nu_f\left(s^t\right)\cdot\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_f\left(c_f^{bachelor}\left(s^r\right)\right)\\
&\left.+\lambda\left(s^t\right)\cdot\left(y_m\left(s_t\right)+y_f\left(s_t\right)-c_m\left(s^t\right)-c_f\left(s^t\right)\right)\right\}.
\end{aligned}
$$

The recursive formulation of this problem is

$$
\begin{aligned}
V\left(\mathbf{y}\left(s^t\right),\mu_m\left(s^{t-1}\right),\mu_f\left(s^{t-1}\right)\right)=&\ \mu_m\left(s^t\right)\cdot u_m\left(c_m\left(s^t\right)\right)+\mu_f\left(s^t\right)\cdot u_f\left(c_f\left(s^t\right)\right)\\
&+\left(\mu_m\left(s^t\right)+\mu_f\left(s^t\right)\right)\cdot\Psi\\
&+\beta\cdot\sum_{s^{t+1}}\pi\left(s^{t+1}\middle|\,s^t\right)\cdot V\left(\mathbf{y}\left(s^{t+1}\right),\mu_m\left(s^t\right),\mu_f\left(s^t\right)\right)\\
&-\nu_m\left(s^t\right)\cdot S_m\left(y_m\left(s^t\right)\right)-\nu_f\left(s^t\right)\cdot S_f\left(y_f\left(s^t\right)\right)\\
&+\lambda\left(s^t\right)\cdot\left(y_m\left(s_t\right)+y_f\left(s_t\right)-c_m\left(s^t\right)-c_f\left(s^t\right)\right).
\end{aligned}
$$

Since in this model, the outside options do not depend on any decision variables of the couple household, they can be skipped from the Lagrangian and the recursive formulation as long as they are considered for the determination of the Kuhn-Tucker multipliers. Thus, the recursive formulation of the maximization problem is finally (the budget constraint is accounted for by restricting the set of feasible choices accordingly):

$$
\begin{aligned}
V\left(\mathbf{y}\left(s^{t}\right),\mu_{m}\left(s^{t-1}\right),\mu_{f}\left(s^{t-1}\right)\right) = {} & \mu_{m}\left(s^{t}\right)\cdot u_{m}\left(c_{m}\left(s^{t}\right)\right)+\mu_{f}\left(s^{t}\right)\cdot u_{f}\left(c_{f}\left(s^{t}\right)\right) \\
& +\left(\mu_{m}\left(s^{t}\right)+\mu_{f}\left(s^{t}\right)\right)\Psi \\
& +\beta\cdot\sum_{s^{t+1}}\pi\left(s^{t+1}\middle|\,s^{t}\right)\cdot V\left(\mathbf{y}\left(s^{t+1}\right),\mu_{m}\left(s^{t}\right),\mu_{f}\left(s^{t}\right)\right).
\end{aligned}
$$

Simplifying notation by omitting the state of nature notation, the recursive formulation becomes

$$
V\left(\mathbf{y},\mu_{m},\mu_{f}\right)=\mu_{m}'\cdot u_{m}\left(c_{m}\right)+\mu_{f}'\cdot u_{f}\left(c_{f}\right)+\left(\mu_{m}'+\mu_{f}'\right)\cdot\Psi+\beta\cdot\mathbb{E}\left[V\left(\mathbf{y}',\mu_{m}',\mu_{f}'\right)\right],
$$

where $\mu_{m}\left(s^{t}\right)=\mu_{m}'$ is notational convention.

Since the optimal allocation in this model depends only on *relative* Pareto weights and couple households do not face any intertemporal decisions, it is possible to normalize Pareto weights such that they sum up to one. In this case, the Bellman equation becomes

$$
V\left(\mathbf{y},\mu\right)=\mu'\cdot u_{m}\left(c_{m}\right)+\left(1-\mu'\right)\cdot u_{f}\left(c_{f}\right)+\Psi+\beta\cdot\mathbb{E}\left[V\left(\mathbf{y}',\mu'\right)\right], \tag{83}
$$

which is identical to (65) in the main text.

A.2 Derivation of the risk-sharing conditions

Full commitment couple household

The full commitment couple household maximizes

$$
V^{fc}\left(\mathbf{y},\mu_{0}^{fc}\right)=\max_{c_{m},c_{f}}\ v\left(c_{m},c_{f},\mu_{0}^{fc}\right)+\beta\cdot\mathbb{E}\left[V^{fc}\left(\mathbf{y}',\mu_{0}^{fc}\right)\right],
$$

subject to the period budget constraints

$$
c_{m}+c_{f}=y_{m}+y_{f},
$$

given the household target function

$$
v\left(c_{m},c_{f},\mu_{0}^{fc}\right)=\mu_{0}^{fc}\cdot u_{m}\left(c_{m}\right)+\left(1-\mu_{0}^{fc}\right)\cdot u_{f}\left(c_{f}\right)+\Psi.
$$

The Lagrangian is

$$
\mathcal{L}=v\left(c_{m},c_{f},\mu_{0}^{fc}\right)+\beta\cdot\mathbb{E}\left[V^{fc}\left(\mathbf{y}',\mu_{0}^{fc}\right)\right]+\lambda\cdot\left(y_{m}+y_{f}-c_{m}-c_{f}\right).
$$

The first-order conditions are

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial c_m} &= \frac{\partial v\left(c_m, c_f, \mu_0^{fc}\right)}{\partial c_m} - \lambda = \mu_0^{fc} \cdot \frac{\partial u_m\left(c_m\right)}{\partial c_m} - \lambda = 0, \\
\frac{\partial \mathcal{L}}{\partial c_f} &= \frac{\partial v\left(c_m, c_f, \mu_0^{fc}\right)}{\partial c_f} - \lambda = \left(1 - \mu_0^{fc}\right) \cdot \frac{\partial u_f\left(c_f\right)}{\partial c_f} - \lambda = 0, \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= y_m + y_f - c_m - c_f = 0.
\end{aligned}$$

Combining the first two first-order conditions yields the risk-sharing condition (63).

Limited commitment couple household

The limited commitment couple household maximization problem is

$$V\left(\mathbf{y}, \mu\right) = \max_{c_m, c_f} \; v\left(c_m, c_f, \mu'\right) + \beta \cdot \mathbb{E}\left[V\left(\mathbf{y}', \mu'\right)\right],$$

subject to the period budget constraints

$$c_m + c_f = y_m + y_f,$$

given the household target function

$$v\left(c_m, c_f, \mu'\right) = \mu' \cdot u_m\left(c_m\right) + \left(1 - \mu'\right) \cdot u_f\left(c_f\right) + \Psi,$$

and the updating rule for the Pareto weights

$$\begin{aligned}
\mu' &= \operatorname{argmin}_{\tilde{\mu}} \left|\tilde{\mu} - \mu\right| \text{such that} \\
u_m\left(c_m\right) + \Psi + \beta \cdot \mathbb{E}\left[V_m\left(\mathbf{y}', \tilde{\mu}\right)\right] &\geq S_m\left(y_m\right) \\
u_f\left(c_f\right) + \Psi + \beta \cdot \mathbb{E}\left[V_f\left(\mathbf{y}', \tilde{\mu}\right)\right] &\geq S_f\left(y_f\right).
\end{aligned}$$

The Lagrangian of the problem is

$$\mathcal{L} = v\left(c_m, c_f, \mu'\right) + \beta \cdot \mathbb{E}\left[V\left(\mathbf{y}', \mu'\right)\right] + \lambda \cdot \left(y_m + y_f - c_m - c_f\right).$$

The first-order conditions are

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial c_m} &= \frac{\partial v\left(c_m, c_f, \mu'\right)}{\partial c_m} - \lambda = \mu' \cdot \frac{\partial u_m\left(c_m\right)}{\partial c_m} - \lambda = 0, \\
\frac{\partial \mathcal{L}}{\partial c_f} &= \frac{\partial v\left(c_m, c_f, \mu'\right)}{\partial c_f} - \lambda = \left(1 - \mu'\right) \cdot \frac{\partial u_f\left(c_f\right)}{\partial c_f} - \lambda = 0, \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= y_m + y_f - c_m - c_f = 0,
\end{aligned}$$

given the updating-rule of the Pareto weight. Combining the first two first-order conditions yields the risk-sharing condition (67):

$$\frac{\partial u_m(c_m)/\partial c_m}{\partial u_f(c_f)/\partial c_f} = \frac{1-\mu'}{\mu'}.$$

A.3 Nash bargaining

We have assumed (initial) Pareto weights, μ_0^{fc} and μ_0, for the analysis of the full commitment and limited commitment couple models in Section 3 which were exogenously given. Thus, the full commitment model was strictly speaking not a collective model but rather a unitary model since the Pareto weight was a parameter and not endogenous. The initial Pareto weights can be endogenized in a household formation step. Household formation takes place before shock in the first period realize. To highlight this, we use the time index 0 for the household formation in the following.

Spouses would only enter in the household if their expected lifetime utility from being married is higher than their corresponding value as a single. Thus, initial Pareto weights $\breve{\mu} = \{\mu_0^{fc}, \mu_0\}$ are determined such that the initial participation constraints

$$\mathbb{E}_0\mathcal{V}_m(\mathbf{y}, \breve{\mu}) \geq \mathbb{E}_0 S_m(y_m), \tag{84}$$

$$\mathbb{E}_0\mathcal{V}_f(\mathbf{y}, \breve{\mu}) \geq \mathbb{E}_0 S_f(y_f), \tag{85}$$

with $\mathcal{V}_i(\mathbf{y}, \breve{\mu}) = \left\{ V_i^{fc}(\mathbf{y}, \breve{\mu}), V_i(\mathbf{y}, \breve{\mu}) \right\}$, are fulfilled.

Several approaches are used in the literature to determine initial Pareto weights (see Browning, Chiappori, and Weiss 2014). While the collective approach of Chiappori (1988) does not require a particular bargaining solution, several studies use a bargaining approach to determine the initial weights, starting with Manser and Brown (1980) and McElroy and Horney (1981). The vast majority of these studies use the Nash (1950) bargaining solution.[2] We follow this literature and determine $\breve{\mu}$ using the symmetric Nash program

$$NP = \text{argmax}_{\breve{\mu}} \; [\mathbb{E}_0 (V_m(\mathbf{y}, \breve{\mu}) - S_m(y_m))]^{0.5} \cdot [\mathbb{E}_0 (V_f(\mathbf{y}, \breve{\mu}) - S_f(y_f))]^{0.5}, \tag{86}$$

The resulting first-order condition is

$$\frac{\mathbb{E}_0 (V_m(\mathbf{y}, \breve{\mu}) - S_m(y_m))}{\mathbb{E}_0 (V_f(\mathbf{y}, \breve{\mu}) - S_f(y_f))} = -\frac{\mathbb{E}_0 (\partial V_m(\mathbf{y}, \breve{\mu})/\partial \breve{\mu})}{\mathbb{E}_0 (\partial V_f(\mathbf{y}, \breve{\mu})/\partial \breve{\mu})}. \tag{87}$$

[2] The Kalai and Smorodinsky (1975) bargaining solution and the Egalitarian bargaining solution of Kalai (1977) are alternatives but have been used rarely.

Equations (86) and (87) can be used to determine $\breve{\mu}$ numerically.

A.4 Solution algorithm bachelor household

Solving the limited commitment household model requires in Step 2 of Algorithm 2 the solution of the bachelor household model to check the participation constraints in Step 6. Algorithm 3 summarizes the numerical solution of the bachelor household maximization problem.

Algorithm 3 (Bachelor households) *The algorithm consists of the following two steps:*

1: Set a convergence criterion ζ. Discretize the exogenous income processes using the approach of Tauchen (1986).

2: Choose initial guesses for the value functions $S_{m,0}(y_m)$, and $S_{f,0}(y_f)$. The optimal policy functions are $\mathcal{C}_i^S(y_i) = y_i$, $i = m, f$.

2a: Given that you are in iteration step j, use the Bellman equation (54) to get:

$$S_{m,j+1}(y_m) = u_m(y_m) + \beta \cdot \mathbb{E}\left[S_{m,j}\left(y'_m\right)\right], \text{ and}$$
$$S_{f,j+1}(y_f) = u_f(y_f) + \beta \cdot \mathbb{E}\left[S_{f,j}\left(y'_f\right)\right].$$

2b: If $\sup|S_{m,j+1}(y_m) - S_{m,j}(y_m)| < \zeta$, and $\sup|S_{f,j+1}(y_f) - S_{f,j}(y_f)| < \zeta$ stop. Otherwise set $j = j+1$ and go back to Step 21a. Iterate until convergence.

A.5 Heterogeneous risk aversion

In this appendix, we investigate analytically how the Kaplan and Violante (2010) insurance coefficients depend on individuals' risk aversion. This is helpful to isolate the effects that result from the efficient allocation of risk between spouses on the one hand and from changes in the strength of commitment problems on the other hand.

The starting point is the risk-sharing condition (67):

$$\frac{\partial u_m(c_m)/\partial c_m}{\partial u_f(c_f)/\partial c_f} = \frac{1-\mu'}{\mu'}.$$

With CRRA utility and gender-specific parameters σ_m and σ_f, one obtains

$$c_f = \left(\frac{\mu'}{1-\mu'}\right)^{-1/\sigma_f} \cdot c_m^{\sigma_m/\sigma_f}.$$

In general, the risk-sharing condition of spouse i is

$$c_i = \left(\frac{1-\mu'_i}{\mu'_i}\right)^{-1/\sigma_i} \cdot c_{-i}^{\sigma_{-i}/\sigma_i},$$

with $i = \{m, f\}$. $-i$ denotes the spouse of person i. The Pareto weights are given by $\mu'_m = \mu'$ and $\mu'_f = 1 - \mu'$.

Insurance coefficients. We now derive the Kaplan and Violante (2010) insurance coefficients from the risk-sharing condition. Taking logs, one obtains

$$\log c_i = -\frac{1}{\sigma_i} \cdot \log\left(\left(\frac{1-\mu'_i}{\mu'_i}\right)\right) + \frac{\sigma_{-i}}{\sigma_i} \cdot \log c_{-i}.$$

Using the relative Pareto weight of person i

$$x'_i = \frac{\mu'_i}{1-\mu'_i},$$

the risk-sharing condition is

$$\begin{aligned}\log c_i &= -\frac{1}{\sigma_i} \cdot \log \frac{1}{x'_i} + \frac{\sigma_{-i}}{\sigma_i} \cdot \log c_{-i} \\ &= \frac{1}{\sigma_i} \cdot \log x'_i + \frac{\sigma_{-i}}{\sigma_i} \cdot \log c_{-i}.\end{aligned}$$

Consumption growth of person i is then given by

$$\Delta \log c'_i = \frac{\sigma_{-i}}{\sigma_i} \cdot \Delta \log c'_{-i} + \frac{1}{\sigma_i} \cdot \Delta \log \log x''_i.$$

The case of full commitment is nested and corresponds to the case where Pareto weights are fixed over time:

$$\Delta \log c'_i = \frac{\sigma_{-i}}{\sigma_i} \cdot \Delta \log c'_{-i},$$

using $x''_i = x'_i = x^{fc}_{i,0}$.

Inserting the expression for consumption growth in the definition of the Kaplan and Violante (2010) insurance coefficients yields

$$\phi_i = 1 - \frac{\operatorname{cov}\left(\Delta \log c'_i, \varepsilon_i\right)}{\operatorname{var}\left(\varepsilon_i\right)} = 1 - \frac{\operatorname{cov}\left(\frac{\sigma_{-i}}{\sigma_i} \cdot \Delta \log c'_{-i} + \frac{1}{\sigma_i} \cdot \Delta \log x''_i, \varepsilon_i\right)}{\operatorname{var}\left(\varepsilon_i\right)}.$$

The covariance rule

$$\operatorname{cov}\left(aX + bY, Z\right) = a \cdot \operatorname{cov}\left(X, Z\right) + b \cdot \operatorname{cov}\left(Y, Z\right),$$

implies

$$
\begin{aligned}
\phi_i &= 1 - \frac{\operatorname{cov}\left(\frac{\sigma_{-i}}{\sigma_i} \cdot \Delta \log c'_{-i} + \frac{1}{\sigma_i} \cdot \Delta \log x''_i, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} \\
&= 1 - \frac{\sigma_{-i}}{\sigma_i} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} - \frac{1}{\sigma_i} \cdot \frac{\operatorname{cov}\left(\Delta \log x''_i, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)}.
\end{aligned}
\tag{88}
$$

Full commitment. In the case of full commitment, the growth rate of the relative Pareto weight of person i, $\Delta \log x'_i$, is equal to zero. Thus, $\operatorname{cov}(\Delta \log x''_i, \varepsilon_i) = 0$ and the insurance coefficient is

$$
\phi_i = 1 - \frac{\sigma_{-i}}{\sigma_i} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)}.
$$

If both partners have identical degrees of risk aversion, $\sigma_m = \sigma_f$, the insurance coefficients are independent of risk aversion, since *cov* $\left(\Delta \log c'_{-i}, \varepsilon_i\right) / \operatorname{var}(\varepsilon_i)$ is symmetric across gender.

If spouses differ with respect to risk aversion, insurance coefficients will depend on risk aversion also under full commitment. The reaction of the insurance coefficient to changes in person i's risk aversion is under full commitment given by

$$
\frac{\partial \phi_i}{\partial \sigma_i} = \frac{\sigma_{-i}}{\sigma_i^2} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} > 0.
$$

Note that, under full commitment, $\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right) > 0$, since an increase (decrease) of household income implies that both spouses consume more (less). Thus, the insurance coefficient of person i is strictly increasing in person i's risk aversion. Moreover, the insurance coefficient of person i is decreasing in person $-i$'s risk aversion,

$$
\frac{\partial \phi_i}{\partial \sigma_{-i}} = -\frac{1}{\sigma_i} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} < 0.
$$

Limited commitment. Under limited commitment, the insurance coefficient is given by the complete expression in (88). Hence, the reaction of the insurance coefficient to changes in person i's risk aversion is under full commitment given by

$$
\frac{\partial \phi_i}{\partial \sigma_i} = \frac{\sigma_{-i}}{\sigma_i^2} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} + \frac{1}{\sigma_i^2} \cdot \frac{\operatorname{cov}\left(\Delta \log x''_i, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)}.
$$

Therefore, the insurance coefficient of person i increases in σ_i if

$$
\sigma_{-i} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} + \frac{\operatorname{cov}\left(\Delta \log x''_i, \varepsilon_i\right)}{\operatorname{var}(\varepsilon_i)} > 0.
\tag{89}
$$

The next derivative shows how person i's insurance coefficient reacts to changes in person $-i$'s risk aversion:

$$\frac{\partial \phi_i}{\partial \sigma_{-i}} = -\frac{1}{\sigma_i} \cdot \frac{\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right)}{\operatorname{var}\left(\varepsilon_i\right)}.$$

The insurance coefficient of person i decreases in σ_{-i} as long as $\operatorname{cov}\left(\Delta \log c'_{-i}, \varepsilon_i\right) > 0$.

Thus, under limited commitment, there can be additional effects that weaken the relation between person i's risk aversion and person i's insurance coefficient. To understand these effects, consider for example the female insurance coefficient, $i = f$. There are, in general, three scenarios for the covariance between male consumption growth and the female income shock. First, $\operatorname{cov}(\Delta \log c'_m, \varepsilon_f) < 0$ implies that a positive (negative) female income shock decreases (increases) male consumption. Under limited commitment, female intra-household decision power increases weakly with her income. Thus, if there is a positive income shock, her Pareto weight tends to increase (or remain unchanged) and the male Pareto weight should thus tend to *decrease*. The first scenario occurs when, on average, income shocks are in a range such that, after a positive (negative) female income shock, the male Pareto weight changes in a way that the husband receives on average less (more) consumption than before. Second, $\operatorname{cov}(\Delta \log c'_m, \varepsilon_f) > 0$ implies that a positive (negative) female income shock tends to increase (decrease) male consumption. This can only occur when, on average, income shocks are in a range such that both spouses receive higher consumption independent of a potential rebargaining. In the third scenario, $\operatorname{cov}(\Delta \log c'_m, \varepsilon_f) = 0$, female income shocks and male consumption growth do not co-vary as both effects cancel out.

Now consider the second term on the left-hand side of (89). The covariance between growth of the female relative Pareto weight and the female income shock is larger than or equal to zero. In a limited-commitment model, the female relative Pareto weight x may either increase after a female income shock, holding male income fixed, or may remain at its previous period value. The covariance can only be equal to zero if wifes have either never an incentive to leave the household or when spouses can fully commit. Thus, we have $\operatorname{cov}\left(\Delta \log x''_f, \varepsilon_f\right) > 0$. In summary, whether condition (89) holds or not is ultimately a quantitative question. The right panel in Figure 6 shows that condition (89) is met in the simulated model.

A.6 Model with altruism

The full commitment couple household maximizes (61) subject to the period budget constraints (60) with

$$v\left(c_m, c_f, \mu_0^{fc}\right) = \left(\mu_0^{fc} + \varrho_f \cdot \left(1 - \mu_0^{fc}\right)\right) \cdot u_m\left(c_m\right) + \left(1 - \mu_0^{fc} + \varrho_m \cdot \mu_0^{fc}\right) \cdot u_f\left(c_f\right) +$$
$$\left(1 + \varrho_f \cdot \left(1 - \mu_0^{fc}\right) + \varrho_m \cdot \mu_0^{fc}\right) \cdot \Psi.$$

The Bellman equation is

$$V^{fc}\left(\mathbf{y}, \mu_0^{fc}\right) = \max_{c_f} \ \left(\mu_0^{fc} + \varrho_f \cdot \left(1 - \mu_0^{fc}\right)\right) \cdot u_m\left(y_m + y_f - c_f\right) + \left(1 - \mu_0^{fc} + \varrho_m \cdot \mu_0^{fc}\right) \cdot u_f\left(c_f\right) +$$
$$\left(1 + \varrho_f \cdot \left(1 - \mu_0^{fc}\right) + \varrho_m \cdot \mu_0^{fc}\right) \cdot \Psi + \beta \cdot \mathbb{E}\left[V^{fc}\left(\mathbf{y}', \mu_0^{fc}\right)\right].$$

The first-order conditions imply the intra-household risk-sharing condition

$$\left(\mu_0^{fc} + \varrho_f \cdot \left(1 - \mu_0^{fc}\right)\right) \cdot \frac{\partial u_m\left(c_m\right)}{\partial c_m} = \left(1 - \mu_0^{fc} + \varrho_m \cdot \mu_0^{fc}\right) \cdot \frac{\partial u_f\left(c_f\right)}{\partial c_f}.$$

The limited commitment couple household problem features state-dependent Pareto weights. Utility is given by

$$v\left(c_m, c_f, \mu'\right) = \left(\mu' + \varrho_f \cdot \left(1 - \mu'\right)\right) \cdot u_m\left(c_m\right) + \left(1 - \mu' + \varrho_m \cdot \mu'\right) \cdot u_f\left(c_f\right) +$$
$$\left(1 + \varrho_f \cdot \left(1 - \mu'\right) + \varrho_m \cdot \mu'\right) \cdot \Psi,$$

and the risk-sharing condition is

$$\left(\mu' + \varrho_f \cdot \left(1 - \mu'\right)\right) \cdot \frac{\partial u_m\left(c_m\right)}{\partial c_m} = \left(1 - \mu' + \varrho_m \cdot \mu'\right) \cdot \frac{\partial u_f\left(c_f\right)}{\partial c_f}.$$

Chapter IV

Estimating Labor-Supply Elasticities when Commitment between Spouses is Limited

1 Introduction

The Frisch elasticity of labor supply is an important concept in both labor economics and macroeconomics. It measures how willingly individuals substitute hours worked intertemporally. The Frisch elasticity determines the labor-supply responses to temporary changes in hourly wage rates and to predictable life-cycle patterns in wage rates. It is also a decisive determinant of the cost of business cycles and the fiscal multiplier. Despite its importance, there is no consensus about a range of values for the Frisch elasticity with estimates from microeconometric studies diverging substantially from the results of macroeconomic studies, see, e.g., Keane and Rogerson 2015.

The Frisch elasticity is defined as the reaction of labor supply to changes in wage rates holding marginal utility of wealth constant. Hence, it can be estimated in a regression of hours worked on wage rates when controlling for consumption, the latter being closely tied to the marginal utility of wealth, see, e.g., Altonji (1986). However, the relevant consumption variable would be consumption of the individual worker and this information is usually not observed in household panel surveys which are needed to account for unobserved heterogeneity in preferences for work. Instead, the consumption information (if available at all) provided by most household panel surveys relates to the level of the household.[1] Household consumption may be insufficient for a consistent estimation of the Frisch elasticity because most workers live in families[2] and there are strong theoretical and empirical arguments that consumption of a household can develop quite differently from consumption of its individual members. Recent papers in family economics have referred to models of limited commitment to explain differential dynamics of individual and household consumption. There is strong empirical

[1] To cope with a number of issues regarding consumption data, the seminal paper by Altonji (1986) proposed to estimate the labor supply condition in growth rates and to identify the Frisch elasticity as the coefficient on expected wage growth. The latter is uncorrelated with changes in marginal utility when households have sufficient access to insurance or borrowing possibilities. In empirical applications, the Altonji (1986) approach faces two challenges. First, consumption insurance possibilities are limited in the real world and a substantial share of households has close to no wealth (see Kuhn and Ríos-Rull 2016) reducing their ability to self-insure. Domeij and Floden (2006) have shown that this results in a substantial downward bias in the estimated Frisch elasticity. Bredemeier et al. (2019) have shown that this challenge can be overcome by exploiting the double-earner structure of most households and have derived an estimation procedure that yields an unbiased estimate of the Frisch elasticity in presence of borrowing constraints. The second challenge is that the Altonji (1986) approach requires information on expected wage growth which is usually obtained using instrumental variables. However, wage growth – in contrast to wage levels – is hard to predict and hence instruments tend to be weak, see Keane (2011) for a discussion.

[2] In 2019, 98 million labor-force participants in the U.S. are part of a married couple (84 million) or of an opposite-sex unmarried couple (14 million). This amounts to almost 77% of the U.S. labor force. Numbers stem from the US Census Bureau, America's Families and Living Arrangements 2019, Tables FG1 and UC1.

evidence for this kind of household behavior, see Dercon and Krishnan (2000), Mazzocco (2007), Robinson (2012), Cesarini et al. (2017) and Lise and Yamada (2019). Because commitment is limited, individual household members' shares in total household consumption vary and reflect members' relative bargaining positions.

In this paper, we analyze the consequence of limited commitment in the family for the estimation of labor-supply elasticities. As intra-household bargaining positions depend on income potentials and thus on wage rates, proxying individual consumption by household consumption in a labor-supply regression yields biased estimates of the Frisch elasticity. We study the quantitative importance of estimation biases arising from limited commitment, develop new estimation approaches that yield unbiased estimates of the Frisch elasticity and apply these approaches to panel data from the United States.[3] In a first part of the paper, we build a quantitative model of decision making with limited commitment between spouses in dual-earner households which we use to document estimation biases arising from limited commitment and to develop improved estimation procedures. In such a model, a rise in an individual's hourly wage rate exerts three effects on the individuals' labor supply. First, labor supply increases due to a conventional substitution effect that is governed by the Frisch elasticity and can be used to recover this elasticity. Second, household consumption increases which induces labor supply to fall due to a wealth effect. Third, the individual's bargaining position in the household may increase reflecting the improvement of its outside option. This leads the household to grant more leisure to the individual thereby reducing its labor supply through a limited-commitment effect.[4] Both the income effect and the limited-commitment effect could be accounted for with data on individual consumption, but data on household consumption can only account for the income effect but not the limited-commitment effect.

We study the impact of this limited-commitment effect on labor-supply estimations in two ways. First, we exploit first-order conditions of our model to show analytically that conventional estimation procedures yield biased results. Second, we run Monte-Carlo experiments

[3] Specifically, we use the Panel Study of Income Dynamics (PSID). Recent waves of the PSID include more detailed information on household consumption than the waves used by, e.g., Altonji (1986) where researchers were restricted to using food expenditures as a proxy for consumption. Still, also in recent waves of the PSID, consumption information is provided on the level of the household. Blundell et al. (2016) and Blundell et al. (2018) use these data to estimate the parameters of life-cycle models with intra-household risk sharing, including labor-supply elasticities. Other papers that have used the consumption data from the PSID are, e.g., Attanasio and Pistaferri (2014), Hokayem and Ziliak (2014), Andreski et al. (2014), and Charles et al. (2014).

[4] While for some applications, the total effect may be a sufficient statistic, other applications require disentangling changes in labor supply due to intertemporal substitution from changes due to wealth effects or changes in bargaining powers, i.e., to identify the Frisch elasticity. For example, in many macroeconomic contexts, one is interested in the reaction of labor supply to temporary changes in aggregate wage rates. This reaction is mostly driven by intertemporal substitution as the wealth effect is small and outside options are affected rather symmetrically.

to document the quantitative importance of these biases using simulated model data from our quantitative model. Estimating a standard labor-supply condition may yield an estimated Frisch elasticity that is up to 30% below its true value when household consumption rather than individual consumption is controlled for.

We then use our model to develop improved estimation procedures that can be applied when data on individual consumption is not available. The key to obtaining an unbiased estimate of the Frisch elasticity in a limited-commitment context without having access to individual consumption data is to use information on other household member's behavior to assess bargaining positions: In dual-earner households, hours worked and wage rates of ones spouse convey information about relative bargaining positions.

A first approach is to use information on hours worked and the wage rate of the spouse as additional controls in a labor-supply regression. We use our model to show analytically that this approach yields almost unbiased estimates when these variables are included in sufficiently flexible form. We use Monte-Carlo experiments to quantify the bias reduction achieved by this approach and find that it can reduce the bias due to limited commitment by more than three quarters when a second-order polynomial is used. Second, we develop an iterative procedure. In this approach, we deduce spouses' relative bargaining positions from relative labor supply and relative wage rates based on a guess of the Frisch elasticity and other parameters, then update this guess in a labor-supply regression where the deduced bargaining position is added as an additional control. This procedure is iterated until convergence. We show in Monte-Carlo experiments that, while more complicated than the first approach, this procedure converges to an unbiased estimate of the Frisch elasticity.

In the empirical part of the paper, we apply our approaches to panel data from the Panel Study of Income Dynamics. The results are strongly supportive of our theoretical results. Applying our improved estimation approaches, we find estimates of the Frisch elasticity of around 0.7. Our estimates are hence considerably larger than the estimates of many other microeconometric studies where the elasticity of labor supply is estimated to be essentially zero (see, for example, Browning et al. 1985, Blundell and Walker 1986, and Ziliak and Kniesner 1999). Thus, controlling for consumption and correcting for biases due to limited commitment between spouses can contribute to reconciling microeconometric evidence on the Frisch elasticity with the results of macroeconomic model evaluations that typically deliver larger Frisch elasticities.

This chapter contributes to both, the literature on estimating labor-supply elasticities and the limited-commitment couple-household literature. Keane (2011) provides a comprehensive survey of the literature on estimating labor-supply elasticities. Specifically, the micro/macro puzzle on the elasticity of labor supply is a central question in this literature, see Keane and Rogerson (2015) for an overview. Several studies have contributed to solving the micro/macro puzzle by pointing out a number of downward biases in microeconometric estimates, see, e.g., Blomquist (1985, 1988),Alogoskoufis (1987), Heckman (1993), Rupert et al. (2000), Imai and Keane (2004), and Domeij and Floden (2006). This chapter adds to this by showing that limited commitment between spouses – if not appropriately corrected for – also biases estimates of the Frisch elasticity downward.

The literature on limited commitment between spouses in marriages has been surveyed by Chiappori and Mazzocco (2017). A large part of this literature documents empirical evidence for limited commitment, i.e., rejects the competing full-commitment paradigm, see, e.g., Mazzocco 2007, Robinson 2012, and Lise and Yamada 2019. A second strand of the literature studies the consequences of limited commitment, see, e.g., Mazzocco et al. (2014), Voena 2015, and Ábrahám and Laczó 2018. While most papers in this second group focus on welfare or policy effects, we study the consequence of limited commitment for labor-supply estimations.

The remainder of this chapter is organized as follows. We develop a limited-commitment couple-household model in Section 2. In Section 3, we derive estimation approaches analytically. In Section 4, we perform Monte-Carlo experiments to evaluate our estimation approaches using synthetic data from a simulation of our calibrated model. In Section 5, we provide an empirical application of our estimation approaches using PSID data. Section 6 concludes.

2 A dual-earner household model with limited commitment

We consider an economy which is populated by a continuum of agents of mass 2. The population is equally divided into men and women. We distinguish between agents (individual persons) and households which consist of two agents. Agents differ by a deterministic wage component (fixed effect) and by an idiosyncratic stochastic component. Wage rates are exogenous in the model. Households differ from one another by the wage fixed effects of household members, by the idiosyncratic productivities of household members, and by direct utility gains from marriage. Intra-household decision-making of couples is characterized by limited commitment among spouses. Specifically, this means that agent's behavior will react to changes in their relative bargaining positions that occur during the household's life. This is responsible for the bias in conventional labor-supply estimations as wage raises improve bargaining positions which tends to increase leisure time and reduce labor supply. Spouses are subject to joint budget constraints. Throughout, variables with a prime denote next period values.

Spouses bargain over how much each of them works and consumes in different states of the world. Bargaining positions are determined by spouses' outside options which we take to be the lifetime utility of single life. Importantly, we take into account that either spouse can leave the household at his or her own choice (unilateral divorce) whenever the outside option is preferable to life within the household. This approach follows, among others, Mazzocco (2007), Voena (2015), and Ábrahám and Laczó (2018). Spouses acknowledge this possibility in their bargaining process and understand that they have no possibility to enforce promises made by the partner. Put differently, commitment to promises made to the partner is limited. As a consequence, any bargained plan over consumption and labor supply has to ensure that partners never actually want to leave the household. This can be understood as a self-enforcing contract[1] between the spouses as, under the given plan, incentives are always as such that spouses want to stay in the household - or, more technically, that participation constraints hold at any time.

We make use of the results of Nash (1950) and Marcet and Marimon (2019) which imply that such a problem can be written as the problem of a Pareto planner with time-varying

[1] In fact, the limited-commitment couple-household literature uses the methods developed by the literature on self-enforcing recursive contracts. Several studies have used the same general approach to study topics from development economics, political economy, international economics, and labor economics. See Golosov et al. (2016) for an introduction into the theoretical backgrounds and for an overview of the literature.

weights on individual utilities wherein changes in individual Pareto weights reflect the bindingness of the respective participation constraint. This formulation of the household decision problem has an intuitive interpretation as one can understand changes in Pareto weights as the household adjusting (or "rebargaining") the weight of an individual member when this member would otherwise want to leave the household.

Preferences. Agents receive utility from consumption of goods and services c, and disutility from labor n. An individual agent is denoted by ij, where $i = \{m, f\}$ denotes the gender of agent ij and j denotes the household. Individual preferences of agent ij are represented by the concave utility function

$$u_{ij}\left(c_{ij}, n_{ij}\right) = \frac{c_{ij}^{1-\sigma} - 1}{1-\sigma} - \alpha_{ij} \cdot \frac{n_{ij}^{1+1/\eta_i}}{1+1/\eta_i} + \Psi_j,$$

when married, and by

$$u_{ij}\left(c_{ij}, n_{ij}\right) = \frac{c_{ij}^{1-\sigma} - 1}{1-\sigma} - \alpha_{ij} \cdot \frac{n_{ij}^{1+1/\eta_i}}{1+1/\eta_i},$$

if ij is single or divorced. $1/\sigma$ measures the elasticity of intertemporal substitution of consumption, α_{ij} is the individual weight on disutility of labor, and η_i is the Frisch labor-supply elasticity which may depend on gender. Agents discount future utility with the discount factor $\beta \in (0, 1)$. The direct utility gain from being married, Ψ_j, captures all gains and losses from marriage not explicitly modeled such as love, companionship, intimacy, children, returns to scale in consumption, home production, etc. Since we model being divorced as an absorbing state, the utility gain Ψ_j also includes the negative value of remarriage possibilities as a divorcee. Direct utility gains from marriage are household-specific and differences in this parameter between households may reflect match quality. In Section 4.1, we calibrate this parameter to match the extent of commitment issues in our model to an empirical target.

Wage rates and budget constraint. Market work, n, is rewarded by the real wage rate w measured in units of the market consumption good c which is the numéraire. Individual wage rates consist of three components: a fixed effect ψ_{ij}, a gender-specific effect ν_i, and a stochastic component z_{ij}:

$$\log w_{ij} = \log \psi_{ij} + \log \nu_i + z_{ij}. \tag{90}$$

The vector of joint wage rate realizations is denoted by $\omega_j = (w_{mj}, w_{fj})$. The budget constraint of the couple is

$$c_{mj} + c_{fj} \leq w_{mj} n_{mj} + w_{fj} n_{fj}. \tag{91}$$

We do not specify a particular wage process at this point since our analytical results in Section 3 do not require a particular specification of the wage process. In our quantitative analysis in Section 4, we assume that the stochastic wage component follows a stationary first-order autoregressive process (AR).

Intra-household decision-making under limited commitment. Intra-household decision-making is characterized by limited commitment as in Mazzocco (2007), Voena (2015), and Ábrahám and Laczó (2018). Following Ábrahám and Laczó (2018), we assume that the outside option of a spouse is being a single for the rest of her life.[2] Formally, the participation constraints of the individual spouses are

$$V_{ij}\left(\omega_j, \mu_{ij}, \mu_{-ij}\right) \geq S_{ij}\left(w_{ij}\right), \tag{92}$$

where V_{ij} denotes the expected lifetime utility of agent ij within the couple household, $-ij$ is individual ij's spouse, and S_{ij} denotes the expected lifetime utility when being single.[3] The behavior of the couple under limited commitment is ex-post efficient such that the household target function can be summarized as a weighted sum of individual utility functions, with Pareto weights μ_{mj} and μ_{fj}, respectively, see Mazzocco (2007). We assume that initial Pareto weights $\mu_{mj,0}$ and $\mu_{fj,0}$ are determined by Nash bargaining.

Compared to a situation with full commitment, where participation constraints must be only fulfilled once at the time of household formation, couples which are subject to limited commitment face the participation constraints as additional constraints in every period and in any state of nature. As a consequence, the maximization problem of the couple household becomes non-stationary. Marcet and Marimon (2019) show that the non-stationary maximization problem with participation constraints of the spouses can be transformed into a stationary problem where the Pareto weights of the spouses become additional time-varying state variables.[4] Individual Pareto weights react to changing outside options of the partners as follows. Whenever one spouse has an incentive to leave the household, i.e., his or her participation constraint is binding, the couple increases his or her Pareto weight such that his or her participation constraint holds with equality. This is achieved by adjusting the Pareto weight of the spouse by exactly the Kuhn-Tucker multiplier on the respective participation

[2]This assumption is for computational simplicity. A more elaborated modeling of the outside option would yield the same implications in our context. Specifically, giving divorcees the possibility to remarry (as in Voena (2015) would raise the value of the outside option and lead us to recalibrate the direct utility value from marriage, Ψ, accordingly.

[3]The maximization problem of the single household can be found in Appendix A.1.1.

[4]A detailed derivation of the problem can be found in Appendix A.1.2.

constraint (see Marcet and Marimon 2019).[5]

We formulate the household target function as

$$v\left(c_{mj}, c_{fj}, n_{mj}, n_{fj}, \mu'_{mj}, \mu'_{fj}\right) = \mu'_{mj} \cdot u_{mj}\left(c_{mj}, n_{mj}\right) + \mu'_{fj} \cdot u_{fj}\left(c_{fj}, n_{fj}\right).$$

The prime notation of the Pareto weights highlights that given the Pareto weights from the previous period as state variables, couples update the Pareto weights and thereby the relative Pareto weight whenever participation constraints are binding.

To summarize, the results of Nash (1950) and Marcet and Marimon (2019) allow to formulate the household decision problem as a simple Pareto planning problem where the couple household maximizes

$$V_j\left(\omega_j, \mu_{mj}, \mu_{fj}\right) = \max_{c_{mj}, c_{fj}, n_{mj}, n_{fj}} v\left(c_{mj}, c_{fj}, n_{mj}, n_{fj}, \mu'_{mj}, \mu'_{fj}\right) + \beta \cdot \mathbb{E}\left[V\left(\omega'_j, \mu'_{mj}, \mu'_{fj}\right)\right], \tag{93}$$

with the updating rule for the Pareto weights

$$\mu'_{ij} = \mu_{ij} + \phi_{ij},$$

where ϕ_{ij} is the Kuhn-Tucker multiplier on the participation constraint (92), subject to the period budget constraints (91). $\mathbb{E}$ denotes the expectation operator which includes all information in the particular period.

The first-order conditions are

$$\lambda_j = \mu'_{mj} \cdot c_{mj}^{-\sigma} = \mu'_{fj} \cdot c_{fj}^{-\sigma}, \tag{94}$$

$$w_{mj} = -\mu'_{mj} \cdot \frac{n_{mj}^{1/\eta_m}}{\lambda_j}, \tag{95}$$

$$w_{fj} = -\mu'_{fj} \cdot \frac{n_{fj}^{1/\eta_f}}{\lambda_j}, \tag{96}$$

and the budget constraint (91), given the vector of wage rate realizations ω_j and the previous period Pareto weights, μ_{mj} and μ_{fj}. λ_j denotes the Lagrange multiplier on the budget constraint (91). The ratio of marginal utilities of consumption of the spouses may vary over time due to potential time variation in the Pareto weights which can be seen from the risk-sharing condition (94). Thus, intra-household risk-sharing is imperfect since commitment is

[5] In the numerical implementation, we use that the problem can be further simplified such that the sole additional state variable is the *relative* Pareto weight of one of the spouses, e.g., the male. In doing so, we adapt the numerical approach of Ábrahám and Laczó (2018) to our model. Whenever the participation constraint of one spouse is binding, the Pareto weight of this spouse increases which implies a change in the relative Pareto weight of the male.

limited. The labor-supply conditions are given by (95) and (96). Also these static first-order conditions reflect spouses' varying bargaining powers μ'_{ij}. We use these first-order conditions to derive estimation equations under limited commitment in Section 3.

Full commitment. As a benchmark to document estimation biases originating from limited commitment, we also consider the special case where spouses can fully commit. In particular, we consider specifications where estimates of the Frisch labor-supply elasticity are biased under limited commitment while they are unbiased under full commitment. In the full-commitment case, Pareto weights are determined once and for all such that participation constraints at the time of household formation are fulfilled. We assume that Pareto weights are determined by a Nash bargaining solution as in the case of limited commitment. Since Pareto weights are fixed under full commitment, we have $\mu'_{ij} = \mu_{ij} = \bar{\mu}_{ij}$. The resulting first-order conditions under full commitment are:

$$\lambda_j = \bar{\mu}_{mj} \cdot c_{mj}^{-\sigma} = \bar{\mu}_{fj} \cdot c_{fj}^{-\sigma}, \tag{97}$$

$$w_{mj} = -\bar{\mu}_{mj} \cdot \frac{n_{mj}^{1/\eta_m}}{\lambda_j} \ , \tag{98}$$

$$w_{fj} = -\bar{\mu}_{fj} \cdot \frac{n_{fj}^{1/\eta_f}}{\lambda_j} \ . \tag{99}$$

Note that the first-order conditions differ from the limited-commitment model only in that the Pareto weights are constant. This has important implications for the consistency of labor-supply estimations as discussed in the next section.

3 Estimation approaches to recover the Frisch elasticity of labor-supply

In this section, we use the first-order conditions of the model to derive conditions which can be used to estimate the Frisch elasticity of labor supply. Whether the estimations yield unbiased estimates depends on the availability of data on individual consumption and on whether there is full or limited commitment between spouses.

3.1 The estimation problem

We introduce time indices to clarify the panel dimension of the estimation. Rearranging the first-order condition for labor supply, the labor-supply condition of agent ij in period t is given by

$$n_{ijt} = \left(\frac{w_{ijt} \cdot \lambda_{jt}}{\alpha_{ij} \cdot \mu_{ijt}} \right)^{\eta_i},$$

where we define $\mu_{ijt} = \mu'_{ij}$ for notational convenience. Using the first-order conditions for individual consumption $\lambda_{jt} = \mu_{ijt} \cdot c_{ijt}^{-\sigma}$ and taking logs, the labor-supply condition becomes

$$\log n_{ijt} = -\eta_i \cdot \log \alpha_{ij} + \eta_i \cdot \log w_{ijt} - \eta_i \cdot \sigma \cdot \log c_{ijt}.$$

Hence, if one had access to data on individual consumption, one could simply estimate

$$\log n_{ijt} = \delta_{0,ij} + \delta_{1,i} \cdot \log w_{ijt} + \delta_{2,i} \cdot \log c_{ijt} + \varepsilon_{ijt}, \tag{100}$$

where $\delta_{1,i}$ and $\delta_{2,i}$ are regression coefficients, $\delta_{0,ij}$ is an individual fixed effect, and ε_{ijt} is a residual stemming from measurement errors. Estimating equation (100) with OLS would yield an unbiased estimate, $\hat{\delta}_{1,i}$ for the Frisch labor-supply elasticity,

$$\mathbb{E}\ \hat{\delta}_{1,i} = \eta_i.$$

However, to estimate condition (100), one needs data on individual consumption. In a couple household, individual consumption can behave differently from household consumption and usually only the latter is observed in household panel data. Then, the estimation equation (100) could only be estimated using household consumption which would result in an estimation bias when household consumption and individual consumption are not perfectly correlated.

We now derive this bias analytically and discuss how it depends on the degree of commitment in the household. Household consumption is given by

$$c_{jt} = c_{mjt} + c_{fjt}.$$

Household consumption can be related to individual consumption of one of the spouses through the risk-sharing condition

$$\mu_{mjt} \cdot c_{mjt}^{-\sigma} = \mu_{fjt} \cdot c_{fjt}^{-\sigma}.$$

Then, household consumption is given by

$$c_{jt} = \left(1 + x_{ijt}^{-1/\sigma}\right) \cdot c_{ijt},$$

where $x_{ijt} = \mu_{ijt}/\mu_{-ijt}$ denotes the relative Pareto weight of individual ij, and $-ij$ denotes the spouse of agent ij. Using these conditions, the labor-supply condition can be written as

$$\log n_{ijt} = -\eta_i \cdot \log \alpha_{ij} + \eta_i \cdot \log w_{ijt} + \eta_i \cdot \sigma \cdot \log\left(1 + x_{ijt}^{-1/\sigma}\right) - \eta_i \cdot \sigma \cdot \log c_{jt}. \tag{101}$$

Hence, if the relative Pareto weight x_{ijt} were observable, one could estimate

$$\log n_{ijt} = \gamma_{0,ij} + \gamma_{1,i} \cdot \log w_{ijt} + \gamma_{2,i} \cdot \log c_{jt} + \gamma_{3,i} \cdot \log\left(1 + x_{ijt}^{-1/\sigma}\right) + \varepsilon_{ijt}, \tag{102}$$

where $\gamma_{0,ij}$ is an individual fixed effect, $\gamma_{1,i}$, $\gamma_{2,i}$, and $\gamma_{3,i}$ are regression coefficients, and ε_{ijt} is a residual, and would obtain an unbiased estimate of the Frisch elasticity,

$$\mathbb{E}\ \widehat{\gamma}_{1,i} = \eta_i.$$

However, while household consumption, c_{jt}, on the right-hand side of (102) is in principle observable, the relative Pareto weight, x_{ijt}, cannot be observed in empirical data. Thus, one would usually estimate the equation

$$\log n_{ijt} = \kappa_{0,ij} + \kappa_{1,i} \cdot \log w_{ijt} + \kappa_{2,i} \cdot \log c_{jt} + \tilde{\varepsilon}_{ijt}, \tag{103}$$

where the composed residual, $\tilde{\varepsilon}_{ijt}$, is given by

$$\tilde{\varepsilon}_{ijt} = \log\left(1 + x_{ijt}^{-1/\sigma}\right) + \varepsilon_{ijt}.$$

Due to the omission of $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ the estimated coefficient on the log wage rate is

$$\mathbb{E}\ \widehat{\kappa}_{i,1} = \eta_i + \eta_i \cdot \sigma \cdot \frac{cov\left(\log w_{ij}, \log\left(1 + x_{ij}^{-1/\sigma}\right)\right)}{var\left(\log w_{ij}\right)}.$$

Hence, such an estimation would still yield an unbiased estimate of the Frisch elasticity if the omitted relative Pareto weight were a constant, i.e., if there were full commitment. Then, the term $\eta_i \cdot \sigma \cdot \log\left(1 + x_{ij}^{-1/\sigma}\right)$ would simply become part of the individual fixed effect, $\kappa_{0,ij}$.

However, under limited commitment, the omitted term $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ is varying over time. The deviation from its individual-specific mean would then be part of the combined residual $\tilde{\varepsilon}_{ijt}$. This is problematic because the relative Pareto weight and hence the combined residual is correlated with the wage rate. Wage raises improve outside options and hence bargaining positions and can therefore lead to an increase in the individual's Pareto weight. Hence, $\log w_{ijt}$ and $\tilde{\varepsilon}_{ijt}$ (which includes a term decreasing in x_{ijt}) are negatively correlated and, as a consequence, an omitted variable bias occurs and the estimate of the Frisch labor-supply elasticity, $\hat{\kappa}_{1,i}$, is downward biased.

3.2 Addressing the estimation problem

The bias in the estimate of the Frisch elasticity originates from the fundamental unobservability of spouses' relative bargaining positions which are part of the true labor-supply condition (101) in form of the term $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$. In this section, we present two approaches that ameliorate this problem. Both approaches have in common that they use the observable behavior of an individual's spouse to gather information about the individual's relative bargaining position. Specifically, information about relative Pareto weights can be obtained from spouses' relative labor supply conditional on wage rates. In particular, dividing the two labor-supply conditions (95) and (96) yields

$$x_{ijt} = \frac{\alpha_{-ij}}{\alpha_{ij}} \cdot \frac{w_{ijt}}{w_{-ijt}} \cdot \frac{n_{-ijt}^{1/\eta_{-i}}}{n_{ijt}^{1/\eta_i}}. \tag{104}$$

Our first approach, outlined in Section 3.2, relies on a log-linear approximation of the term $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ in equation (101) and substituting in condition (104). The advantage of this approach is that it yields a linear condition that can easily be estimated. Its disadvantage stems from the necessary approximation which implies that the resulting coefficients are subject to approximation errors.

Our second approach, described in Section 3.2, is to make $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ in equation (101) effectively observable by using condition (104) and then to estimate (102). This approach delivers an unbiased estimate of the Frisch elasticity but requires iterative estimation with updated guesses for relative bargaining weights and has to be performed jointly for both men and women.

Linear estimation approach

To obtain a linear estimation condition, we have to replace $\log\left(1+x_{ijt}^{-1/\sigma}\right)$ in equation (102) by a linear term in logs that is log linear in observables. We apply a first-order Taylor approximation of $\log\left(1+x_{ijt}^{-1/\sigma}\right)$ around a given value for the relative Pareto weight, $x_{ijt}=\overline{x}$,

$$\begin{aligned}\log\left(1+x_{ijt}^{-1/\sigma}\right) &= \log\left(1+\overline{x}^{-1/\sigma}\right)+\frac{1}{1+\overline{x}^{-1/\sigma}}\cdot\frac{1}{\sigma}\cdot\left(x_{ijt}-\overline{x}\right)+\Phi_{ijt} \\ &= \log\kappa+\frac{1}{\kappa\sigma}\cdot\frac{1}{\overline{x}}\left(\log x_{ijt}+\log\overline{x}\right)+\Lambda_{ijt},\end{aligned} \tag{105}$$

where $\kappa = 1+\overline{x}^{-1/\sigma}$, Φ_{ijt} and Λ_{ijt} are higher-order terms, and the last step uses that $x_{ijt}-\overline{x} = 1/\overline{x}\cdot\left(x_{ijt}\overline{x}-1\right)$ is approximately $1/\overline{x}\cdot\log(x_{ijt}\overline{x})$ in the vicinity of the point of approximation.[1]

Using the condition for the relative Pareto weight (104) in logs, we obtain

$$\log\left(1+x_{ijt}^{-1/\sigma}\right) = \log\kappa+\frac{1}{\kappa\sigma}\cdot\begin{pmatrix}\log w_{-ijt}-\log w_{ijt}+\log\alpha_{ij}-\log\alpha_{-ij}\\ +\frac{1}{\eta_i}\log n_{ijt}-\frac{1}{\eta_{-i}}\log n_{-ijt}+\log\overline{x}\end{pmatrix}+\Lambda_2.$$

Inserting this equation in (101) and solving for n_{ijt} yields the approximated labor-supply condition

$$\begin{aligned}\log n_{ijt} = {} & \Theta_{ij}+\eta_i\cdot\log w_{ijt}+\eta_i\cdot\frac{\vartheta}{1-\vartheta}\cdot\log w_{-ijt}+\frac{\eta_i}{\eta_{-i}}\cdot\frac{\vartheta}{1-\vartheta}\cdot\log n_{-ijt} \\ & -\frac{1}{1-\vartheta}\cdot\eta_i\cdot\sigma\cdot\log c_{jt}-\eta_i\cdot\frac{\vartheta}{1-\vartheta}\cdot\log\alpha_{ij}-\eta_i\cdot\frac{\vartheta}{1-\vartheta}\cdot\log\alpha_{-ij} \\ & +\eta_i\cdot\sigma\cdot\left(\log\kappa+\vartheta\log\overline{x}\right)+\eta_i\cdot\sigma\cdot\frac{\vartheta}{1-\vartheta}\cdot\Lambda_{ijt},\end{aligned} \tag{106}$$

with $\vartheta = 1/\left(\kappa\sigma\right)$. Hence, estimating

$$\log n_{ijt} = \zeta_{0,ij}+\zeta_{1,i}\cdot\log w_{ijt}+\zeta_{2,i}\cdot\log w_{-ijt}+\zeta_{3,i}\cdot\log n_{-ijt}+\zeta_{4,i}\cdot\log c_{jt}+\varepsilon_{ijt}, \tag{107}$$

where $\zeta_{0,ij}$ is an individual fixed effect, $\zeta_{1,i}$, $\zeta_{2,i}$, $\zeta_{3,i}$, and $\zeta_{4,i}$ are regression coefficients, and ε_{ijt} is a residual delivers an estimate of the Frisch elasticity of labor supply,

$$\mathbb{E}\,\widehat{\zeta}_{i,1}\approx\eta_i,$$

which is close to the true Frisch elasticity if the first-order approximation in (105) is sufficiently accurate. More accuracy can be achieved by adding the non-linear terms in Λ_2, i.e., higher orders and interactions of spouses' wage rates and hours. We quantify the accuracy of this

[1] The particular point of approximation chosen, $\overline{x}$, is irrelevant for the main results as it does not affect the relation between wage rates and labor supply, see below.

approach in our Monte-Carlo study in Section 4.2.

Iterative estimation

In this approach, we want to deduce the Pareto weights and ultimately the term $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ in equation (102) from observed information on relative labor supply and wage rates. From equation (104), one can deduce the relative Pareto weight x_{ijt} from observations for wage rates w_{ijt} and w_{-ijt} and hours n_{ijt} and n_{-ijt} when one has a guess for the parameters α_{ij}, α_{-ij}, η_i, and η_{-i}. From this deduced value for x_{ijt}, one can infer $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ with a guess for σ. The key to an unbiased estimate for the Frisch elasticity is that estimated coefficients from a regression of condition (102) also contain information about all these parameters which can be used to update guesses. Therefore, we propose an iterative approach where we guess parameter values, infer relative bargaining weights, run a regression of condition (102), update the guesses and then iterate until convergence.

Specifically, the estimation approach proceeds as follows. First, one makes a guess for η_m, η_f, and σ as well as for the household-specific preference weights, α_{ij}. Then, one repeats the following steps until convergence.

1. Use the observable information on w_{ijt}, w_{-ijt}, n_{ijt}, and n_{-ijt} as well as the guesses for η_m, η_f, σ, and the α_{ij}'s to calculate $x_{ijt} = \alpha_{-ij}/\alpha_{ij} \cdot w_{ijt}/w_{-ijt} \cdot n_{-ijt}^{1/\eta_{-i}}/n_{ijt}^{1/\eta_i}$ and $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ for every individual ij and every period t.
2. Using the values for $\log\left(1 + x_{ijt}^{-1/\sigma}\right)$ run gender-specific regressions of condition (102). Use the estimated coefficient to update the guesses as follows: $\eta_i = \widehat{\gamma}_{1,i}$, $\sigma = \widehat{\gamma}_{2,i}/\eta_i$, and $\alpha_{ij} = \exp\left(-\widehat{\gamma}_{0,ij}/\eta_i\right)$.

After convergence, the estimate on the log wage rate in the regression for gender i is an unbiased estimate of the Frisch elasticity,

$$\mathbb{E}\, \widehat{\gamma}_{1,i} = \eta_i.$$

3.3 Graphical illustration

We now illustrate graphically estimation biases resulting from limited commitment and how exploiting observable behavior of spouses can help correcting these biases.For this, we use our calibrated economy which is solved numerically (see Section 4 below for a description of our calibration strategy). Here, we provide a graphical illustration of the different effects changes

Figure 10: Policy functions.

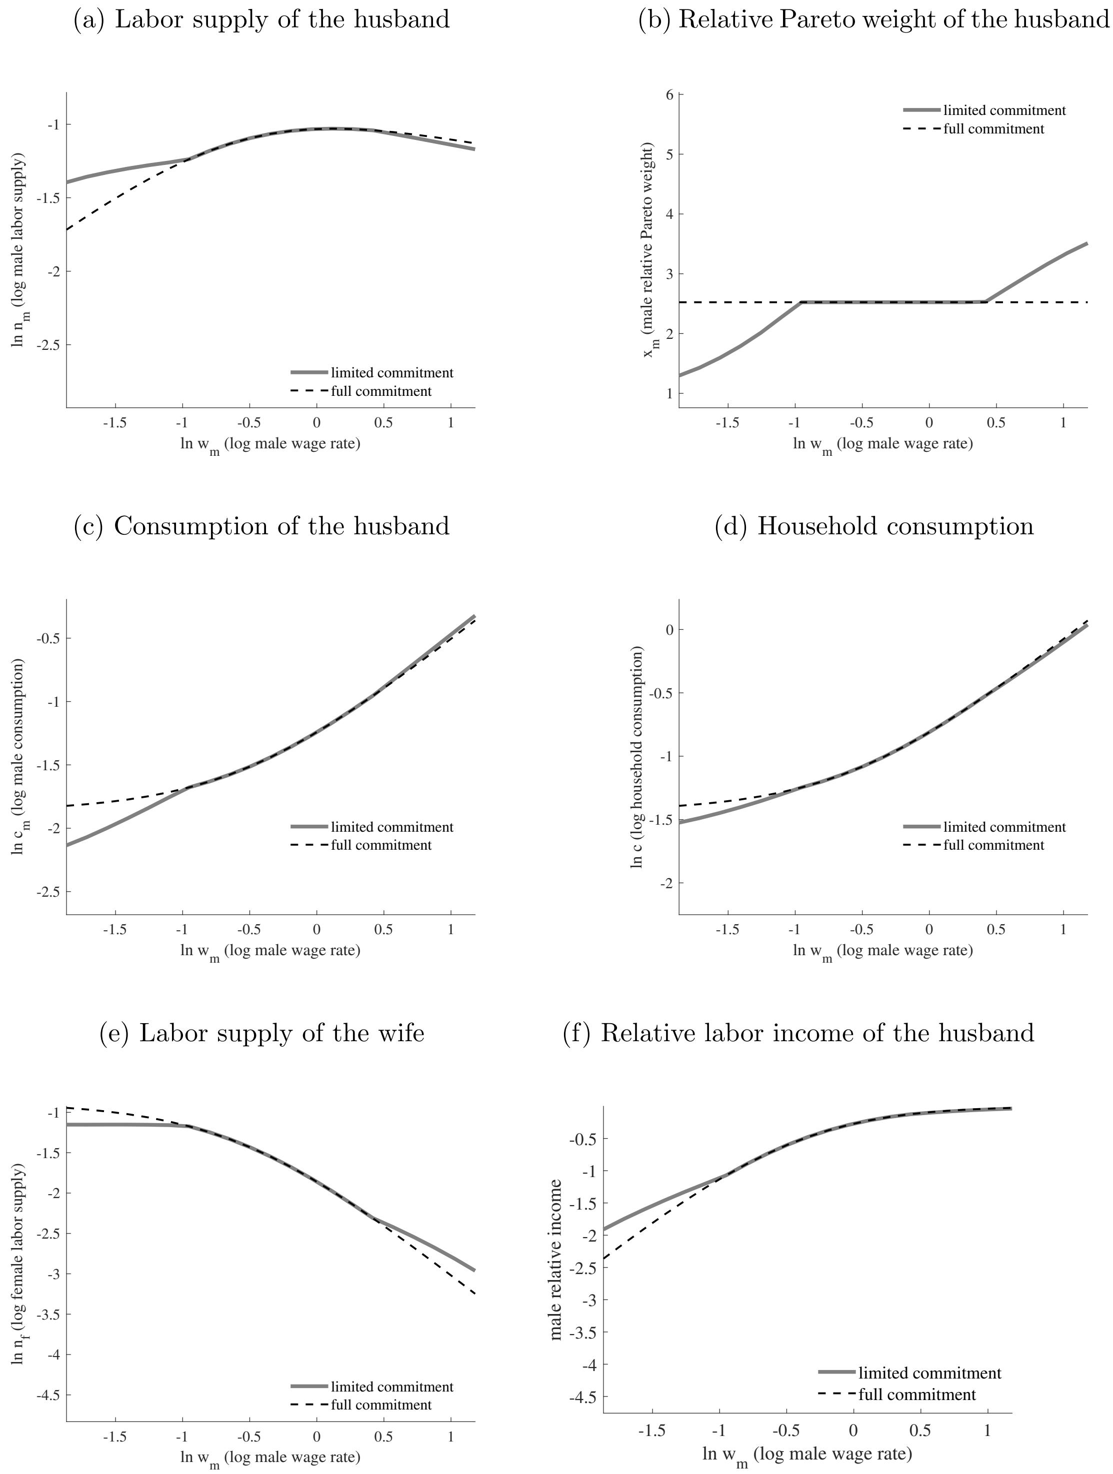

NOTE.–See Section 4.1 for a description of the calibration of the quantitative model. The wage rate of spouse ij is given by $\log w_{ij} = \log \psi_{ij} + \log \nu_i + z_{ij}$, where z_{ij} follows a first-order autoregressive process with autocorrelation 0.9136 and innovation variance 0.0426. The gender wage gap is $\frac{\nu_f}{\nu_m} = 0.8$. In the figure, we consider a household where both spouses have the same fixed effect, $\psi_{ij} = -0.3428$. Thus, (logged) mean wage rates are -0.3428 (male) and -0.5659 (female). Wife's wage rate w_{fj} and relative Pareto weight from the previous period, x_{mj}, fixed at their respective mean levels.

Figure 11: Illustration of estimation approaches.

(a) Labor supply, controlled for individual consumption

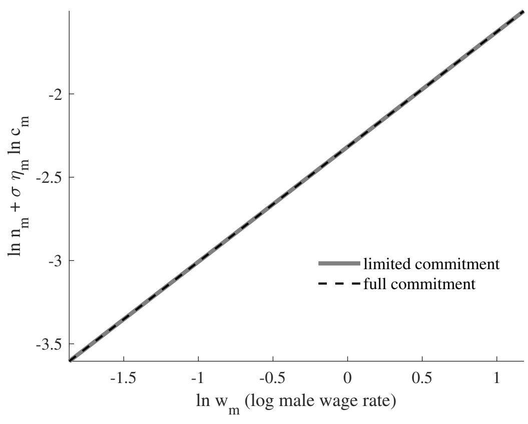

(b) Labor supply, controlled for household consumption

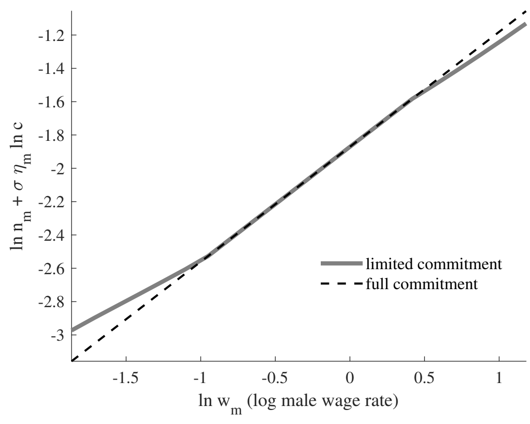

(c) Labor supply, controlled for household consumption and spouse's hours and wage rate

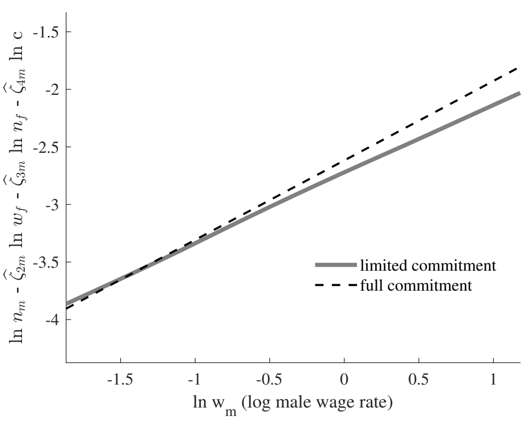

(d) Labor supply, controlled for household consumption and deduced bargaining weights

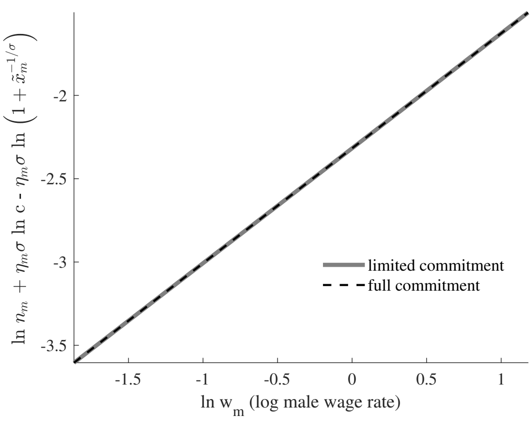

NOTE.–See Section 4.1 for a description of the calibration of the quantitative model. The wage rate of spouse ij is given by $\log w_{ij} = \log \psi_{ij} + \log \nu_i + z_{ij}$, where z_{ij} follows a first-order autoregressive process with autocorrelation 0.9136 and innovation variance 0.0426. The gender wage gap is $\frac{\nu_f}{\nu_m} = 0.8$. In the figure, we consider a household where both spouses have the same fixed effect, $\psi_{ij} = -0.3428$. Thus, (logged) mean wage rates are -0.3428 (male) and -0.5659 (female). Wife's wage rate w_{fj} and relative Pareto weight from the previous period, x_{mj}, fixed at their respective mean levels.

in wage rates exert on labor supply in our model, how they bias labor-supply estimations, and how the observable behavior of the spouse can help correcting this bias.

Panel (a) of Figure 1 shows the labor supply of the husband as a function of his wage rate, hence, his labor-supply function. For the figures, we have fixed the wife's wage rate w_{fj} and the relative Pareto weight from the previous period, x_{mj}, at their respective mean levels. Under full commitment (see dashed line), the labor-supply function of the husband is smoothly upward sloping with a flattening out at high wage rates where income effects are strong. Under limited commitment (see solid line), the labor-supply function has two additional kinks. When the wage rate is high, the individual works less than under full commitment and, reversely, the individual works more than under full commitment when the wage rate is low. These kinks visualize the limited-commitment effect of changes in wage rates on labor supply. They are the result of adjustments in the husband's relative Pareto weight x_{mj} which is shown in Panel (b) of Figure 1. When the wage rate increases, the outside option of the husband improves. Eventually, he is tempted to leave the household and can only be prevented from doing so when the household increases the weight of the husband in household preferences. This induces the household to value his leisure time more strongly and to reduce his labor supply accordingly. When the husband's wage rate is low, it is the wife who may be tempted to leave the household because she now has to give much of her earnings to the husband which makes her outside option relatively more attractive. Then, the Pareto weight of the husband is reduced and, accordingly, his labor supply is larger compared to the case of full commitment.

To correct for the limited-commitment effect in a labor-supply regression, one needs to include variables that carry sufficient information about spouses' relative bargaining positions or, put differently, also show the respective kinks which reflect adjustment of relative bargaining weights. Panel (c) shows that individual consumption carries this information. When commitment between spouses is limited, consumption of the husband kinks upward when his relative Pareto weight increases and, reversely, his consumption kinks downward when his relative weight is reduced. As discussed before, information on individual consumption is mostly not available in household panel data and in empirical estimations one therefore needs to rely on total household consumption. In contrast to individual consumption, household consumption is substantially less informative about changes in bargaining positions, see panel (d). Household consumption shows no kinks but only a slight and smooth reduction when Pareto weights need to be adjusted.[2]

[2]In these situations, the spouse with the relatively low wage works relatively much and the other spouse enjoys relatively much consumption in order to be kept in the household. Hence, labor earnings and thus

Panel (e) shows the wife's labor supply as a function of the husband's wage rate. While the cross-wage effect is generally negative reflecting income effects, we also observe kinks that reflect limited-commitment effects. Thus, the cross-wage effect does contain some information about changes in relative bargaining weights. When the husband's wage rate is high and his Pareto weight is increased, labor supply of the wife rises because the household values her leisure time relatively less.[3]. This shows that controlling for the spouse's hours worked in a labor-supply regression picks up some of the limited-commitment effect which one wants to correct for in order to isolate the substitution effect and thus helps to identify the Frisch elasticity. Note that this approach is feasible in empirical applications as household panel surveys usually do report variables for the partner. Finally, panel (f) of Figure 10 shows the husband's relative labor income which is connected to the term in condition (104) which we use to deduce relative bargaining weights in our iterative procedure. Also this variable shows kinks where Pareto weights are adjusted under limited commitment. Note that the latter two variables, the cross-wage labor supply and relative labor earnings, are informative about changes in relative bargaining weights because the kinks in those variables go into opposite directions depending on whether the husband's relative Pareto weight is adjusted upward or downward. By contrast, household consumption is fundamentally inept to account for the limited-commitment effect because household consumption is reduced whenever weights are adjusted, independent of the direction of the adjustment.

In Figure 11, we illustrate our estimation approaches graphically. Specifically, we plot rearranged versions of the estimation conditions (100), (101), (103), and (107) where we subtract terms from the right-hand side which are in principle observable or can be made observable, respectively.[4] In panel (a), we plot the term $\log n_{mj} + \sigma\eta_m \log c_{mj}$, i.e., the husband's labor supply controlled for the husbands' consumption, against the log wage rate of the husband. Under both full and limited commitment, the result is a straight line with slope 0.69 which is the Frisch elasticity of husband's labor supply in our a calibration. This illustrates that, if one had access to data on individual consumption, the Frisch elasticity could be inferred from a simple labor-supply regression with individual consumption as a control. Under limited commitment, the policy function for labor supply has kinks but these kinks are straightened out by controlling for individual consumption which shows the opposite kinks, see panels (a) and (c) of Figure 10.

consumption decrease relative to the full-commitment case.

[3]The opposite happens for low wage rates of the husband, where the wife's Pareto weight is increased and her labor supply is thus reduced

[4]The resulting lines can be understood as the analogs to scatterplots where one shows residuals from a regression of the dependent variable on control variables against the regressor of interest.

In panel (b) of Figure 11, we take into account that individual consumption data is usually not available and control for household consumption instead, i.e., we plot the term $\log n_{mj} + \sigma\eta_m \log c_j$ against the husband's log wage rate. If there were full commitment between spouses, the result would still be a straight line and the slope would equal the Frisch elasticity η_m. However, under limited commitment, the kinks from the policy function for labor supply are not straightened out by controlling for household consumption since the latter has no kinks.

In Panel (c) of Figure 11, we consider our linear estimation approach from Section 3.2. We plot the husband's labor supply controlling for household consumption as well as hours worked and the wage rate of the wife, i.e., the term $\log n_{ijt} - \widehat{\zeta}_{2,i} \cdot \log w_{-ijt} - \widehat{\zeta}_{3,i} \cdot \log n_{-ijt} - \widehat{\zeta}_{4,i} \cdot \log c_{jt}$ from a rearranged version of the estimation condition (107), using the estimated coefficients from the Monte-Carlo regressions.[5] Under full commitment, the result is a straight line with a slope close to the true Frisch elasticity which is 0.69. This is not surprising as already controlling for household consumption alone delivers this result, see panel (b). Under limited commitment, where controlling for household consumption alone is insufficient to straighten out the kinks in the labor-supply function, additionally controlling for the partner's hours and wage rate with their informative kinks does achieve a substantial straightening. Furthermore, the resulting line is only slightly flatter than its full-commitment counterpart, illustrating that an estimate close to the true Frisch elasticity can be achieved in a labor-supply regression that controls for household consumption as well as hours worked by the partner and the partner's wage rate.

Panel (d) of Figure 11 illustrates the final step of our iterative estimation procedure. When parameter guesses have converged sufficiently, the deduced value $\widetilde{x}$ for the relative Pareto weight is close to its true value x. Then, plotting the term $\log n_{mj} + \sigma\eta_m \log c_j - \sigma\eta_m \log(1 + \widetilde{x_{mj}}^{-1/\sigma})$, which is the rearranged left-hand side of condition (102) when the latter two terms are added, against the log wage rate of the wife results in a straight line for both full and limited commitment. In both cases, the slope of the line is equal to the true Frisch elasticity for women.

[5] Here, we use regression coefficients instead of model parameters because the parameter κ in the linearized condition (106) depends on the point of approximation, which we do not need to specify for the regression, and is infinite under full commitment.

4 Monte-Carlo study

In this section, we use our model to perform Monte Carlo studies to evaluate the performance of the estimation approaches discussed in Section 3.

4.1 Calibration and simulation set-up

Table 22 summarizes the parameter values. The parametrization is a combination of setting some parameters to values taken from the literature and calibrating others. Our calibration strategy focuses on targets in the labor market and the extent of commitment issues since these aspects are the most relevant for our analysis.

The stochastic wage rate component is assumed to follow a first-order autoregressive process

$$z_{ijt} = \rho \cdot z_{ijt-1} + \epsilon_{ijt},$$

with autocorrelation ρ and innovations $\epsilon \sim N\left(0, \sigma_\epsilon^2\right)$. Floden and Lindé (2001) estimate such a process for a joint sample of men and women from PSID data. We use their estimates and set $\rho = 0.9136$ and $\sigma_\epsilon^2 = 0.0426$. Wage-rate shocks are assumed to be uncorrelated among spouses which is in line with Ortigueira and Siassi (2013) who document a very low correlation of employment shocks for spouses who report different occupations. The parametrization of the fixed-effect component also follows Floden and Lindé (2001) who estimate $\sigma_\psi^2 = 0.1175$. As Floden and Lindé (2001), we assume that the fixed effect can take two values, $\psi_{ij} = \exp(-\sigma_\psi)$ and $\psi_{ij} = \exp(\sigma_\psi)$, with equal probability. We account for heterogeneity in mean wages to generate heterogeneity in mean Pareto weights across households which allows us to calibrate the degree of commitment problems, see below. Hyslop (2001) documents a correlation of wage fixed effects in households of 0.5 which we achieve by setting the shares of households where both spouses have the same fixed effect to 37.5% each. The shares of households where spouses have different fixed effects are set to 12.5% each. We calibrate the gender-specific components of the wage rates to match the observed gender wage gap of about 20% (see, e.g., Blau and Kahn 2017). To match this target, we set $\nu_m = 1$ and $\nu_f = 0.8$.

The true Frisch labor-supply elasticities in our model are given by the parameters η_m and η_f. In line with the empirical results of Bredemeier et al. (2019), we set $\eta_m = 0.69$ and $\eta_f = 1.05$.

In line with the parametrization of the wage process, we consider two different average

Table 22: Parameter values, baseline model

Description	Parameter	Value	Source
Autocorrelation wage shock	ρ	0.9136	Floden and Lindé (2001)
Variance wage shock	σ_{ε}^2	0.0426	Floden and Lindé (2001)
Variance fixed effect	σ_{ψ}^2	0.1175	Floden and Lindé (2001)
Gender-specific wage component ...			
... men	ν_m	1	calibrated
... women	ν_f	0.8	calibrated
Frisch labor-supply elasticity ...			
... men	η_m	0.69	Bredemeier et al. (2019)
... women	η_f	1.05	Bredemeier et al. (2019)
Disutility weight labor ...			
... highly educated men	$\alpha_{m,\text{high}}$	13.69	calibrated
... less educated men	$\alpha_{m,\text{low}}$	28.75	calibrated
... highly educated women	$\alpha_{f,\text{high}}$	19.79	calibrated
... less educated women	$\alpha_{f,\text{low}}$	54.27	calibrated
Marriage utility gain parameter ...			
... households with commitment issues	Ψ_{low}	0.0286	calibrated
... households without commitment issues	Ψ_{high}	5	set
Discount factor	β	0.96	set
Risk aversion	σ	1.5	set

labor-supply targets by gender. Aguiar and Hurst (2007) document average hours worked of different education groups, both for men and women. We distinguish between individuals with a high-school degree or less and individuals with some college education or more. These groups' average weekly hours are 37.04 and 43.05, respectively, for men and 21.84 and 29.86, respectively, for women. Dividing by a weekly endowment of freely disposable time of 112 hours gives the calibration targets $n_{low,m} = 0.33$, $n_{high,m} = 0.38$, $n_{low,f} = 0.20$, and $n_{high,f} = 0.27$. We set the group-specific disutility weights to match these targets.[1]

We distinguish between two different values for the direct utility gains from marriage, Ψ_j.[2] Del Boca and Flinn (2012) find that 93% of all households behave only *ex-post* efficient while 7% of households appear to behave in a way that is also ex-ante efficient, i.e., spouses seem to be able to fully commit to promises made to each other. Accordingly, for 7% of households, we set the direct utility gain from marriage to a value that implies that Pareto weights never need to be adjusted.[3] For the remaining 93% of households, we calibrate direct utility gains from marriage in order to match the ratio of the within-household variance of relative Pareto weights to the between-household variance of one to four as documented by Lise and Yamada (2019). This ratio is informative for the strength of commitment problems within households as the within-household variance of relative Pareto weights reflect revisions or rebargaining of these weights which only happen under limited commitment. By contrast, the between-household variance reflects heterogeneity across households at time of marriage.

The remaining model parameters are set to standard values from the literature. The discount factor is $\beta = 0.96$ and we assume $\sigma = 1.5$.

Having solved the model[4], we simulate synthetic data for 5,000 households of each household type, i.e., 40,000 households in total. For each household, we simulate 1,500 periods and discard the first 500 periods to avoid dependence on initial conditions which gives a total of 40,000,000 household-year observations. In our Monte-Carlo experiments, we draw 10,000 samples of 15,000 year-household combinations taking into account the the population shares of the household types. We report average point estimates and average standard errors in our estimations.

[1] In Aguiar and Hurst (2007)'s sample from the American Time Use Survey, nearly 50% of all men and women have a high school degree or less which is consistent with our calibration of the wage fixed effects.

[2] Together with the four household types in the wage fixed-effect dimension, there are in total eight household types.

[3] We set $\Psi_{high} = 5$.

[4] See Appendix A.2 for details on the numerical solution of the model.

4.2 Simulation results

In this section, we present our Monte-Carlo results. We use simulated data to estimate the four estimation approaches discussed above, i.e., conditions (100), (102), (103), and (107). Table 23 summarizes the results for men and Table 25 in the Appendix shows the results for women. When describing the results, we focus on those for men. Qualitatively, results are the same for women. In both tables, the first two columns refer to the reference economy with full commitment while columns 3 through 6 refer to the model with limited commitment.

We start by estimating condition (100), i.e., we run a regression of log hours worked on an individual fixed effect, the log wage rate, and log individual consumption. This is possible in the Monte-Carlo set-up where we can observe individual consumption while this is usually not possible with real-world data as household panel data usually lack information on individual consumption. Columns (1) and (3) of Table 23 show that, under both full and limited commitment, such an estimation delivers an unbiased estimate of the Frisch elasticity (recall that the true Frisch elasticity is $\eta_m = 0.69$ for men).

We now estimate condition (102) where individual consumption is replaced by household consumption. If spouses in a couple were able to fully commit to promises made to each other, controlling for household consumption would be sufficient to obtain an unbiased estimate of the Frisch elasticity, see column (2) of Table 23. The reason is that, with constant Pareto weights under full commitment, individual consumption is proportional to household consumption and the former can hence be perfectly proxied by the latter. However, taking into account that commitment between spouses is limited in most households shows that controlling for household consumption instead of individual consumption leads to a substantial bias in the estimated Frisch elasticity of about 30%, see column (4) of Table 23. Put differently, the true Frisch elasticity is about 50% larger than the results of such a regression suggest.

We now turn to our improved estimation approaches. We apply them only for the model with limited commitment since, in the reference model with full commitment, already the simple regression (102) yields unbiased estimates. We start with our linear approach (107), where the labor-supply regression is augmented by information on the wage rate and labor supply of the individual's spouse. Since the estimation approach is based on an approximation, its accuracy has to be assessed quantitatively. We consider two variants of our linear estimation approach. First, we estimate condition (107) without any higher-order terms. Second, we add the second-order terms from a Taylor expansion of the term $\log(1 + x^{-1/\sigma})$

Table 23: Monte-Carlo results for men, from synthetic household panel data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	full commitment		limited commitment				
log wage rate $\log w_{mjt}$	0.69	0.69	0.69	0.48	0.60	0.74	0.69
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)
log individual	-1.04	–	-1.04	–	–	–	–
consumption $\log c_{mjt}$	(0.00)	–	(0.00)	–	–	–	–
log household	–	-1.04	–	-0.80	-1.46	-1.62	-1.04
consumption $\log c_{jt}$	–	(0.00)	–	(0.00)	(0.00)	(0.01)	(0.00)
log wage rate	–	–	–	–	0.34	0.34	–
of the spouse, $\log w_{fjt}$	–	–	–	–	(0.00)	(0.00)	–
log hours worked	–	–	–	–	-0.38	-0.34	–
by the spouse, $\log n_{fjt}$	–	–	–	–	(0.00)	(0.01)	–
$\log w_{mjt} \times \log w_{fjt}$	–	–	–	–	–	-0.07	–
	–	–	–	–	–	(0.00)	–
$\log w_{mjt} \times \log n_{fjt}$	–	–	–	–	–	0.05	–
	–	–	–	–	–	(0.00)	–
$\log w_{fjt} \times \log n_{fjt}$	–	–	–	–	–	-0.05	–
	–	–	–	–	–	(0.00)	–
$(\log w_{mjt})^2$	–	–	–	–	–	0.03	–
	–	–	–	–	–	(0.00)	–
$(\log w_{fjt})^2$	–	–	–	–	–	0.04	–
	–	–	–	–	–	(0.00)	–
$(\log n_{fjt})^2$	–	–	–	–	–	0.02	–
	–	–	–	–	–	(0.00)	–
deduced relative Pareto	–	–	–	–	–	–	1.04
weight term $\log(1 + x_{mjt}^{-1/\sigma})$	–	–	–	–	–	–	(0.00)
bias in %	–	–	–	-30	-13	7	–
fixed effects	yes	yes	yes	yes	yes	yes	yes
empirically feasible	no	yes	no	yes	yes	yes	yes

NOTE.–Results are for men. The dependent variable is (logged) hours of the male in household j in period t. A constant and household-type fixed effects are included but not reported. Table shows average estimates from 10,000 Monte-Carlo draws, with average standard errors in parentheses.

which reflects the role of bargaining positions in the labor-supply equation. [5]

The results reported in column (5) of Table 23 show that extending the labor-supply regression by the partner's wage rate and hours reduces the estimation bias by more than half. However, there still remains a bias of about 13%. When we add the second-order terms in column (6), we obtain another reduction of the bias by about half. The estimated Frisch elasticity is now 0.74 and hence differs from the true value of 0.69 by only 7%.[6]

This remaining bias can be eliminated by our iterative procedure. In this approach, we alternate between deducing relative Pareto weights based on parameter guesses and using results from estimating condition (102) to update the guesses. This procedure converges to a regression where relative bargaining weights are adequately controlled for and which hence yields an unbiased estimate of the Frisch elasticity. Our quantitative results reported in column (7) of Table 23 show that our iterative approach delivers an unbiased estimate when applied to the calibrated economy.

In order to analyze robustness to initial parameter guesses, we report the convergence properties of our iterative approach for four different sets of initial parameter guesses in Figure 12.[7] Recall that each iteration includes two regressions (one for each gender) with individual fixed effects. Our iterative approach is found to be robust to initial parameter guesses. Even for the worst set of initial guesses, #4, the estimator converges after about 200 iterations.[8] The starting values for the disutility weights α_m and α_f are relatively unimportant for the convergence of the estimates for the Frisch elasticities η_m and η_f.[9]

In an application to empirical data, our iterative procedure is complicated by the fact that one has to estimate a labor-supply regression also for women even if one is only interested in estimating the Frisch elasticity for men. In the real world, female labor supply is subject to a

[5] Since $\log x$ can be expressed as a linear function of $\log w_m,\ \log w_f,\ \log n_m$, and $\log n_f$, see (104), the second-order term of the Taylor expansion includes the quadratic terms $(\log w_m)^2,\ (\log w_f)^2,\ (\log n_m)^2$, and $(\log n_f)^2$, the multiplicative terms $\log w_m \times \log w_f,\ \log w_m \times \log n_m\ \log w_m \times \log n_f,\ \log w_f \times \log n_m,\ \log w_f \times \log n_f$, and $\log n_m \times \log n_f$, as well as constants. Of these terms, we include in the regression for men all but the ones that contain the left-hand side variable $\log w_m$. We have checked that alternatively including also these terms confounds the estimation substantially.

[6] Including also third-order terms (not shown) yields an estimated Frisch elasticity of 0.71. Hence, the bias is reduced by another four percentage points which comes, however, at the costs of including ten additional terms (terms not including $\log n_m$). We view the regression including second-order terms shown in column (6) of Table 23 as a reasonable compromise between consistency and applicability.

[7] We use identical initial guesses for the labor-disutility weights for each individual. We then update them individually as described in Section 3.2. In the figure, we focus on the estimates for η_m and η_f. The convergence behavior of the other parameters is similar.

[8] For the set of initial guesses #4, the estimates after 100 iterations are 0.7038 for η_m and 1.0777 for η_f. After 200 iterations, the estimates are almost unbiased, 0.6913 for η_m and 1.0526 for η_f (true values: $\eta_m = 0.69,\ \eta_f = 1.05$).

[9] In initial-guess set #3, we consider the case where initial guesses for α_m and α_f are far away from their true values. Nevertheless, the lines in both panels are virtually indistinguishable to the lines for the more elaborated initial guesses in set #1 .

Figure 12: Convergence of the iterative procedure.

(a) Estimate of the labor-supply elasticity for men, η_m

(b) Estimate of the labor-supply elasticity for women, η_f

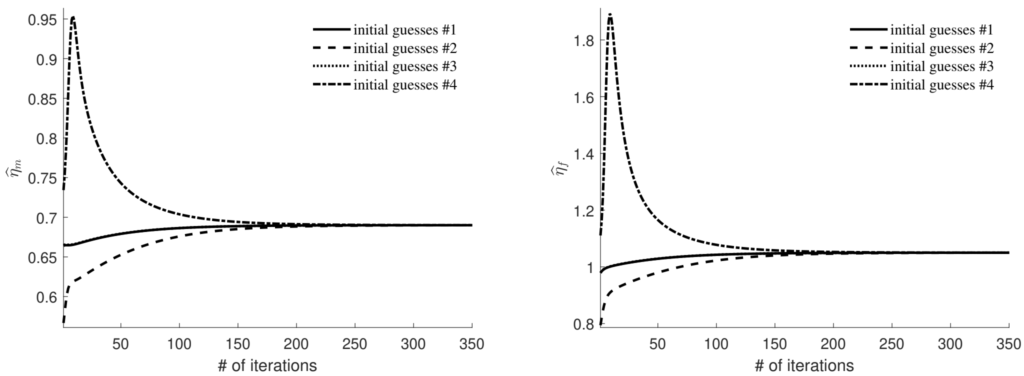

NOTE.–Convergence properties of our iterative approach for four different sets of initial guesses.
Initial guesses #1: $\eta_m = 0.8,\ \eta_f = 0.8,\ \alpha_{ij} = 20,\ \sigma = 2.$
Initial guesses #2: $\eta_m = 0.3,\ \eta_f = 2,\ \alpha_{ij} = 20,,\ \sigma = 2.$
Initial guesses #3: $\eta_m = 0.8,\ \eta_f = 0.8,\ \alpha_{ij} = 50,\ \sigma = 2.$
Initial guesses #4: $\eta_m = 30,\ \eta_f = 30,\ \alpha_{ij} = 1200,\ \sigma = 5.$

number of issues which are substantially less relevant for men and which we did not account for in our simulation study. Examples are selection into the labor force (see e.g., Kimmel and Kniesner 1998) and costs of child care (see e.g., Domeij and Klein 2013). These issues would have to be taken into account in order to obtain an unbiased estimate of women's Frisch elasticity which is needed to identify men's Frisch elasticity in the iterative approach. For this reason, and given the good performance of the linear approach when second-order terms are included, we will focus on the linear approach in the empirical analysis in the next section.

5 Empirical analysis

5.1 Sample selection

We use PSID data to estimate our approximated linear-regression approach. Our sample selection closely follows Blundell et al. (2016) who use PSID waves 1999–2009 to investigate family labor supply as an insurance mechanism. As discussed by Blundell et al. (2016), a key advantage of recent waves of the PSID is the newly included detailed information on household consumption. The new design of the PSID covers over 70% of all consumption items covered in the Consumer Expenditure Survey (CEX).[1] We update the Blundell et al. (2016) sample and consider the extended sample period 1999-2017. During this period, the PSID data has been collected biennially.

We drop the Survey of Economic Opportunity sample and the immigrant sample. We consider married households where both spouses are 30-57 years old as the intensive margin of labor supply is the dominating one in this age group, see Blundell et al. (2016). To apply our approach, we have to select dual-earner households where both spouses are working. As in Blundell et al. (2016), we drop couples in the period where they dissolve but include these household heads when they marry again.

We use the same definition of non-durable consumption and services as Blundell et al. (2016).[2] Missing values in the subcategories of consumption are treated as zeros. Also following Blundell et al. (2016), we construct an education variable with three categories that is consistent over time.[3]. As is standard in the labor-supply literature, the hourly wage is constructed by dividing annual earnings of the spouses by their annual hours worked.[4]

We apply similar sample selection criteria to reduce the impact of measurement error as Blundell et al. (2016). Specifically, we drop observations with wages below half the hourly minimum wage, observations where couples have very high asset values ($20 million and

[1]See Blundell et al. (2016) for a comparison of consumption data from the PSID with the CEX and NIPA.

[2]Nondurable consumption consists of food at home and gasoline consumption and services consists of food out of home, health insurance, health services, utilities, transportation, education, child care, home insurance, and rent (or rent equivalent). Some sub-items of consumption (e.g., within the health services item) are not reported after 2011.

[3]Category 1 is 0-11 grades (includes those with no schooling); category 2 is 12 grades, i.e. high school with and without nonacademic training; category 3 is college dropout, college degree, or college and advanced/professional degree.

[4]Our regressions will take into account the so-called division bias resulting from the constructed wage rate variable. Earnings consist of labor earnings, the labor part of business income, and the labor part of farm income of household head.

more), couples who receive transfers higher than twice the total household earnings and we do not use data displaying extreme jumps from one PSID wave to the next.[5]

5.2 Empirical results

Table 24 summarizes our baseline estimation results from the PSID. In all specifications, we use the same sample of household heads for which we have non-missing data on education, consumption, wage rates, and hours worked for husband and wife (see bottom line in the table). As is well known in the labor-supply literature (see, e.g., Altonji 1986; Borjas 1980; Pencavel 1986 and Keane 2011), OLS results are subject to a division bias as the wage rate is computed as earnings divided by hours worked which generates a spurious negative correlation of the calculated hourly wage with hours worked because measurement error in hours worked occurs on both sides of the regression equation. We therefore run an initial wage regression and use it to calculate predicted wage rates which are uncorrelated with the measurement error in hours worked. We then use predicted wage rates $\widetilde{w}_{ijt}$ in the labor-supply regression. In the wage regression, we include age, age squared and interactions of these two variables with education, household consumption, two binary variables indicating whether the individual is African American or white, a full set of state dummies, as well as time effects. By taking into account the interaction terms in age and education, predicted wages are identified through the age profiles of the returns to education, taking into account a set of control variables as well as time effects that exploit trend and cyclicality of wages for prediction.[6] We also considered extended specifications for the wage regression, where we for instance include age and education of the individual's spouse or where we include additional interactions between year and education. We found that this has no strong effect on our baseline results. After estimating the wage regression, we eliminate the observations that fall into the 0.25 percentile of the distribution of estimated wage residuals by year (corresponding to 33 observations in total) and re-estimate the wage regression for the cleaned sample. This way, we drop individuals who tend to over-report hours most strongly.

Table 24 shows the results of our labor-supply regressions. In all specifications, we include the (log) wage rate, individual and time fixed effects and a quadratic in age. By including time effects in the labor-supply regression we filter out effects related to the economy-wide business cycle, i.e., demand-determined fluctuations in hours worked. By including age, which is also included in the wage equation, we control for the most important labor-supply taste

[5] A jump is defined as an extremely positive (negative) change from $t-2$ to t, followed by an extreme negative (positive) change from t to $t+2$, see Blundell et al. (2016), footnote 16, for details.

[6] Such a specification for the wage regression is widely used in the literature. In our case, the wage regression has an R^2 of about 33%.

Table 24: Estimation results for men, from PSID data.

	(1)	(2)	(3)	(4)	(5)	(6)
log wage	0.08	0.26	0.71	0.70	0.71	0.65
rate $\log \tilde{w}_{mjt}$	(0.02)	(0.13)	(0.38)	(0.38)	(0.38)	(0.37)
	p=0.00	*p=0.05*	*p=0.06*	*p=0.07*	*p=0.06*	*p=0.08*
log household	–	-0.11	-0.13	-0.13	-0.13	-0.04
consumption $\log c_{jt}$	–	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
		p=0.17	*p=0.09*	*p=0.10*	*p=0.12*	*p=0.66*
log wage rate	–	–	0.38	0.38	0.40	0.60
of the spouse $\log w_{fjt}$	–	–	(0.32)	(0.32)	(0.32)	(0.33)
			p=0.24	*p=0.24*	*p=0.22*	*p=0.07*
log hours worked by	–	–	0.09	0.08	0.09	0.20
the spouse $\log n_{fjt}$	–	–	(0.14)	(0.14)	(0.14)	(0.14)
			p=0.53	*p=0.54*	*p=0.50*	*p=0.16*
number of children	–	–	–	0.00	0.02	–
	–	–	–	(0.01)	(0.01)	–
				p=0.83	*p=0.05*	
family size	–	–	–	–	-0.01	–
	–	–	–	–	(0.01)	–
					p=0.05	
age of youngest child	–	–	–	–	0.00	–
child	–	–	–	–	(0.00)	–
					p=0.14	
quadratic in age	yes	yes	yes	yes	yes	yes
individual fixed effects	yes	yes	yes	yes	yes	yes
time effects	yes	yes	yes	yes	yes	yes
second-order terms	no	no	yes	yes	yes	yes
predicted consumption	no	no	no	no	no	yes
Observations	12,086	12,086	12,086	12,086	12,086	11,517

NOTE.–Results are for men. The dependent variable is (logged) hours of the male in household j in period t. $\log \tilde{w}_{mjt}$ is the predicted (logged) wage rate (see main text). A constant, individual fixed effects, time effects and a quadratic in age are included in all specifications but coefficients are not reported. Column 3 shows our preferred approach where the wage rate and hours worked of the individual's partner are included as well as higher-order terms. Columns 4 and 5 show robustness checks. Second-order terms are $(\log w_m)^2$, $(\log w_f)^2$, $(\log n_f)^2$, $\log w_m \times \log w_f$, $\log w_m \times \log n_f$, and $\log w_f \times \log n_f$. Standard errors in parentheses. p is the p-value of rejection the null hypothesis that the coefficient is zero.

shifter in our context. When the disutility from labor varies in a way that is correlated with (predicted) wage rates, this could bias the estimated labor-supply elasticity. As a consequence, the estimated effects of wages on labor supply could pick up a life-cycle profile in the taste for work which we avoid by including age in the labor-supply regression.[7]

Column 1 refers to a simple regression where we regress (log) hours worked on (log) wage rates, taking into account individual and time fixed effects as well as age. As discussed before, such a regression is unable to identify the Frisch elasticity as no variable is included that attempts to control for the marginal utility of wealth and hence income effects. As expected, we estimate the wage sensitivity of hours worked to be very small in this specification. The second column shows that the coefficient on the wage rate increases substantially when total household consumption is included and hence the income effect is taken into account. Such a regression is in the spirit of Altonji (1986) and yields an estimate for the Frisch elasticity of 0.26. In line with the presence of wealth effects, we estimate a negative coefficient on household consumption but the estimate is imprecise in this specification. Our analysis has shown that total household consumption is an imperfect proxy for individual consumption since commitment between spouses is limited. As discussed before, a regression using total household consumption leads to an underestimation of the Frisch elasticity. The relatively small value of 0.26 for the estimated Frisch elasticity is indicative for this downward bias.

Our Monte-Carlo studies have shown that time variation in bargaining weights can be captured to some extent in a regression that controls for hours worked and the wage rate of the individual's spouse, next to household consumption. A more accurate estimation of the Frisch elasticity is obtained when also higher-order terms are included. Our empirical results for the PSID are in line with these theoretical predictions. When we control for hours worked and the wage rate of the partner as well as for higher-order terms of these variables, which is our preferred approach, we estimate a Frisch elasticity of 0.71, see column 3 in Table 24. This estimate is substantially larger than the estimate from the Altonji (1986) regression in column 2 which does not take into account changes in bargaining weights. In our preferred specification where the household consumption variable is augmented by variables that capture changes in bargaining weights, the coefficient on household consumption turns statistically significant at the 10% level. Note that our preferred estimate for the Frisch elasticity is also very similar to our own estimate that we obtain using a different estimation approach in first-differences where one has to control for a bias due to borrowing constraints,

[7]While the level term in age is collinear to time effects and is hence already controlled for, we add the quadratic term in age to account for a non-linear pattern in the disutility of work over the life-cycle.

see Chapter II of this Thesis.[8]

Columns 4 and 5 report results for robustness checks that take into account additional potential taste shifters in labor supply. Other variables than age used in the wage regression may also be considered taste shifters in labor supply, an example is education. However, as they are time-invariant they are already accounted for in the labor-supply regression by the individual fixed effects. As potential taste shifters that display time variation, we control for the number of children, family size, and the age of the youngest child in the household. We find that this has close to no additional effect on our preferred estimate for the Frisch elasticity reported in column 3.

To make our setting directly comparable to Altonji (1986), we also consider a specification where we use predicted consumption rather than reported consumption. Altonji (1986) discusses that variation in the relative tastes for leisure and consumption can lead to a correlation between consumption and labor supply that does not necessarily reflect wealth effects only.[9] We follow Altonji (1986) and predict consumption using information on household wealth, family size, children, health and also include year effects (that capture the interest rates used by Altonji), and individual fixed effects (that capture race and characteristics of the place of residence which are used by Altonji).[10] In this specification, we estimate a similar Frisch elasticity as in our preferred specification, see column 6 of Table 24.

[8]The estimated coefficient on the partner's wage rate is similar to the one in our Monte-Carlo analysis but it is estimated imprecisely. The estimated coefficient on hours worked by the partner shows the opposite sign than in our Monte-Carlo studies which may be due to complementarities in leisure between spouses. Such complementarity would not impact on the informativeness of spousal hours about relative bargaining positions.

[9]Note that these arguments relate to the coefficient on consumption and not directly to the one on the wage rate. In addition, only those idiosyncratic fluctuations in taste shifters are concerned that are not already captured in time- or individual fixed effects.

[10]We construct assets as the sum of cash, bonds, stocks, the value of any business, vehicles, housing and other real estate, pension funds, and substract mortgages and other debts.

6 Conclusion

We have analyzed the consequence of limited commitment in the family for the estimation of labor-supply elasticities. In principle, the Frisch elasticity can be estimated in a regression of hours worked on the hourly wage rate where one controls for consumption as a measure of the marginal utility of wealth. We show that such a regression yields an unbiased estimate only when information on individual consumption is available or when there is commitment in the household. We have developed and calibrated a dual-earner household model with limited commitment to analyze the quantitative importance of the bias that results when commitment between household members is limited and household consumption is used as a proxy for household consumption. In such a misspecified model, we have found that estimated Frisch elasticities are around a third below their true values.

We have developed two improved estimation approaches that reduce or fully eliminate this bias. Our approaches exploit information on other household member's behavior to assess bargaining positions. We have performed Monte-Carlo experiments to study the quantitative performance of our estimation approaches. The best combination of close to unbiased estimates and empirical applicability is achieved by a simple linear labor-supply regression that includes hours worked and the wage rate of the individual's partner as well as second-order terms in these variable and the individual's own wage as additional controls. This approach is easy to apply and reduces the bias in the estimated Frisch elasticity to 7%.

We have applied this approach to U.S. panel data from the PSID and found Frisch elasticities for men of about 0.7. In line with the predictions of our model, these estimates are substantially larger than those obtained from labor-supply regressions that do not control for behavior of the partner and hence fail to account for the effects of limited commitment.

A Appendix

A.1 Intra-household decision-making under limited commitment

A.1.1 The maximization problem of the single household (outside option)

The outside option of agent ij in household j in a particular period is given by the expected lifetime utility of being single. The maximization problem of ij being single can be stated as

$$S_{ij}(w_{ij}) = \max_{c_{ij}, n_{ij}} \; u_{ij}(c_{ij}, n_{ij}) + \beta \cdot \mathbb{E}\left[S_{ij}\left(w'_{ij}\right)\right], \tag{108}$$

subject to the period budget-constraint

$$c_{ij} \leq w_{ij} n_{ij}. \tag{109}$$

The first-order conditions of the maximization problem are

$$w_{ij} = \frac{\partial u_{ij}(c_{ij}, n_{ij}) / \partial n_{ij}}{\lambda_{ij}}, \tag{110}$$

$$\lambda_{ij} = \frac{\partial u_{ij}(c_{ij}, n_{ij})}{\partial c_{ij}}, \tag{111}$$

where λ_{ij} denotes the Lagrange multiplier on the budget constraint, and the budget constraint. The resulting expected lifetime utility of agent ij, $S_{ij}(w_{ij})$, is used as the outside option in his or her participation constraint (92).

A.1.2 The maximization problem of the couple household under limited commitment

For the detailed presentation of the intra-household decision-making under limited commitment from Section 2, we put more structure in the model notation. A (sub-)index t of a variable denotes the value of the particular variable in period t. As described in Section 2, the wage rates of the spouses are stochastic. The current state of nature is denoted by s_t and implies a wage rate pair $\omega_j\left(s^t\right) = \left(w_{mj}\left(s^t\right), w_{fj}\left(s^t\right)\right)$. Note that the wage rates may depend on the history of states, for example when the wage rates include an autoregressive component. All decisions of the spouses depend also on the history of states and the current state of nature. The unconditional probability of a particular state is given by $\pi(s_t)$. The history of events until time t is $s^t = (s_0, ..., s_t)$. Thus, the unconditional probability of a particular history of states is $\pi\left(s^t\right)$. Therefore, the conditional probability of a particular history of states in period r, given the history of states in period t is $\pi\left(s^r\middle|\, s^t\right)$ with $\pi\left(s^t\middle|\, s^t\right)=1$.

The expected lifetime utility of spouse ij in couple household j at household formation is

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\left[u_{ij}\left(c_{ij}\left(s^t\right),n_{ij}\left(s^t\right)\right)\right],\text{ for } i=m,f.$$

The couple-household target function can be stated as

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\left[\mu_0\cdot u_{mj}\left(c_{mj}\left(s^t\right),n_{mj}\left(s^t\right)\right)+\left(1-\mu_0\right)\cdot u_{fj}\left(c_{fj}\left(s^t\right),n_{fj}\left(s^t\right)\right)\right],$$

using the fact that any (*ex-post*) efficient allocation of a couple household can be retrieved as the solution of a maximization problem where the target function of the couple household is a weighted sum of individual utility functions. $\mu_{mj}\left(s^{-1}\right)=\mu_0$ and $\mu_{fj}\left(s^{-1}\right)=1-\mu_0$ denote the initial Pareto weights of the male and the female, respectively. For now, we take the initial Pareto weights as given, we consider the determination of the Pareto weights below in the household formation. Note that initial Pareto weights are determined such that they are compatible with initial participation constraints:

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\, s^{-1}\right)\cdot\left[u_{ij}\left(c_{ij}\left(s^t\right),n_{ij}\left(s^t\right)\right)-u_{ij}^{single}\left(c_{ij}^{single}\left(s^t\right),n_{ij}^{single}\left(s^t\right)\right)\right]\geq 0,\forall i.$$

i.e., spouses have a higher expected lifetime utility within the couple household than being single. Note, the conditional probability $\pi\left(s^t\middle|\, s^{-1}\right)$ indicates that the participation constraints must be fulfilled once, at household formation. Further, note that the subindex s^{-1} implies that the initial participation constraint must hold only in expectations. The particular shock realizations in the first period, period $t=0$, does not matter.

If commitment between spouses is limited, expected lifetime utility of both spouses must be at least as high within the couple as in the outside option. Thus, there is an infinite number of additional participation constraints:

$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\, s^t\right)\cdot\left[\begin{array}{c}u_{ij}\left(c_{ij}\left(s^r\right),n_{ij}\left(s^r\right)\right)-\\ u_{ij}^{single}\left(c_{ij}^{single}\left(s^r\right),n_{ij}^{single}\left(s^r\right)\right)\end{array}\right]\geq 0,\forall s^t,\ \forall t,\ \forall i.$$

The conditional probability $\pi\left(s^r\middle|\, s^t\right)$ shows that the participation constraints must be fulfilled in every period, after any history of states, and in any state of nature. Using these participation constraints, and the budget constraint

$$c_{mj}\left(s^t\right)+c_{fj}\left(s^t\right)\leq w_{mj}\left(s^t\right)n_{mj}\left(s^t\right)+w_{fj}\left(s^t\right)n_{fj}\left(s^t\right)$$

we can finally state the maximization problem using the Lagrangian approach:

$$\begin{aligned}\mathcal{L} = & \sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\left\{\mu_0\cdot u_{mj}\left(c_{mj}\left(s^t\right),n_{mj}\left(s^t\right)\right)+\left(1-\mu_0\right)\cdot u_{fj}\left(c_{fj}\left(s^t\right),n_{fj}\left(s^t\right)\right)\right.\\ & +\phi_{mj}\left(s^t\right)\cdot\left(\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r|\,s^t\right)\cdot\left[\begin{array}{c}u_{mj}\left(c_{mj}\left(s^r\right),n_{mj}\left(s^r\right)\right)-\\ u_{mj}^{single}\left(c_{mj}^{single}\left(s^r\right),n_{mj}^{single}\left(s^r\right)\right)\end{array}\right]\right)\\ & +\phi_{fj}\left(s^t\right)\cdot\left(\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r|\,s^t\right)\cdot\left[\begin{array}{c}u_{fj}\left(c_{fj}\left(s^r\right),n_{fj}\left(s^r\right)\right)-\\ u_{fj}^{single}\left(c_{fj}^{single}\left(s^r\right),n_{fj}^{single}\left(s^r\right)\right)\end{array}\right]\right)\\ & \left.+\lambda_j\left(s^t\right)\cdot\left(w_{mj}\left(s^t\right)n_{mj}\left(s^t\right)+w_{fj}\left(s^t\right)n_{fj}\left(s^t\right)-c_{mj}\left(s^t\right)-c_{fj}\left(s^t\right)\right)\right\},\end{aligned}$$

where $\phi_{ij}\left(s^t\right)$ denotes the Kuhn-Tucker multiplier of spouse ij's participation constraint, and $\lambda_j\left(s^t\right)$ denotes the Lagrangian multiplier of the budget constraint. The initial Pareto weights μ_0 and $1-\mu_0$ guarantee that the initial participation constraints hold. However, in contrast to a situation of full commitment, participation constraints must be fulfilled even in the first period after the first shock realizations.

Marcet and Marimon (2019) have shown that this Lagrangian can be restated as

$$\begin{aligned}\mathcal{L} = & \sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\left\{\mu_{mj}\left(s^t\right)\cdot u_{mj}\left(c_{mj}\left(s^t\right),n_{mj}\left(s^t\right)\right)+\mu_{fj}\left(s^t\right)\cdot u_{fj}\left(c_{fj}\left(s^t\right),n_{fj}\left(s^t\right)\right)\right.\\ & -\phi_{mj}\left(s^t\right)\cdot\left(\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r|\,s^t\right)\cdot u_{mj}^{single}\left(c_{mj}^{single}\left(s^r\right),n_{mj}^{single}\left(s^r\right)\right)\right)\\ & -\phi_{fj}\left(s^t\right)\cdot\left(\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r|\,s^t\right)\cdot u_{fj}^{single}\left(c_{fj}^{single}\left(s^r\right),n_{fj}^{single}\left(s^r\right)\right)\right)\\ & \left.+\lambda_j\left(s^t\right)\cdot\left(w_{mj}\left(s^t\right)n_{mj}\left(s^t\right)+w_{fj}\left(s^t\right)n_{fj}\left(s^t\right)-c_{mj}\left(s^t\right)-c_{fj}\left(s^t\right)\right)\right\},\end{aligned}$$

with $\mu_{ij}\left(s^t\right)=\mu_{ij}\left(s^{t-1}\right)+\phi_{ij}\left(s^t\right)$, see Marcet and Marimon (2019) for details. The intuitive idea is that the first term of the participation constraints, which is the lifetime utility of being member of the couple household, can be summarized with the individual utility functions in the household target function, which is also the lifetime utility of being member of the couple household. Both terms can be summarized into one term when one allows the Pareto weight of the particular spouse to change whenever his or her participation constraint binds causing the adjustment "rule" of the Pareto weight $\mu_{ij}\left(s^t\right)=\mu_{ij}\left(s^{t-1}\right)+\phi_{ij}\left(s^t\right)$.

The recursive formulation of this problem is

$$
\begin{aligned}
V_j\left(\omega_j\left(s^t\right),\mu_{mj}\left(s^{t-1}\right),\mu_{fj}\left(s^{t-1}\right)\right) = {} & \mu_{mj}\left(s^t\right)\cdot u_{mj}\left(c_{mj}\left(s^t\right),n_{mj}\left(s^t\right)\right)+\\
& \mu_{fj}\left(s^t\right)\cdot u_{fj}\left(c_{fj}\left(s^t\right),n_{fj}\left(s^t\right)\right)+\\
& \beta\cdot\mathbb{E}\left[V\left(\omega_j\left(s^{t+1}\right),\mu_{mj}\left(s^t\right),\mu_{fj}\left(s^t\right)\right)\right]\\
& -\phi_{mj}\left(s^t\right)\cdot S_m\left(w_m\left(s^t\right)\right)-\phi_{fj}\left(s^t\right)\cdot S_f\left(w_f\left(s^t\right)\right)\\
& +\lambda_j\left(s^t\right)\cdot\begin{pmatrix} w_{mj}\left(s^t\right)n_{mj}\left(s^t\right)+w_{fj}\left(s^t\right)n_{fj}\left(s^t\right)- \\ c_{mj}\left(s^t\right)-c_{fj}\left(s^t\right)\end{pmatrix},
\end{aligned}
$$

where we use that

$$
S_{ij}\left(w_{ij}\left(s^t\right)\right)=\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_{mj}^{single}\left(c_{mj}^{single}\left(s^r\right),n_{mj}^{single}\left(s^r\right)\right).
$$

Since the outside options, the second terms of the participation constraints, do not depend on any decision of the couple household, they can be skipped from the Lagrangian as long as they are considered in the determination of the updated Pareto weights, $\mu_{ij}\left(s^t\right)$, i.e., as long as one controls whether the Kuhn-Tucker multipliers ϕ_{ij} are positive or zero. Finally, the maximization problem can be written as

$$
\begin{aligned}
V_j\left(\omega_j\left(s^t\right),\mu_{mj}\left(s^{t-1}\right),\mu_{fj}\left(s^{t-1}\right)\right) = {} & \mu_{mj}\left(s^t\right)\cdot u_{mj}\left(c_{mj}\left(s^t\right),n_{mj}\left(s^t\right)\right)+\\
& \mu_{fj}\left(s^t\right)\cdot u_{fj}\left(c_{fj}\left(s^t\right),n_{fj}\left(s^t\right)\right)+\\
& \beta\cdot\mathbb{E}\left[V\left(\omega_j\left(s^{t+1}\right),\mu_{mj}\left(s^t\right),\left(s^t\right)\right)\right]\\
& +\lambda_j\left(s^t\right)\cdot\begin{pmatrix} w_{mj}\left(s^t\right)n_{mj}\left(s^t\right)+w_{fj}\left(s^t\right)n_{fj}\left(s^t\right)- \\ c_{mj}\left(s^t\right)-c_{fj}\left(s^t\right)\end{pmatrix}.
\end{aligned}
$$

Skipping the state formulation, and using the notational convention $\mu'_{ij}=\mu_{ij}\left(s^t\right)$, the maximization problem becomes as in Section 2:

$$
V_j\left(\omega_j\left(s^t\right),\mu_{mj},\mu_{fj}\right)=\max_{c_{mj},c_{fj},n_{mj},n_{fj}} v\left(c_{mj},c_{fj},n_{mj},n_{fj},\mu'_{mj},\mu'_{fj}\right)+\beta\cdot\mathbb{E}\left[V\left(\omega'_j,\mu'_{mj},\mu'_{fj}\right)\right],
$$

with the updating rule for the Pareto weights

$$
\mu'_{ij}=\mu_{ij}+\phi_{ij},
$$

subject to the period budget

$$
c_{mj}+c_{fj}\leq w_{mj}n_{mj}+w_{fj}n_{fj}.
$$

A.1.3 Nash bargaining

We endogenize initial Pareto weights using Nash (1950)'s bargaining approach in a household formation step. Pareto weights are determined before shocks in the first period realize.[1]

Spouses only enter the couple household when there is a gain from marriage such that their expected lifetime utility of being married is higher than their corresponding value as a single. Formally, initial Pareto weights are determined such that the participation constraints

$$\mathbb{E}_0 V_{mj}\left(\omega_j, \mu_{ij0}, \mu_{-ij0}\right) \geq \mathbb{E}_0 S_m\left(w_m\right), \tag{112}$$

$$\mathbb{E}_0 V_f\left(\omega_j, \mu_{ij0}, \mu_{-ij0}\right) \geq \mathbb{E}_0 S_f\left(w_f\right), \tag{113}$$

are fulfilled. The resulting Nash program is

$$NP = \max_{\mu_{ij0}, \mu_{-ij0}} \left[\mathbb{E}_0\left(V_{mj}\left(\omega_j, \mu_{ij0}, \mu_{-ij0}\right) - S_m\left(w_m\right)\right)\right]^{0.5} \cdot \left[\mathbb{E}_0\left(V_{fj}\left(\omega_j, \mu_{ij0}, \mu_{-ij0}\right) - S_f\left(w_f\right)\right)\right]^{0.5}.$$

We assume that initial Pareto weights sum up to one such that $\mu_{-ij0} = 1 - \mu_{ij0}$.[2] Then, the Nash program simplifies to

$$NP = \max_{\mu_{ij0}} \left[\mathbb{E}_0\left(V_{mj}\left(\omega_j, \mu_{ij0}\right) - S_m\left(w_m\right)\right)\right]^{0.5} \cdot \left[\mathbb{E}_0\left(V_{fj}\left(\omega_j, \mu_{ij0}\right) - S_f\left(w_f\right)\right)\right]^{0.5},$$

where we have skipped the second Pareto weight argument since it is a function of the first Pareto weight argument. The first-order condition is

$$\frac{\mathbb{E}_0\left(V_{mj}\left(\omega_j, \mu_{ij0}\right) - S_m\left(w_m\right)\right)}{\mathbb{E}_0\left(V_{fj}\left(\omega_j, \mu_{ij0}\right) - S_f\left(w_f\right)\right)} = -\frac{\mathbb{E}_0\left(\partial V_{mj}\left(\omega_j, \mu_{ij0}\right)/\partial \mu_{ij0}\right)}{\mathbb{E}_0\left(\partial V_{mj}\left(\omega_j, \mu_{ij0}\right)/\partial \mu_{ij0}\right)}.$$

The Nash program and the resulting first-order condition can be used in the numerical solution to determine μ_{ij0}.

[1] We highlight this by using the time index 0.

[2] This assumption does not change the results, however, notation becomes easier.

A.2 Computation

Consumption and labor supply of the spouses can be written as a function of wage rates and Pareto weights. $\mathcal{C}_{ij}(\omega_j, \mu_{ij}, \mu_{-ij})$ denotes the consumption policy function of spouse ij, and $\mathcal{N}_{ij}(\omega_j, \mu_{ij}, \mu_{-ij})$ denotes the corresponding labor-supply policy function. Analogously, the policy functions for Bachelor households and for limited-commitment couples are given by $\mathcal{C}_{ij}^{S}(w_{ij})$, $\mathcal{N}_{ij}^{S}(w_{ij})$, and $\mathcal{C}_{ij}^{fc}(\omega_j, \bar{\mu}_{ij}, \bar{\mu}_{-ij})$, $\mathcal{N}_{ij}^{fc}(\omega_j, \bar{\mu}_{ij}, \bar{\mu}_{-ij})$, respectively. $\bar{\mu}_{ij}$ and $\bar{\mu}_{-ij}$ denote the Pareto weights under full commitment.

The first-order conditions of the maximization problem of the couple household can be rewritten as

$$c_{fj} = \left(\frac{\mu'_{mj}}{\mu'_{fj}}\right)^{-1/\sigma} \cdot c_{mj} = \left(x'_{mj}\right)^{-1/\sigma} \cdot c_{mj}, \tag{114}$$

$$w_{mj} = \frac{n_{mj}^{1/\eta_m}}{c_{mj}^{-\sigma}}, \tag{115}$$

$$w_{fj} = \frac{n_{fj}^{1/\eta_f}}{c_{fj}^{-\sigma}}, \tag{116}$$

$$c_{mj} + c_{fj} = w_{mj}n_{mj} + w_{fj}n_{fj}, \tag{117}$$

where x'_{mj} denotes the relative Pareto weight of the male. One can determine the intra-household allocation (114)-(117) using the relative Pareto weight instead of using two individual Pareto weights. Therefore, the state space can be reduced by one dimension such that one has to consider the wage pair ω_j and the relative Pareto weight of one spouse x_{ij}.

Algorithm 4 summarizes the numerical solution of the Bachelor household maximization problem from Section A.1.1. We need the solution to the Bachelor household problem as the outside options in the solution of the limited commitment couple model. Algorithm 5 describes the algorithm to solve the full commitment case. Since we abstract from savings opportunities, the policy functions $\mathcal{C}_{ij}^{fc}(\omega_j, \bar{x}_{ij})$ and $\mathcal{N}_{ij}^{fc}(\omega_j, \bar{x}_{ij})$, where $\bar{x}_{ij}$ denotes the relative Pareto weight of spouse ij under full commitment, describe the optimal behavior of unconstrained households. We can use these policy functions in the limited commitment algorithm to check where participation constraints bind. Algorithm 6 summarizes the numerical solution of the limited-commitment model.

Algorithm 4 (Bachelor households) *The algorithm consists of the following two steps:*

1: Set a convergence criterion ξ. Discretize the exogenous wage rate processes using the approach of Tauchen (1986).

2: Solve the non-linear first-order conditions (109)-(111) given the grids from step 1. The results are the optimal policy function $\mathcal{C}_{ij}^{S}(w_{ij})$, $\mathcal{N}_{ij}^{S}(w_{ij})$, $i \in \{m, f\}$.

3: Choose initial guesses for the value functions $S_{mj,0}(w_{mj})$, and $S_{fj,0}(w_{fj})$.

3a: Given that you are in iteration step k, use the Belman equation (108) to get:

$$S_{mj,k+1}(w_{mj}) = u_{mj}\left(\mathcal{C}_{mj}^{S}(w_{mj}), \mathcal{N}_{mj}^{S}(w_{mj})\right) + \beta \cdot \mathbb{E}\left[S_{mj,k}\left(w_{mj}'\right)\right], \text{ and}$$
$$S_{fj,k+1}(w_{fj}) = u_{fj}\left(\mathcal{C}_{fj}^{S}(w_{fj}), \mathcal{N}_{fj}^{S}(w_{fj})\right) + \beta \cdot \mathbb{E}\left[S_{fj,k}\left(w_{fj}'\right)\right].$$

3b: If sup $|S_{mj,k+1}(w_{mj}) - S_{mj,k}(w_{mj})| < \xi$, and sup $|S_{fj,k+1}(w_{fj}) - S_{fj,k}(w_{fj})| < \xi$ stop. Otherwise set $k = k + 1$ and go back to Step 3a. Iterate until convergence.

Algorithm 5 (Full commitment) *The algorithm consists of the following three steps:*

1: Set a convergence criterion ξ. Discretize the exogenous wage rate processes using the approach of Tauchen (1986). Build an equi-distant grid for the relative Pareto weight $\{x_1, ..., x_{N_x}\}$.

2: Solve the non-linear first-order conditions (91), and (97)-(99) given the grids from step 1. The results are the policy function $\mathcal{C}_{ij}^{fc}(\omega_j, \bar{x}_{ij})$ and $\mathcal{N}_{ij}^{fc}(\omega_j, \bar{x}_{ij})$.

3: Choose initial guesses for the value functions $V_{mj,0}^{fc}(\omega_j, \bar{x}_{ij})$, and $V_{fj,0}^{fc}(\omega_j, \bar{x}_{ij})$.

3a: Given that you are in iteration step k, use the Belman equation

$$V_{mj,k+1}^{fc}(\omega_j, \bar{x}_{ij}) = u_{mj}\left(\mathcal{C}_{mj}^{fc}, \mathcal{N}_{mj}^{fc}\right) + \Psi_j + \beta \cdot \mathbb{E}\left[V_{mj,k}^{fc}\left(\omega_j', \bar{x}_{ij}\right)\right], \text{ and}$$
$$V_{fj,k+1}^{fc}(\omega_j, \bar{x}_{ij}) = u_{fj}\left(\mathcal{C}_{fj}^{fc}, \mathcal{N}_{fj}^{fc}\right) + \Psi_j + \beta \cdot \mathbb{E}\left[V_{fj,k}^{fc}\left(\omega_j', \bar{x}_{ij}\right)\right],$$

to get updates of the value functions.[3]

3b: If sup $\left|V_{mj,k+1}^{fc}(\omega_j, \bar{x}_{ij}) - V_{mj,k}^{fc}(\omega_j, \bar{x}_{ij})\right| < \xi$, and sup $\left|V_{fj,k+1}^{fc}(\omega_j, \bar{x}_{ij}) - V_{fj,k}^{fc}(\omega_j, \bar{x}_{ij})\right| < \xi$ stop. Otherwise set $k = k + 1$ and go back to Step 3a. Iterate until convergence.

Algorithm 6 (Limited commitment) *The algorithm consists of the following eight steps:*

1: Set convergence criteria, ξ_1 and ξ_2. Discretize the exogenous income processes using the approach of Tauchen (1986). Build an equi-distant grid for the relative Pareto weight $\{x_1, ..., x_{N_x}\}$.

2: Bachelor household solution. Perform Step 2 and Step 3 of Algorithm 4 with $\xi = \xi_1$.

3: Full commitment solution. Perform Step 2 and Step 3 of Algorithm 5 with $\xi = \xi_1$.

4: Initialization limited commitment solution.

[3] We have skipped the arguments of the policy functions for notational convenience.

4a: Guess initial values for the value functions $V_{mj,0}(\omega_j, x_{ij})$*, and* $V_{fj,0}(\omega_j, x_{ij})$*.*

4b: Guess initial values for a state-dependent Pareto weight grid $\Theta_0(\omega_j, x_{ij})$*.*

4c: Construct a state-independent Pareto weight grid $\tilde{\Theta}$ *for the unconstrained case.*

5: Unconstrained solution (ignoring participation constraints). Given that you are in iteration step k*, use the Belman equation (93) to get:*

$$V_{mj,k+1}(\omega_j, x_{ij}) = u_{mj}\left(\mathcal{C}^{fc}_{mj}, \mathcal{N}^{fc}_{mj}\right) + \Psi_j + \beta \cdot \mathbb{E}\left[V_{mj,k}\left(\omega'_j, x_{ij}\right)\right], \text{ and}$$
$$V_{fj,k+1}(\omega_j, x_{ij}) = u_{fj}\left(\mathcal{C}^{fc}_{fj}, \mathcal{N}^{fc}_{fj}\right) + \Psi_j + \beta \cdot \mathbb{E}\left[V_{fj,k}\left(\omega'_j, x_{ij}\right)\right].$$

6: Checking the participation constraints.

6a: Find values on the grid of μ *such that the participation constraints hold with equality, using the value function* $V_{mj,k+1}(\omega_j, x_{ij})$ *and* $V_{fj,k+1}(\omega_j, x_{ij})$ *from Step 5. To do so, solve for the roots of*

$$V_{mj,k+1}(\omega_j, x_{ij}) - S_{mj}(w_{mj}) = 0,$$
$$V_{fj,k+1}(\omega_j, x_{ij}) - S_{fj}(w_{fj}) = 0,$$

where $S_{ij}(w_{ij})$ *denotes the outside option of spouse* i*.* $\underline{\nu}(\omega_j)$ *denotes the Pareto weight such that participation constraint of the husband holds with equality.* $\overline{\nu}(\omega_j)$ *denotes the Pareto weight such that participation constraint of the wife holds with equality.*

6b: Build a new grid $\Theta_{k+1}(\omega_j, x_{ij}) = \tilde{\Theta}$*. Replace all values for a particular* ω_j *which are smaller than* $\underline{\nu}(\omega_j)$ *by* $\underline{\nu}(\omega_j)$*. Replace all values for a particular* ω_j *which are larger than* $\overline{\nu}(\omega_j)$ *by* $\overline{\nu}(\omega_j)$*.*

7: Constrained case (respecting participation constraints). Constrained case (respecting participation constraints). Find all values of $V_{mj,k+1}(\omega_j, x_{ij})$ *smaller than* $S_{mj}(w_{mj})$*. Replace these values by* $S_{mj}(w_{mj})$*. Find all values of* $V_{fj,k+1}(\omega_j, x_{ij})$ *smaller than* $S_{fj}(w_{fj})$*. Replace these values by* $S_{fj}(w_{fj})$*.*

8: If sup $|V_{mj,k+1}(\omega_j, x_{ij}) - V_{mj,k}(\omega_j, x_{ij})| < \xi_1$*, sup* $|V_{fj,k+1}(\omega_j, x_{ij}) - V_{fj,k}(\omega_j, x_{ij})| < \xi_1$*, and sup* $|\Theta_{k+1}(\omega_j, x_{ij}) - \Theta_k(\omega_j, x_{ij})| < \xi_2$ *stop. Otherwise set* $k = k+1$ *and go back to Step 5. Iterate until convergence.*

A.3 Additional Monte-Carlo results

Table 25 shows the Monte-Carlo results for women and is the analogue to Table 23 which shows the results for men.

Table 25: Monte-Carlo results for women, from synthetic household panel data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	full commitment		limited commitment				
log wage rate $\log w_{fjt}$	1.05	1.05	1.05	0.71	0.90	1.16	1.05
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)
log individual	-1.58	–	-1.58	–	–	–	–
consumption $\log c_{fjt}$	(0.00)	–	(0.00)	–	–	–	–
log household	–	-1.58	–	-1.18	-3.38	-3.40	-1.58
consumption $\log c_{jt}$	–	(0.00)	–	(0.00)	(0.01)	(0.03)	(0.00)
log wage rate	–	–	–	–	1.29	1.16	–
of the spouse, $\log w_{mjt}$	–	–	–	–	(0.01)	(0.02)	–
log hours worked	–	–	–	–	-2.10	-1.69	–
by the spouse, $\log n_{mjt}$	–	–	–	–	(0.01)	(0.05)	–
$\log w_{fjt} \times \log w_{mjt}$	–	–	–	–	–	-0.03	–
	–	–	–	–	–	(0.01)	–
$\log w_{fjt} \times \log n_{mjt}$	–	–	–	–	–	0.23	–
	–	–	–	–	–	(0.01)	–
$\log w_{mjt} \times \log n_{mjt}$	–	–	–	–	–	-0.12	–
	–	–	–	–	–	(0.01)	–
$(\log w_{fjt})^2$	–	–	–	–	–	0.04	–
	–	–	–	–	–	(0.00)	–
$(\log w_{mjt})^2$	–	–	–	–	–	0.00	–
	–	–	–	–	–	(0.00)	–
$(\log n_{mjt})^2$	–	–	–	–	–	0.20	–
	–	–	–	–	–	(0.02)	–
deduced relative Pareto	–	–	–	–	–	–	1.58
weight term $\log(1 + x_{fjt}^{-1/\sigma})$	–	–	–	–	–	–	(0.00)
bias in %	–	–	–	-32	-14	11	–
fixed effects	yes	yes	yes	yes	yes	yes	yes
empirically feasible	no	yes	no	yes	yes	yes	yes

NOTE.–Results are for women. The dependent variable is (logged) hours of the female in household j in period t. A constant and household-type fixed effects are included but not reported. Table shows average estimates from 10,000 Monte-Carlo draws, with average standard errors in parentheses.

Chapter V

Savings and Risk Sharing in a Quantitative Dual-Earner Incomplete-Markets Economy with Limited Commitment

1 Introduction

Chiappori (1988) has extended Becker (1965)'s seminal model of the family by recognizing conflicts between household members and introduced collective models of the family that take into account household bargaining. The subsequent literature on collective models has assumed that family decision making is subject to full commitment between spouses. Full commitment means that, at the time of household formation, spouses choose a contingency plan for their future lives and can make binding promises to stick to this plan. This assumption has been challenged strongly by empirical studies, see, e.g., Dercon and Krishnan (2000), Duflo and Udry (2004), Mazzocco (2007), Robinson (2012), Voena (2015), Blau and Goodstein (2016), Cesarini et al. (2017), Lise and Yamada (2019), and Chiappori et al. (2019).

A number of recent papers have taken a new direction in modelling decision-making by taking into account that commitment within the household is limited. Chiappori and Mazzocco (2017) provide a first survey of the limited-commitment literature focusing on the general theoretical modelling strategy and applications. Full and limited-commitment models have in common that spouses join the household only if their expected lifetime utility within the couple is at least as high as their corresponding expected lifetime utility outside the couple, i.e., their participation constraints must be fulfilled. A key difference between both model types is the frequency of relevant participation constraints. In the full commitment case, participation constraints must be fulfilled only at the stage of household formation. Under limited commitment, by contrast, participation constraints must be fulfilled in every period and in any state of the world because household members are free to leave the household at any time (unilateral divorce). Taking into account these additional participation constraints affects both, behavior and welfare of couple households because spouses cannot commit to plans which may be beneficial to both spouses ex ante but later in life impose incentives for one of the partners to leave the household.

Ligon et al. (2002) have shown that the inability to fully commit to agreed-upon plans implies that intra-household allocations are *ex-ante* inefficient. Limited commitment can affect household decisions in many aspects of life, ranging from education choices to fertility decisions. Accordingly, limited commitment can induce various sources of inefficiencies. Some inefficiencies are of particular importance when markets are incomplete and individuals have an incentive to self-insure against idiosyncratic risk. Such a world is described by the heterogeneous-agents incomplete-markets model, the full-commitment version of which has become the leading tool for quantitative analysis in macroeconomics, see Heathcote et al. (2009) for a survey of this literature. When markets are incomplete, commitment problems between spouses can be expected to have a particular relevance for consumption risk sharing and for savings behavior.

In this paper, we introduce limited commitment in a dual-earner heterogeneous-agents incomplete-markets economy and investigate inefficiencies due to limited commitment in saving and risk-sharing behavior. This chapter is closely related to Ábrahám and Laczó (2018) who analyze the optimal savings behavior in a couple-household model with limited commitment between spouses. Ábrahám and Laczó (2018) show that couples save more when commitment between spouses is limited. Voena (2015) and Angelini et al. (2019) provide empirical evidence for this and document that unilateral divorce leads to increased savings by married couples. Theoretically, the higher savings incentives under limited commitment reflect a strategic motive to make divorce costly to the spouse who may be tempted to file for divorce in the future. The wealthier the household, the more this individual would lose upon divorce. Therefore, higher household savings relax future participation constraints. From an ex-ante perspective, this behavior is inefficient because it would not have been necessary if spouses could fully commit to plans in which case spouses could save less and enjoy more consumption earlier in life.[1]

The baseline setting considered by Ábrahám and Laczó (2018) is one where income shocks of spouses are perfectly negatively correlated such that total household income is constant, where spouses lose all assets upon divorce and where divorcees cannot save in the outside option. The assumption of constant household income is useful for their purposes because it shuts down precautionary savings under full commitment and hence isolates savings due to limited commitment. The assumptions regarding asset losses and saving possibilities upon divorce facilitate the derivation of analytical results and the numerical solution of the model. While Ábrahám and Laczó (2018) discuss that their main results carry over to more general set-ups, there is a still lack of quantitative results on savings and risk sharing under limited commitment. In our paper, we contribute by developing a generalized version of the Ábrahám and Laczó (2018) set-up where we introduce general stochastic income processes, relax the

[1] In the second part of their paper, Ábrahám and Laczó (2018) analyze hidden savings. While household-public assets are lost upon divorce, hidden private savings remain upon divorce. Ábrahám and Laczó (2018) show that hidden private savings do not occur in equilibrium since all incentives to save hiddenly are internalized by the optimal savings decision in household-public asset holdings.

assumption that all household assets are lost upon divorce, and allow divorcees to save. We use our model to provide a *quantitative* analysis of risk-sharing and precautionary savings under incomplete markets and limited commitment.

First, we quantify the theoretical result of Ábrahám and Laczó (2018) that limited commitment induces couple households to save more and and compare the quantitative importance of this saving motive to the well-studied precautionary saving motive. The latter is precluded in the Ábrahám and Laczó (2018) set-up because, there, precautionary savings are shut off under full commitment. Second, we perform a quantitative analysis of the study of property division rules upon divorce of Ábrahám and Laczó (2016) and how they affect behavior within the marriage.[2] While Ábrahám and Laczó (2016) relax the assumptions of Ábrahám and Laczó (2018) regarding asset losses and saving possibilities upon divorce, they continue to assume that income is constant at the household level and varies only at the individual level. Hence, the results from our generalized model complement their results quantitatively.

For our quantitative analysis of couple's savings behavior under limited commitment in a generalized set-up, we develop a practically feasible and fast numerical approach which constitutes another contribution of this chapter to the literature. Under limited commitment, the additional participation constraints complicate the numerical solution. Marcet and Marimon (2019) have shown that this can be handled by including the Pareto weights of spouses in the household's planner solution as an additional time-varying state variable. Previous papers on limited commitment have used this procedure in restrictive settings (Ábrahám and Laczó 2016, Mazzocco 2007) or have applied the simplifying restriction that the sum of spouses' Pareto weights is constant over time (Oikonomou and Siegel 2015, Foerster 2019, Obermeier 2019). For a number of applications, including ours, it is crucial to study the behavior of couple households with limited commitment in less restrictive environments. For example, Chiappori and Mazzocco (2017) call for recognizing the endogeneity of the outside option that arises from the accumulation of wealth within the household. This endogeneity is ruled out by assumption in the Ábrahám and Laczó (2018) framework. In particular with regard to savings behavior, the simplifying assumptions that Pareto weights sum up to one in each period used by Oikonomou and Siegel (2015), Foerster (2019) and Obermeier (2019) can lead to substantial biases in the model-predicted behavior of couple households.[3] In order to combine endogenous outside options and unrestricted endogenous bargaining weights in

[2]Ábrahám and Laczó (2016) is a follow-up paper to Ábrahám and Laczó (2018).

[3]For intertemporal decision problems, Marcet and Marimon (2019) have shown that Pareto weights grow over time such that they do not add up to one.

an infinite-horizon model, we extend Carroll (2006)'s endogenous grid method (EGM) to the case of incomplete-markets limited-commitment couple models.

Our results show that the savings effect of limited commitment is quite substantial. In our calibrated model, limited commitment induces couples to increase their yearly saving by about 4% and leads to an average rise in accumulated wealth of about 11%. This leads to two counteracting effects of limited commitment on consumption insurance. Higher wealth which is accumulated under limited commitment leads to a lower volatility of household consumption compared to the full-commitment case. However, time-varying individual consumption shares which reflect changes in relative bargaining positions induce individual consumption to vary more strongly than under full commitment. In our model simulations, the second effect dominates and individuals are effectively insured worse under limited commitment despite the fact that they live in wealthier households. Quantitatively, limited commitment reduces the insurance value of marriage by one fifth. In terms of welfare, we find that limited commitment reduces married couples' well-being by a consumption equivalent of about 0.3% – which corresponds to about one third of the costs of business cycles estimated by Krusell et al. (2009). Oversaving is responsible for about 40% of the welfare loss due to limited commitment.

Regarding the rules governing asset division upon divorce, we find that asset losses upon divorce and penalties for the spouse who files for divorce can have substantial impacts on decision making within the marriage. Both measures have the potential to reduce intra-household inefficiencies substantially. Large asset losses or penalties of about 70% can even raise consumption insurance within the marriage to the efficient full-commitment level. However, the largest efficiency gains are obtained for moderate losses or penalties. Losses or penalties in the range of 20% already close half of the efficiency gap in consumption insurance due to limited commitment. By contrast, uneven property division rules that do not depend on events within the marriage or on who files for divorce, affect efficiency within the marriage only to a small degree.

The remainder of this chapter is organized as follows. In Section 2, we develop a dual-earner incomplete markets economy with limited commitment. Section 3 provides a detailed discussion of the numerical issues and procedures to solve this model. Section 4 provides a quantitative comparison of savings behavior, risk sharing, and welfare under full and limited commitment. Section 5 studies rules determining property division upon divorce and their effects on saving and risk-sharing behavior within the marriage. Finally, Section 6 concludes.

2 A dual-earner incomplete markets economy with limited commitment

In this section, we present a dual-earner heterogeneous-agents incomplete-markets economy with limited commitment. In this economy, there are two channels of risk-sharing: intra-household risk-sharing and precautionary savings. We will show that limited commitment deteriorates insurance, causing welfare losses compared to an environment of full commitment.

Bachelor households. Households can save in a single non-state contingent asset with risk-free interest rate r. Borrowing is excluded. The maximization problem of a bachelor household is

$$S_i(\Omega_i) = \max_{c_i, a_i'} \; u_i(c_i) + \beta \cdot \mathbb{E}\left[S_i\left(\Omega_i'\right)\right], \tag{118}$$

subject to the period budget constraint

$$c_i + a_i' \leq y_i + (1+r)\, a_i, \tag{119}$$

and the borrowing constraint

$$a_i' \geq 0. \tag{120}$$

The state variables of an individual with gender i are summarized by $\Omega_i = (a_i, y_i)$. The first-order conditions of the maximization problem are

$$\lambda = \frac{\partial u_i(c_i)}{\partial c_i}, \tag{121}$$

$$\lambda - \xi = \beta \cdot (1+r) \cdot \mathbb{E}\lambda', \tag{122}$$

$$\xi \geq 0,\; a_i' \geq 0,\; \xi \cdot a_i' = 0, \tag{123}$$

and the budget constraint (119), given the income realizations and the initial asset values. λ is the Lagrange multiplier on the budget constraint. ξ denotes the Kuhn-Tucker multiplier on the borrowing constraint (120). $\mathcal{C}_i^S(\Omega_i)$ and $\mathcal{A}_i^S(\Omega_i)$ denote the optimal policy functions for consumption, labor supply, and savings. Lifetime utilities result from the Bellman equation (118) using the optimal policy functions

$$S_i(\Omega_i) = u_i\left(\mathcal{C}_i^S(\Omega_i)\right) + \beta \cdot \mathbb{E}\left[S_i\left(\Omega_i'\right)\right],$$

with $\Omega_i' = \left(\mathcal{A}_i^S\left(\Omega_i\right), y_i'\right)$. These lifetime utilities are used as outside options in the participation constraints of couples.

Full-commitment couple. Spouses face idiosyncratic income shocks, y_i, which follow a Markov process. The vector of joint income realizations is $\mathbf{y} = (y_m, y_f)$. Household preferences are given by the weighted sum of individual preferences

$$v\left(c_m, c_f, \mu_0^{fc}\right) = \mu_0^{fc} \cdot u_m\left(c_m\right) + \left(1 - \mu_0^{fc}\right) \cdot u_f\left(c_f\right) + \Psi. \tag{124}$$

The full-commitment couple household maximizes

$$V^{fc}\left(\Omega^{fc}, \mu_0^{fc}\right) = \max_{c_m, c_f, a'} \; v\left(c_m, c_f, \mu_0^{fc}\right) + \beta \cdot \mathbb{E}\left[V^{fc}\left(\Omega^{fc\prime}, \mu_0^{fc}\right)\right], \tag{125}$$

subject to the period budget constraint

$$c_m + c_f + a' \leq y_m + y_f + (1+r)\,a, \tag{126}$$

and subject to the joint borrowing constraint

$$a' \geq 0. \tag{127}$$

The first-order conditions are[1]

$$\lambda = \mu_0^{fc} \cdot \frac{\partial u_m\left(c_m\right)}{\partial c_m}, \tag{128}$$

$$\lambda = \left(1 - \mu_0^{fc}\right) \cdot \frac{\partial u_f\left(c_f\right)}{\partial c_f}, \tag{129}$$

$$\lambda - \xi = \beta \cdot (1+r) \cdot \mathbb{E}\lambda', \tag{130}$$

$$\xi \geq 0,\; a' \geq 0,\; \xi \cdot a' = 0,$$

and the budget constraint (126), given the vector of income realizations $\mathbf{y}$ and the initial asset values. λ is the Lagrange multiplier on the joint budget constraint. ξ denotes the Kuhn-Tucker multiplier on the joint borrowing constraint (127). Condition (130) is the Euler equation of a potentially borrowing constrained couple household.

(128) and (129) yield the risk-sharing condition

$$\frac{\partial u_m\left(c_m\right)/\partial c_m}{\partial u_f\left(c_f\right)/\partial c_f} = \frac{1 - \mu_0^{fc}}{\mu_0^{fc}}. \tag{131}$$

Under full commitment, risk-sharing is perfect since the right-hand side of (131) is constant

[1] A derivation can be found in Appendix A.2.1.

over time and over states. The full-commitment model is used as the Pareto efficient benchmark.

Limited-commitment couple. The opportunity to save implies that households take intertemporal decisions. Then, following Marcet and Marimon (2019), not only the relative Pareto weights matter but also the absolute value of Pareto weights. In a model with savings opportunities, Pareto weights need not sum up to one which reflects that the couple does not weigh each period equally. Specifically, periods where participation constraints are binding are weighted higher. This implies that couples have an incentive to transfer resources to these periods to avoid binding participation constraints. To capture these incentives, it is important to account for the absolute values of individual Pareto weights and the sum of the Pareto weights.

The timing is as follows. Spouses enter the period with state variables Ω. Then, income shocks realize. The updating of the Pareto weights depends on incomes and on asset holdings. In consequence, a given income combination may cause an updating of the Pareto weights in some situations but not in others, depending on asset holdings and potentially binding participation constraints.

The recursive formulation of the couple problem under limited commitment can be derived using the results of Marcet and Marimon (2019).[2] The key step is to include the participation constraints directly in the problem. The participation constraints are

$$\begin{aligned} u_m(c_m) + \Psi + \beta \cdot \mathbb{E}\left[V_m\left(\Omega', \mu_m', \mu_f'\right)\right] &\geq S_m\left(y_m, (1-\varphi)\cdot\kappa\cdot a\right) \\ u_f(c_f) + \Psi + \beta \cdot \mathbb{E}\left[V_f\left(\Omega', \mu_m', \mu_f'\right)\right] &\geq S_f\left(y_f, (1-\varphi)\cdot(1-\kappa)\cdot a\right), \end{aligned}$$

where $\Omega' = (a', \mathbf{y}')$ summarizes all state variables except the Pareto weights. Couples lose a constant share, φ, of all assets if they file for divorce. The remaining joint asset holdings are distributed with shares κ and $1-\kappa$ to the former husband and the former wife, respectively. Whether commitment issues become less or more severe due to the opportunity to save depends on the two parameters φ and κ. Commitment issues become on average less severe with higher asset losses upon divorce.[3]

[2] The transformation of the maximization problem to its Marcet and Marimon (2019) version can be found in Appendix A.1.

[3] Voena (2015) discusses different divorce law regimes with respect to their implications for the division of property upon divorce. We focus on unilateral divorce.

The recursive couple maximization problem including the participation constraints is

$$V(\Omega, \mu_m, \mu_f) = \max_{c_m, c_f, a} v\left(c_m, c_f, \mu_m', \mu_f'\right) - \nu_m \cdot S_m\left(y_m, (1-\varphi)\cdot\kappa\cdot a\right) - \\ \nu_f \cdot S_f\left(y_f, (1-\varphi)\cdot(1-\kappa)\cdot a\right) + \beta\cdot\mathbb{E}\left[V\left(\Omega', \mu_m', \mu_f'\right)\right], \tag{132}$$

subject to the period budget constraints (126) and the joint borrowing constraint (127), with

$$v\left(c_m, c_f, \mu_m', \mu_f'\right) = \mu_m' \cdot u_m(c_m) + \mu_f' \cdot u_f(c_f) + \left(\mu_m' + \mu_f'\right)\cdot\Psi,$$

where ν_g denotes the Kuhn-Tucker multiplier on the participation constraint of a spouse with gender g. In problem (132), the two additional terms correspond to the right-hand sides of the participation constraints. The left-hand sides of the participation constraints are included in the first term $v\left(c_m, c_f, \mu_m', \mu_f'\right)$ due to the update of the Pareto weights.

Marcet and Marimon (2019) show that the current period Pareto weight is the previous period Pareto weight plus the Kuhn-Tucker multiplier on the current period participation constraint. Thus, Pareto weights evolve according to $\mu_g' = \mu_g + \nu_g$, with $g = \{m, f\}$. Initial Pareto weights are given by $\mu_{m0} = \mu_0$ and $\mu_{f0} = 1 - \mu_0$. Although the Pareto weights are unbounded, the relative Pareto weight is bounded. Using the results of Ligon et al. (2002), one can show that state-dependent optimal intervals of the relative Pareto weight exist implying an updating rule for the relative Pareto weight. The relative male Pareto weight is defined as

$$x' = \frac{\mu_m'}{\mu_f'} = x \cdot \frac{1-\upsilon_f}{1-\upsilon_m}, \tag{133}$$

with $\upsilon_g = \nu_g/\mu_g'$. Following Ábrahám and Laczó (2018), We use the relative Pareto weight x as a co-state variable instead of the two Pareto weights. This allows to reduce the number of relevant state variables from two, μ_m and μ_f, to one, x. This approach is feasible if one uses additionally the term υ_g to characterize the Euler equation. We will discuss this approach in detail in Section 3.

The optimal state-dependent intervals of the relative Pareto weight x are given by

$$x' = \begin{cases} \underline{x}'(\Omega), & \text{if } x < \underline{x}'(\Omega), \\ x, & \text{if } x \in \left[\underline{x}'(\Omega), \overline{x}'(\Omega)\right], \\ \overline{x}'(\Omega), & \text{if } x > \overline{x}'(\Omega), \end{cases} \tag{134}$$

with the lower and upper bounds of the intervals implicitly defined by

$$u_f(c_f) + \Psi + \beta \cdot \mathbb{E}\left[V_f\left(\Omega', \overline{x}'(\Omega)\right)\right] = S_f\left(y_f, (1-\varphi)\cdot(1-\kappa)\cdot a\right), \text{ and}$$
$$u_m(c_m) + \Psi + \beta \cdot \mathbb{E}\left[V_m\left(\Omega', \underline{x}'(\Omega)\right)\right] = S_m\left(y_m, (1-\varphi)\cdot\kappa\cdot a\right).$$

Thus, $\overline{x}'(\Omega)$ is the maximum value of the relative male Pareto weight such that the wife is indifferent between staying in and leaving the couple. $\underline{x}'(\Omega)$ is analogously the minimum value of the relative male Pareto weight such that he is indifferent between staying in and leaving the couple. We use the optimal intervals in the numerical implementation of the model.

The first-order conditions of the couple maximization problem are[4]

$$\lambda = \mu_m' \cdot \frac{\partial u_m(c_m)}{\partial c_m}, \tag{135}$$
$$\lambda = \mu_f' \cdot \frac{\partial u_f(c_f)}{\partial c_f}, \tag{136}$$
$$\lambda - \xi = \beta \cdot \mathbb{E}\left[(1+r)\cdot\lambda' - \nu_m' \cdot \frac{\partial S_m}{\partial a'} - \nu_f' \cdot \frac{\partial S_f}{\partial a'}\right], \tag{137}$$
$$\xi \geq 0,\ a' \geq 0,\ \xi \cdot a' = 0,$$

and the budget constraint (126), given the vector of income realizations $\mathbf{y}$ and the initial asset value. λ is the Lagrange multiplier on the budget constraint. ξ denotes the Kuhn-Tucker multiplier on the joint borrowing constraint (120). In comparison to the model with full commitment, there are two additional terms on the right-hand side of the Euler equation (137). Spouses take into account in their expectations that savings may have an impact on the Pareto weights in the next period. There are two different saving incentives. First, a short-run incentive depends on whether spouses expect that the Pareto weight will change in the next period or not. If for example a spouse with a relatively high Pareto weight expects that her Pareto weight will decrease in the next period, the person has an incentive to save less since she would receive only a smaller share of these future asset holdings than in the current period. Thus, an expected decrease of the Pareto weight may reduce savings. This incentive is primarily captured by the first term in the brackets on the right-hand side of the Euler equation. The second, longer-term, saving incentive is caused by the impact of savings on future commitment issues. When higher savings reduce the likelihood of binding participation constraints in the future, spouses have an incentive to save more. In contrast, more assets may also increase the likelihood of binding participation constraints. This longer

[4]A derivation can be found in Appendix A.2.2.

term saving incentive is captured by the second and third term in the brackets on the right-hand side of the Euler equation. Whether more or less savings increase or decrease future commitment issues is a quantitative question. In any case, however, savings behavior is inefficient due to the additional wedge in the Euler equation which is absent in the Pareto efficient full-commitment benchmark case.

Conditions (135) and (136) imply the risk-sharing condition

$$\frac{\partial u_m\left(c_m\right)/\partial c_m}{\partial u_f\left(c_f\right)/\partial c_f}=\frac{\mu_f'}{\mu_m'}. \tag{138}$$

Risk-sharing is imperfect as the Pareto weights vary over time and states, resulting in a welfare loss compared to the case of full commitment.

Individual lifetime utilities result from the Bellman equation (132) by

$$V_i\left(\Omega,\mu_m,\mu_f\right)=u_i\left(\mathcal{C}_i\left(\Omega,\mu_m,\mu_f\right)\right)+\Psi+\beta\cdot\mathbb{E}\left[V_i\left(\Omega',\mu_m',\mu_f'\right)\right], \tag{139}$$

using the policy functions $\mathcal{C}_i\left(\Omega,\mu_m,\mu_f\right)$ and $\mathcal{A}_i\left(\Omega,\mu_m,\mu_f\right)$, with $\Omega'=\left(\mathcal{A}_i\left(\Omega,\mu_m,\mu_f\right),\mathbf{y}'\right)$. These individual lifetime utilities are used to check the participation constraints and for welfare evaluations.

3 Numerical algorithm

Knowing current state variables and the optimal savings decision described by the Euler equation, all other decisions follow from intratemporal conditions. Thus, the algorithm is based on the idea to find optimal savings first and thereafter determine all other decisions. The state space of the couple maximization problem is $\mathbf{N} = N_{y_m} \times N_{y_f} \times N_x \times N_a$, where N_{y_i} denotes the number of income grid points of spouse i, N_x is the number of grid points of the relative Pareto weight, and N_a denotes the number of asset grid points.

3.1 Full commitment

Algorithm 7 summarizes the numerical solution to the full-commitment maximization problem.[1]

Algorithm 7 (Full commitment) *The algorithm consists of the following four steps:*

1: Set convergence criteria ζ_1 and ζ_2. Discretize the exogenous income processes using the approach of Tauchen (1986). Build a grid for the relative Pareto weight $\{x_1, ..., x_{N_x}\}$. Build a grid for current period assets a, $\{a_1, ..., a_{N_a}\}$. Build a grid for next period assets a', $\{\tilde{a}_1, ..., \tilde{a}_{N_a}\}$.

2: EGM step.

2a: Choose an initial guess for the asset policy function $\mathcal{A}_0^{fc}\left(\Omega^{fc}, x_0^{fc}\right)$.

2b: Given that you are in iteration step j, find male or female consumption in the next period as a function of next period incomes and next period asset holdings using the budget constraint (126) and the risk-sharing condition (131), given next period assets a' by the exogenous grid $\{\tilde{a}_1, ..., \tilde{a}_{N_a}\}$ and the guess for the asset policy function $\mathcal{A}_j^{fc}\left(\Omega^{fc}, x_0^{fc}\right)$.

2c: Use the Euler equation (130) to calculate current period male or female consumption as a function of current period incomes and next period assets. Use the risk-sharing condition (131) to find the consumption of the other spouse.

2d: Calculate the endogenous grid of current period assets $\hat{a}$ using the current period budget constraint (126). This yields a mapping of the endogenous grid to the exogenous grid of a': $a'(\hat{a})$.

2e: Use $a'(\hat{a})$ to interpolate a' on the exogenous grid $\{a_1, ..., a_{N_a}\}$ to find the asset policy function $\mathcal{A}_{j+1}^{fc}\left(\Omega, x_0^{fc}\right)$.

2f: Check whether the borrowing constraint (127) is binding. If the borrowing constraint is binding, replace all values of $\mathcal{A}_{j+1}^{fc}\left(\Omega^{fc}, x_0^{fc}\right)$ smaller than the borrowing limit by the borrowing limit.

[1] The algorithm for Bachelor households can be found in Appendix A.3.

2g: If $\sup\left|\mathcal{A}^{fc}_{j+1}\left(\Omega^{fc}, x^{fc}_0\right) - \mathcal{A}^{fc}_j\left(\Omega^{fc}, x^{fc}_0\right)\right| < \zeta$ *stop, set* $\mathcal{A}^{fc}\left(\Omega^{fc}, x^{fc}_0\right) = \mathcal{A}^{fc}_{j+1}\left(\Omega^{fc}, x^{fc}_0\right)$. *Otherwise set* $j = j + 1$ *and go back to step 2b. Iterate until convergence.*

3: Find policy functions for consumption, $\mathcal{C}^{fc}_i\left(\Omega^{fc}, x^{fc}_0\right)$*, using the period budget constraint (126) and the risk-sharing condition (131).*

4: Choose initial guesses for the value functions $V^{fc}_{m,0}\left(\Omega^{fc}, x^{fc}_0\right)$ *and* $V^{fc}_{f,0}\left(\Omega^{fc}, x^{fc}_0\right)$.

4a: Given that you are in iteration step j*, use the Bellman equation (125) to get individual value functions:*

$$V^{fc}_{m,j+1}\left(\Omega^{fc}, x^{fc}_0\right) = u_m\left(\mathcal{C}^{fc}_m\left(\Omega^{fc}, x^{fc}_0\right)\right) + \beta \cdot \mathbb{E}\left[V^{fc}_{m,j}\left(\Omega^{fc\prime}, x^{fc}_0\right)\right], \text{ and}$$
$$V^{fc}_{f,j+1}\left(\Omega^{fc}, x^{fc}_0\right) = u_f\left(\mathcal{C}^{fc}_f\left(\Omega^{fc}, x^{fc}_0\right)\right) + \beta \cdot \mathbb{E}\left[V^{fc}_{f,j}\left(\Omega^{fc\prime}, x^{fc}_0\right)\right].$$

4b: If $\sup\left|V^{fc}_{m,j+1}\left(\Omega^{fc}, x^{fc}_0\right) - V^{fc}_{m,j}\left(\Omega^{fc}, x^{fc}_0\right)\right| < \zeta_2$*, and* $\sup\left|V^{fc}_{f,j+1}\left(\Omega^{fc}, x^{fc}_0\right) - V^{fc}_{f,j}\left(\Omega^{fc}, x^{fc}_0\right)\right| < \zeta_2$ *stop. Otherwise set* $j = j + 1$ *and go back to Step 4a. Iterate until convergence.*

In Step 1 of Algorithm 7, we construct a grid with an odd number of grid points and monotonic increasing grid points for the Pareto weight such that $x_1 = 1/x_{N_x}$, $x_2 = 1/x_{N_x-1}$, and so on.[2] We use a grid with monotonic increasing grid points for the assets to achieve more grid points near the borrowing constraint to capture the non-linearities of the policy functions near the borrowing constraint.

Steps 2b to Step 2g are a policy function version of Carroll (2006)'s endogenous grid method (EGM, henceforth). While Carroll (2006) uses a value function approach, our approach combines the speed advantages of both a policy function approach (see Heer and Maussner 2009) and EGM. In this approach, we iterate over the asset policy function until convergence. The EGM approach is based on the idea that the dynamic problem is inversed. Typically, one searches for the solution to a dynamic problem given current period exogenous and endogenous state variables. Carroll (2006) inverses the dynamic problem and discretizes the next period endogenous state variable.[3] Taken this next period endogenous state and the current period exogenous state as given, Carroll (2006) solves for the current period endogenous state as the endogenous grid. The idea is: given my next period asset stock is a_{next}

[2] The advantage of such a grid is that solutions are symmetric if there is no parameter heterogeneity between spouses. Other grid choices are possible. We checked that our numerical results do not depend on a specific choice for the grids.

[3] Carroll (2006) considers a dynamic problem with a single endogenous state variable. Generalizations for more endogenous state variables are available, see Hintermaier and Koeniger (2010), Druedahl and Jørgensen (2017), and Ludwig and Schön (2018). However, several endogenous state variables reduce the speed gain from EGM compared to more traditional solution methods like value function iteration substantially.

and my current period exogenous income is w, how high must my current asset stock be if I behave optimally.

When the asset policy function has converged, all other policy functions can be calculated using the period budget constraint (126) and the risk-sharing condition (131). Finally, individual value functions are recovered by iterating the Bellman equation (125).

We illustrate the EGM step (Step 2 of Algorithm 7) using the utility function

$$u\left(c_i\right)=\frac{c_i^{1-\sigma}-1}{1-\sigma}. \tag{140}$$

Given an arbitrary initial guess for the asset policy function in iteration step j $\mathcal{A}_j^{fc}\left(\Omega^{fc},x_0^{fc}\right)$, one has to solve for male consumption in the next period as a function of next period state variables in Step 2b of Algorithm 7. The utility function (140) implies that female consumption, male labor supply, and female labor supply can be stated as function of male consumption using the first-order conditions.[4]

The risk-sharing condition (131) implies

$$c_f=\left(x_0^{fc}\right)^{-1/\sigma}\cdot c_m, \tag{141}$$

where x_0^{fc} denotes the relative Pareto weight under full commitment. Using the policy function $\mathcal{A}_j^{fc}$[5], the budget constraint (126) becomes

$$c_m+\left(x_0^{fc}\right)^{-1/\sigma}\cdot c_m+\mathcal{A}_j^{fc}=y_m+y_f+(1+r)\,a. \tag{142}$$

Given the beginning-of-period state variables y_m, y_f, and a and the asset policy function $\mathcal{A}_j^{fc}$, one can determine male consumption as the solution to the non-linear equation (142). In Step 2b of Algorithm 7, we determine next period male consumption as a function of next period incomes and asset holdings, given the asset policy function $\mathcal{A}_j^{fc}$. To do this, we solve for male consumption in (142) updated one period ahead

$$c_m'+\left(x_0^{fc}\right)^{-1/\sigma}\cdot c_m'+\mathcal{A}_j^{fc}=y_m'+y_f'+(1+r)\,a'.$$

We use a fast vectorized bisection root-finding routine to determine this consumption policy function.

We search for current period consumption as a function of current period exogenous states and next period asset holdings in Step 2c of Algorithm 7. Given the utility function, the Euler

[4] More general, three of the four variables, male consumption, female consumption, male labor supply, and female labor supply, can be stated as a function of the fourth variable.

[5] We skip the functional notation to simplify the notation.

equation is

$$c_m^{-\sigma} = \beta \cdot (1+r)\,\mathbb{E}\left[\left(c_m'\right)^{-\sigma}\right]. \tag{143}$$

The borrowing constraint is ignored in this step. Potential negative savings are handled in Step 2f of Algorithm 7. Given c_m' as a function of next period state variables and using the Markovian structure of the stochastic component of the income, one can calculate c_m as a function of current period exogenous states and next period asset holdings. Then, using conditions (141), one gets female consumption as a function of current period exogenous states and next period asset holdings.

Finally, one can use the budget constraint (126) to calculate the endogenous grid $\hat{a}$ analytically in Step 2d of Algorithm 7:

$$\hat{a} = \frac{c_m + c_f + a' - y_m - y_f}{1+r}.$$

Since all variables are functions of current period exogenous states and next period asset holdings, we have a one-to-one mapping for each income combination from the endogenous grid $\hat{a}$ to next period assets: $a'(\hat{a})$.

This mapping can be used in Step 2e of Algorithm 7 to find a new asset policy function using interpolation methods. We use linear interpolation and interpolate $a'(\hat{a})$ on the exogenous grid $\{a_1, ..., a_{N_a}\}$ to get $\mathcal{A}_{j+1}^{fc}\left(\Omega, \mu_0^{fc}\right)$.

In Step 2f of Algorithm 7 it is checked whether borrowing constraints are binding. All values $\mathcal{A}_{j+1}^{fc} < 0$ are replaced by 0. In Step 2g of Algorithm 7 it is checked whether the asset policy function has converged or not. If not, a new iteration step is started using the new asset policy function as initial guess and Step 2b to Step 2f are performed again. If yes, one calculates the final policy functions for c_m and c_f in Step 3 of Algorithm 7 using the budget constraint (126) and the risk-sharing condition 141). In the final Step 4 of Algorithm 7, value functions are obtained by iterating the Bellman equation (125) for given policy functions.

3.2 Limited commitment

In the limited commitment model, the participation constraints have to be checked in each iteration step and the optimal savings decisions, the value functions and the state-dependent grids of the relative Pareto weight have to be determined simultaneously. In addition, iterating on the Euler equation (137) requires initial values of the redefined Kuhn-Tucker multipliers of the participation constraints, υ_g, in each iteration step. Thus, one has to iterate until six value and policy functions have jointly converged.[6] In the full-commitment case, one can iterate on the optimal savings decisions and on the value functions separately which is much faster. Algorithm 8 summarizes the numerical solution of the limited-commitment maximization problem.

Algorithm 8 (Limited commitment) *The algorithm consists of the following eight steps:*

1: *Set convergence criteria ζ_1, ζ_2, ζ_3, and ζ_4. Discretize the exogenous income processes using the approach of Tauchen (1986). Build a grid for the relative Pareto weight $\{x_1, ..., x_{N_x}\}$. Build a grid for current period assets a, $\{a_1, ..., a_{N_a}\}$. Build a grid for next period assets a', $\{\tilde{a}_1, ..., \tilde{a}_{N_a}\}$.*

2: *Bachelor household solution.*

3: *Full-commitment solution.*

   3a: *Perform Step 2 of Algorithm 7.*

   3b: *Perform Step 3 of Algorithm 7.*

   3c: *Perform Step 4 of Algorithm 7.*

4: *Initialization limited-commitment solution.*

   4a: *Guess initial value functions $V_{m,0}(\Omega, x)$, and $V_{f,0}(\Omega, x)$.*

   4b: *Guess initial values for a state-dependent relative Pareto weight grid $\Theta_0(\Omega, x)$.*

   4c: *Construct a state-independent relative Pareto weight grid $\tilde{\Theta}$ for the unconstrained case.*

   4d: *Guess initial asset policy function $\mathcal{A}_0(\Omega, x)$.*

   4e: *Guess initial values for the multipliers υ_m and υ_f, $\upsilon_{m,0}(\Omega, x) = \frac{\nu_m(\Omega,x)}{\mu'_m(\Omega,x)}$ and $\upsilon_{f,0}(\Omega, x) = \frac{\nu_f(\Omega,x)}{\mu'_f(\Omega,x)}$.*

5: *Unconstrained solution (ignoring participation constraints). Given that you are in iteration step j, perform the following steps:*

   5a: *Given next period assets a' by the exogenous grid $\{\tilde{a}_1, ..., \tilde{a}_{N_a}\}$, the guess for the asset policy function $\mathcal{A}_j(\Omega, x)$, and the guess for the state-dependent relative Pareto weight grid $\Theta_j(\Omega, x)$, find male or female consumption in the next period as function of next period incomes and next period asset holdings using the budget constraint (126) and the risk-sharing condition (138).*

   5b: *Use the Euler equation (137) with $\upsilon_{m,j}(\Omega, x)$ and $\upsilon_{f,j}(\Omega, x)$ to calculate current period male or female consumption as a function of current period incomes and next period assets. Use the risk-sharing condition (138) to find consumption of the other spouse.*

[6] Convergence of the redefined Kuhn-Tucker multipliers is guaranteed by the convergence of the state-dependent grids of the relative Pareto weight.

5c: *Calculate the endogenous grid of current period assets* $\hat{a}$ *using the current period budget constraint (126). This gives a mapping from the endogenous grid to the exogenous grid of* a': $a'(\hat{a})$.

5d: *Upper envelope step. Perform Algorithm 9 to find the monotone endogenous grid* $\acute{a}$.

5e: *Use* $a'(\acute{a})$ *to interpolate* a' *on the exogenous grid* $\{a_1, ..., a_{N_a}\}$ *to find the asset policy function* $\mathcal{A}^{uncon}(\Omega, x)$.

5f: *Check whether the borrowing constraint (127) is binding. Replace all values of* $\mathcal{A}^{uncon}(\Omega, x)$ *smaller than the borrowing limit by the borrowing limit.*

5g: *Use the budget constraint (126) and the risk-sharing condition (138) to find* $\mathcal{C}_m^{uncon}(\Omega, x)$ *and* $\mathcal{C}_f^{uncon}(\Omega, x)$.

5h: *Use the Bellman equation to get:*

$$V_m^{uncon}(\Omega, x) = u_m\left(\mathcal{C}_m^{uncon}(\Omega, x)\right) + \beta \mathbb{E} V_{m,j}\left(\Omega', x\right),$$
$$V_f^{uncon}(\Omega, x) = u_f\left(\mathcal{C}_f^{uncon}(\Omega, x)\right) + \beta \mathbb{E} V_{f,j}\left(\Omega', x\right).$$

6: *Check participation constraints.*

6a: *Find values of* x *such that the participation constraints hold with equality, using the value functions* $V_m^{uncon}(\mathbf{y}, x)$, *and* $V_f^{uncon}(\mathbf{y}, x)$ *from Step 5. To do so, solve for the roots of*

$$V_m^{uncon}(\Omega, x) - S_m\left(y_m, (1-\varphi)\cdot\kappa\cdot a\right) = 0,$$
$$V_f^{uncon}(\Omega, x) - S_f\left(y_f\,(1-\varphi)\cdot(1-\kappa)\cdot a\right) = 0,$$

where S_i *denotes the outside option of spouse* i. $\underline{\nu}(\Omega)$ *denotes the relative Pareto weight such that the participation constraint of the husband holds with equality.* $\overline{\nu}(\Omega)$ *denotes the relative Pareto weight such that the participation constraint of the wife holds with equality.*

6b: *Build a new grid* $\Theta_{j+1}(\Omega, x) = \tilde{\Theta}$. *Replace all values for a particular state combination* Ω *which are smaller than* $\underline{\nu}(\Omega)$ *by* $\underline{\nu}(\Omega)$. *Replace all values for a particular* Ω *which are larger than* $\overline{\nu}(\Omega)$ *by* $\overline{\nu}(\Omega)$.

7: *Constrained case (respecting participation constraints). Given that you are in iteration step* j, *perform the following steps:*

7a: *Set* $V_{m,j+1}(\Omega, x) = V_m^{uncon}(\Omega, x)$ *and* $V_{f,j+1}(\Omega, x) = V_f^{uncon}(\Omega, x)$. *Find all values of* $V_{m,j+1}(\Omega, x)$ *smaller than* $S_m(y_m, (1-\varphi)\cdot\kappa\cdot a)$ *and replace these values by* $S_m(y_m, (1-\varphi)\cdot\kappa\cdot a)$. *Adjust the corresponding values of* $V_{f,j+1}(\Omega, x)$ *by interpolating* $V_f^{uncon}(\Omega, x)$ *on the corresponding relative Pareto weights. Find all values of* $V_{f,j+1}(\Omega, x)$ *smaller than* $S_f(y_f\,(1-\varphi)\cdot(1-\kappa)\cdot a)$ *and replace these values by* $S_f(y_f\,(1-\varphi)\cdot(1-\kappa)\cdot a)$. *Adjust the corresponding values of* $V_{m,j+1}(\Omega, x)$ *by interpolating* $V_m^{uncon}(\Omega, x)$ *on the corresponding relative Pareto weights.*

7b: *Interpolate* $\mathcal{A}^{uncon}(\Omega, x)$ *on the new state-dependent relative Pareto weight grid* $\Theta_{j+1}(\Omega, x)$ *to get* $\mathcal{A}_{j+1}(\Omega, x)$.

7c: Update policy functions for υ_m and υ_f using the definition of the relative Pareto weight (133):

$$\Theta_{j+1}(\Omega,x)=\frac{\mu'_m(\Omega,x)}{\mu'_f(\Omega,x)}=\tilde{\Theta}\cdot\frac{1-\upsilon_{f,j+1}(\Omega,x)}{1-\upsilon_{m,j+1}(\Omega,x)},$$

for $\upsilon_{m,j+1}(\Omega,x)$ and $\upsilon_{f,j+1}(\Omega,x)$.

8: If sup $|\mathcal{A}_{j+1}(\Omega,x)-\mathcal{A}_j(\Omega,x)|<\zeta_1$, sup $|V_{m,j+1}(\Omega,x)-V_{m,j}(\Omega,x)|<\zeta_2$, sup $|V_{f,j+1}(\Omega,x)-V_{f,j}(\Omega,x)|$ ζ_2, sup $|\Theta_{j+1}(\Omega,x)-\Theta_j(\Omega,x)|<\zeta_3$, sup $|\upsilon_{m,j+1}(\Omega,x)-\upsilon_{m,j}(\Omega,x)|<\zeta_4$, and sup $|\upsilon_{f,j+1}(\Omega,x)-\upsilon_{f,j}($ ζ_4 stop. Otherwise set $j=j+1$ and go back to Step 5. Iterate until convergence.

Step 1, Step 2, and Step 3 of Algorithm 8 follow from Algorithm 7 and Appendix A.3. We use the same grids for the relative Pareto weight and assets as for the full-commitment solution. The solution of the bachelor model in Step 2 of Algorithm 8 is used in Step 5b in the Euler equation and to check borrowing constraints in Step 6a.

We use the lifetime utilities of the full-commitment couple as initial guesses for the value functions in the limited-commitment solution. We use the grid of the relative Pareto weight from Step 1 of Algorithm 8, $\{x_1,...,x_{N_x}\}$, as initial guess of the state-dependent relative Pareto weight grid in Step 4b of Algorithm 8. The same grid should be chosen for the state-independent relative Pareto weight grid, $\tilde{\Theta}$. We use the full commitment optimal savings decision as initial guess for the limited-commitment model in Step 4d of Algorithm 8, $\mathcal{A}_0(\Omega,x)=\mathcal{A}^{fc}(\Omega,x)$. We initialize the redefined Kuhn-Tucker multipliers υ_m and υ_f compatible to the full-commitment model where both multipliers are always equal to zero, independent of the particular state variables, $\upsilon_{m,0}(\Omega,x)=0$ and $\upsilon_{f,0}(\Omega,x)=0$.[7]

Steps 5 to Step 8 of Algorithm 8 are our EGM solution approach to the limited-commitment couple model. The solution algorithm is based on the idea that beginning with the next period commitment between spouses is limited. To illustrate the differences between the full commitment and the limited-commitment solution, we use the same utility function, see (140).

In Step 5a of Algorithm 8 one has to calculate next period male or female consumption as a function of next period state variables. Since commitment between spouses is limited in the next period and thereafter, one needs the value of the relative Pareto weight with which the couple starts in the next period. We use the initialized Pareto weight grid, $\{x_1,...,x_{N_x}\}$, as starting values for the next period. Thus, we consider N_x different starting values of the relative Pareto weight in the next period, and the couple starts in the next period with the

[7]This initialization strategy follows Ábrahám and Laczó (2018).

relative Pareto weight x. Thus, we have $x' = x$ since there was no updating of the relative Pareto weight in the current period, i.e., the next period starting value of the relative Pareto weight is the current period relative Pareto weight. All other state variables are summarized by Ω'. Given the initial guess for the state-dependent relative Pareto weight grids, $\Theta_j(\Omega, x)$, the relative Pareto weight is updated when participation constraints are binding. Thus, the Pareto weight after the update is given by $x'' = \Theta_j(\Omega', x') = \Theta_j(\Omega', x)$. The risk-sharing condition (138) implies

$$c'_f = \left(x''\right)^{-1/\sigma} \cdot c'_m. \tag{144}$$

Using this condition, the budget constraint in the next period becomes

$$c'_m + \left(x''\right)^{-1/\sigma} \cdot c'_m + \mathcal{A}_j = y'_m + y'_f + (1+r)\, a'.$$

Thus, given the asset policy function, $\mathcal{A}_j(\Omega, x)$ and the state-dependent relative Pareto weight grids, $\Theta_j(\Omega, x)$, one can calculate next period male consumption as a function of next period state variables and the current period Pareto weight. We use a fast vectorized bisection root-finding routine to determine this policy function.

In Step 5b of Algorithm 8, one finds current period consumption and current period labor supply as a function of current period incomes, next period asset holdings and the current period relative Pareto weight. These functions can be determined using the Euler equation (137) and the initial guesses of the re-defined Kuhn-Tucker multipliers of the participation constraints $\upsilon_{m,j}(\Omega, x)$ and $\upsilon_{f,j}(\Omega, x)$. Using male marginal utility of consumption, $\partial u_m(c_m)/\partial c_m = c_m^{-\sigma}$, the Euler equation (137) can be rewritten as[8]

$$c_m^{-\sigma} = \beta \cdot \mathbb{E}\left[(1+r) \cdot \frac{1}{1-\upsilon'_m} \cdot \left(c'_m\right)^{-\sigma} - \frac{\upsilon'_m}{1-\upsilon'_m} \cdot \frac{\partial S_m}{\partial a'} - \frac{\upsilon'_f}{\left(1-\upsilon'_m\right) \cdot x''} \cdot \frac{\partial S_f}{\partial a'}\right].$$

The partial derivatives $\partial S_m/\partial a'$ and $\partial S_f/\partial a'$ can be obtained from the Envelope conditions of the bachelor household problem with

$$\begin{aligned} \frac{\partial S_m}{\partial a'} &= (1-\varphi) \cdot \kappa \cdot \left(c_m^{\text{bach},\prime}\right)^{-\sigma}, \text{ and} \\ \frac{\partial S_f}{\partial a'} &= (1-\varphi) \cdot (1-\kappa) \cdot \left(c_f^{\text{bach},\prime}\right)^{-\sigma}, \end{aligned}$$

where φ denotes the constant share of assets which is lost when a couple files for divorce, and κ denotes the share of the remaining joint asset holdings of the former husband. Thus, the solutions of the bachelor household problems are needed to capture the impact of savings

[8] A derivation of the equation can be found in Appendix A.3.

decisions on future outside options in the Euler equation of the limited-commitment problem. To find the consumption values of the bachelor households $c_m^{\text{bach},\prime}$ and $c_f^{\text{bach},\prime}$, we interpolate the corresponding policy functions on $a'_m = (1-\varphi)\cdot\kappa\cdot a'$ and $a'_f = (1-\varphi)\cdot(1-\kappa)\cdot a'$. Then, one can find current period male consumption as a function of current period incomes, next period asset holdings, and the current period relative Pareto weight using the Markovian structure of the model. Thereafter, one can use condition (144) to find current period female consumption as functions of the current period incomes, next period asset holdings, and the current period relative Pareto weight.

Step 5c of Algorithm 8 is similar to the full-commitment case, see Step 2c of Algorithm 7. One can calculate the endogenous grid $\hat{a}$ analytically by using the budget constraint (126):

$$\hat{a} = \frac{c_m + c_f + a' - y_m - y_f}{1+r}.$$

Since all variables are functions of current period incomes, current period relative Pareto weight, and next period asset holdings, we have a one-to-one mapping for each income and relative Pareto weight combination from the endogenous grid $\hat{a}$ to next period assets: $a'(\hat{a})$.

Upper envelope step. A particular challenge for the EGM-approach is that the presence of an endogenous state variable in the participation constraints may induce non-convexities. This can be problematic for two reasons. First, the first-order conditions are only sufficient for a global maximum when the target function is concave and all constraints are convex, see Stokey et al. (1989).[9] To deal with this potential issue, we have solved the model several times with different initial guesses. The results were always the same. The second issue is that non-convex constraints may induce ranges in the endogenous grid $\hat{a}$ where the one-to-one mapping $a'(\hat{a})$ is not monotone. Thus, potentially there can be grid points of the endogenous grid which are mapped to two or more different values for next period assets. Only one of these values can be optimal. To deal with this issue, we adapt the upper envelope approach of Iskhakov et al. (2017) to be compatible with our limited-commitment model.[10] Their approach can be used in all cases where non-convexities cause non-monotonic ranges of the endogenous grid as long as there is only one endogenous state variable.[11] Algorithm 9 summarizes the adapted upper envelope approach of Iskhakov et al. (2017).

Algorithm 9 (Upper envelope algorithm) *The algorithm consists of the following four*

[9] Strictly speaking, constraints must be non-empty, compact-valued, convex and continuous.

[10] Fella (2014) provides an alternative approach to deal with non-monotonic endogenous grids. Druedahl and Jørgensen (2017) generalize the approach of Iskhakov et al. (2017) to multiple endogenous state variables and non-convexities.

[11] In the original set-up considered by Iskhakov et al. (2017), non-convexities are caused by discrete choices.

steps:

1: Calculate the lifetime utility of both spouses as a function of current period incomes, the current period relative Pareto weight, and next period asset holdings. This gives for each ω-x combination a value function $\tilde{V}_{m,\omega-x}$ and $\tilde{V}_{f,\omega-x}$ as a function of the endogenous grid $\hat{a}$. These value functions imply that the household target function can be written as $\tilde{V}_{\omega-x} = \tilde{V}_{m,\omega-x} + x^{-1} \cdot \tilde{V}_{f,\omega-x}$.

2: For each ω-x combination, find non-monotonic ranges of the endogenous grid. In particular, for each ω-x combination and each $j = 1, ..., N_a - 1$ check whether $\hat{a}_j > \hat{a}_{j+1}$. If yes, perform Step 3. If no, go to Step 4.

3: Find $j' > j$ such that $\hat{a}_{j'} > \hat{a}_j$. Define partitions, $N_1 = \{1, ..., j\}$, $N_2 = \{j, ..., j'\}$, and $N_3 = \{j', N_a\}$. Interpolate the household target function from Step 1 on these partitions. Using these interpolated values of the target function, detect suboptimal values of $\hat{a}$. Delete these suboptimal values.

4: Set $\acute{a} = \hat{a}$.

After resolving the problem of non-monotone ranges in the endogenous grid, one can use the mapping $a'(\acute{a})$ to find the new asset policy functions using interpolation in Step 5e of Algorithm 8. In particular, we use linear interpolation and interpolate $a'(\acute{a})$ on the exogenous grid $a_1, ..., a_{N_a}$ to get $\mathcal{A}^{\text{uncon}}(\Omega, x)$.

We search for $\Omega - x$ combinations where the borrowing constraint is binding in Step 5f of Algorithm 8. We replace all values $\mathcal{A}^{\text{uncon}}(\Omega, x) < 0$ by 0.

In Step 5g of Algorithm 8, we use the budget constraint (126) and the risk-sharing condition (138), to calculate unconstrained male and female consumption. For the given utility function, we solve

$$c_m + x^{-1/\sigma} \cdot c_m + \mathcal{A}^{\text{uncon}} = y_m + y_f + (1 + r)\, a$$

for c_m to get $\mathcal{C}_m^{\text{uncon}}(\Omega, x)$. Thereafter, one can use (144) to calculate $\mathcal{C}_f^{\text{uncon}}(\Omega, x)$.

In Step 5h, we calculate the unconstrained value functions. As before, this reflects the basic idea of the algorithm that beginning with the next period commitment between spouses is limited. Accordingly, the continuation value on the right-hand side of the Bellman equations is the initial guess of the limited-commitment value functions.

These value functions are used in Step 6a to check the participation constraints. The lifetime utilities of the outside options are obtained by interpolating the value functions $S_m(y_m, a_m)$ and $S_f(y_f, a_f)$ on the asset values $(1 - \varphi) \cdot \kappa \cdot a$ and $(1 - \varphi) \cdot (1 - \kappa) \cdot a$. $\underline{\nu}(\Omega)$ denotes the relative Pareto weight such that the participation constraint of the husband holds with equality. $\overline{\nu}(\Omega)$ denotes the relative Pareto weight such that the participation constraint of the wife holds with equality. Thus, for any income and asset combination, there exists an optimal interval $[\underline{\nu}, \overline{\nu}]$ where both participation constraints are fulfilled at least with equality. These intervals can be used in the simulation of the model to update the relative Pareto weight.

In Step 7a of Algorithm 8, we calculate the new value functions under limited commitment. To do so, we set $V_{m,j+1}(\Omega, x) = V_m^{\text{uncon}}(\Omega, x)$ and $V_{f,j+1}(\Omega, x) = V_f^{\text{uncon}}(\Omega, x)$. Then, we replace all values which are smaller than the outside options by the values of the outside option. Note that also the values of the value function of the spouse have to be replaced. We interpolate both value functions on the new Pareto weight grid $\Theta_{j+1}(\Omega, x)$. This guarantees that both participation constraints are always fulfilled and that the lifetime utilities of spouses with non-binding participations constraints correspond to the correct Pareto weight value. These interpolated value functions become the initial guesses in the next iteration of the algorithm.

To get the new asset policy function, we interpolate the unconstrained asset policy function $\mathcal{A}^{\text{uncon}}(\Omega, x)$ on the state-dependent relative Pareto weight grid $\Theta_{j+1}(\Omega, x)$ to get $\mathcal{A}_{j+1}(\Omega, x)$ in Step 7b of Algorithm 8. The policy function $\mathcal{A}^{\text{uncon}}(\Omega, x)$ describes the optimal savings behavior of the couple knowing that beginning with the next period commitment is limited. To find the new constrained policy function, one has to consider that the relative Pareto weight may also be updated in the current period. The couple starts the period with the relative Pareto weight x. Then, the relative Pareto weight is updated such that both participation constraints are fulfilled. This is done by the policy function $\Theta_{j+1}(\Omega, x)$. To find the constrained asset policy function, one has to interpolate the unconstrained asset policy function on the relative Pareto weight updating policy function. This gives the asset policy function of a constrained household which starts with relative Pareto weight x. If the relative Pareto weight is not updated in the current period the constrained asset policy function is simply the unconstrained asset policy function from Step 5f. If the relative Pareto weight is updated in the current period the constrained asset policy function is the unconstrained asset policy function from Step 5f evaluated at the new value of the relative Pareto weight.

Finally, the new values of the redefined Kuhn-Tucker multipliers on the participation constraints $\upsilon_{m,j+1}$ and $\upsilon_{f,j+1}$ are calculated in Step 7c. Since the new values of the state-dependent Pareto weight grids $\Theta_{j+1}(\Omega, x)$ are derived from the state-independent grid $\tilde{\Theta}$, one can derive these values implicitly through

$$\Theta_{j+1}(\Omega, x) = \frac{\mu'_m(\Omega, x)}{\mu'_f(\Omega, x)} = \tilde{\Theta} \cdot \frac{1 - \upsilon_{f,j+1}(\Omega, x)}{1 - \upsilon_{m,j+1}(\Omega, x)}. \tag{145}$$

Since we abstract from divorces, only one participation constraint can be violated at the same time. Thus, to calculate $\upsilon_{m,j+1}$ one can set $\upsilon_{f,j+1} = 0$. Then, $\upsilon_{m,j+1}$ is given by

$$\upsilon_{m,j+1}(\Omega, x) = \begin{cases} 0, & \text{if } \Theta_{j+1}(\Omega, x) \leq \tilde{\Theta}, \\ 1 - \frac{\tilde{\Theta}}{\Theta_{j+1}(\Omega,x)}, & \text{if } \Theta_{j+1}(\Omega, x) > \tilde{\Theta}. \end{cases}$$

The re-defined Kuhn-Tucker multiplier on the male participation constraint is equal to 0 whenever the relative Pareto weight does not increase since the male participation constraint is not binding in this case. This can be seen from the definition of the re-defined Kuhn-Tucker multiplier:

$$\upsilon_m(\Omega, x) = \frac{\nu_m(\Omega, x)}{\mu'_m(\Omega, x)},$$

where $\nu_m(\Omega, x)$ is zero whenever the constraint is not binding. When the male participation constraint is binding and the female participation constraint is not binding, the value of $\upsilon_{m,j+1}$ follows from (145) with $\upsilon_{f,j+1} = 0$. Analogously, one can calculate $\upsilon_{f,j+1}$ by

$$\upsilon_{f,j+1}(\Omega, x) = \begin{cases} 0, & \text{if } \Theta_{j+1}(\Omega, x) \geq \tilde{\Theta}, \\ 1 - \frac{\Theta_{j+1}(\Omega,x)}{\tilde{\Theta}}, & \text{if } \Theta_{j+1}(\Omega, x) < \tilde{\Theta}, \end{cases}$$

using $\upsilon_{m,j+1} = 0$.

In Step 8 of Algorithm 8, we check the convergence of the constrained value functions $V_m(\Omega, x)$ and $V_f(\Omega, x)$, the constrained asset policy function $\mathcal{A}(\Omega, x)$, the state-dependent relative Pareto weight grids $\Theta(\Omega, x)$, and the re-defined Kuhn-Tucker multipliers $\upsilon_m(\Omega, x)$ and $\upsilon_f(\Omega, x)$. If all functions have converged, the solution to the limited-commitment model is found. Otherwise, one has to repeat Steps 5 to 7 of Algorithm 8 until convergence.

4 Quantitative results for savings, risk sharing and welfare under limited commitment

We now investigate the savings behavior of couples under limited commitment quantitatively. In Section 4.1, we introduce functional forms and the parameterization. In Section 4.2, we discuss policy functions and Section 4.3 presents simulation results.

4.1 Functional forms and parametrization

Table 26 summarizes the parameter values. Individual utility functions are given by

$$u\left(c_i\right) = \frac{c_i^{1-\sigma} - 1}{1-\sigma}, \tag{146}$$

where $1/\sigma$ measures the elasticity of intertemporal substitution of consumption. We assume $\sigma = 1$, such that (146) becomes

$$u\left(c_i\right) = \log c_i. \tag{147}$$

We set the discount factor β to 0.96, which is a standard value for a yearly calibration of an incomplete-markets model, see, for example, Huggett (1993), Aiyagari (1994), and Krusell and Smith (1998). We calibrate the interest rate r such that the bottom 40% of the asset distribution own 0.6% of total wealth in line with the results of Kuhn and Ríos-Rull (2016) for the U.S. economy.

We use an empirical target from Lise and Yamada (2019) to calibrate the extent of commitment issues. In the model, commitment issues can be strengthened or weakened by adjusting the direct utility gains from marriage, Ψ. Lise and Yamada (2019) find that the intra-household variance of the (logged) Pareto weight of the male is 0.0061.[1] As discussed in Section 2, Pareto weights do not sum up to one when one uses the Marcet and Marimon (2019) approach. When individual Pareto weights grow over time their variance is infinite. However, one can re-define relative Pareto weights such that they become compatible with the target of Lise and Yamada (2019). The *intra*temporal allocation depends only on the *relative* Pareto weight and individual Pareto weights do not matter for the *intra*temporal allocation. Therefore, we construct individual weights, ς_m and ς_f, from the relative Pareto

[1] The analysis by Lise and Yamada (2019) is for Japan. In the U.S., divorce rates are higher than in Japan, see OECD (2019), so that the model calibration is conservative in terms of the strength of commitment problems.

weight which sum up to one, $\varsigma_m + \varsigma_f = 1$.[2] In particular, the individual weight of the male is given by

$$\varsigma_m = \frac{x'}{1+x'}. \tag{148}$$

This can be used to determine the individual weights which sum up to one in each period for all simulated households. Then, we calibrate the direct utility gains from marriage, Ψ, such that the intra-household variance of ς_m is 0.0061.

Log income is given by a standard AR(1) process

$$\begin{aligned} \log y_i &= z_i, \text{ with} \qquad (149) \\ z_i' &= \rho \cdot z_i + \varepsilon_i, \end{aligned}$$

where ρ denotes the autocorrelation, and the income shock ε is normally distributed with variance σ_ε^2, $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$. The parameters of the income process are taken from Aiyagari (1994) and shown in Table 26. In comparison to (149), Ábrahám and Laczó (2018) consider a special case of an income process where total (household) income is constant. Individual incomes are i.i.d. in their analysis and can take three values with equal probability. As total income is constant, there is a perfectly negative correlation between spouses' incomes, i.e., when an individual has the high income, its spouse has the low income and vice versa. Thus, Ábrahám and Laczó (2018) consider a total of three different exogenous stochastic states in their model. In contrast, we consider a standard income process from the incomplete-markets literature that can be calibrated realistically. Because total household income is constant in Ábrahám and Laczó (2018), households have no income risk and there is hence no motive for precautionary savings under full commitment. By contrast, our model features household income risk and hence precautionary savings under full commitment which is important because we want to compare the quantitative role of the savings motive due to limited commitment to the well-studied precautionary savings motive.

In the baseline parameterization, we assume that divorcees do not lose any assets upon divorce, following Mazzocco et al. (2013) or Voena (2015). Hence, we set $\varphi = 0$. In line with Voena (2015), we assume a community property regime in the baseline model. This implies that assets are divided equally among spouses upon divorce, $\kappa = 0.5$. Of course, other predetermined shares are possible which will be investigated in Section 4.

In the numerical solution, we use 9 income grid points per spouse, 101 grid points for

[2] Note that these individual weights are not equal to the Pareto weights in our model which are unbounded. However, given the optimal savings decision, the resulting intratemporal decisions are identical.

Table 26: Parameter values

Description	Parameter	Value	Source
CRRA parameter	σ	1	set
Discount factor	β	0.96	set
Interest rate	r	0.035	calibrated
Marriage utility gain	Ψ	0.018	calibrated
Autocorrelation income shock	ρ	0.9	Aiyagari (1994)
Standard deviation income shock	σ_ε	0.0872	Aiyagari (1994)
Asset losses upon divorce	φ	0	Voena (2015)
Division of assets upon divorce	κ	0.5	Voena (2015)

Figure 13: Relative Pareto weight of the wife

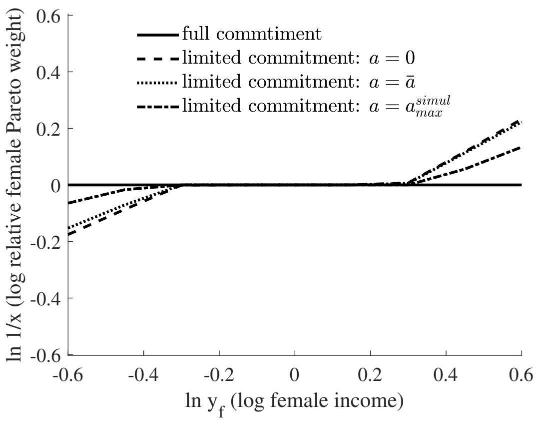

NOTE.–The figure shows policy functions for the relative Pareto weight of the wife as a function of her income. Male income is hold constant at its mean value. $\bar{a}$ denotes mean asset holdings of couples under limited commitment from the simulated model. $a_{\max}^{\text{simul}}$ denotes the maximum asset level in the simulated economy.

the Pareto weight grid, and 201 grid points for the asset grid, resulting in a total of $9 \times 9 \times 101 \times 201 = 1,644,381$ grid points. After having solved the model, we simulate a panel of $5,000$ couple households, both under full and limited commitment, using the same stochastic shocks. We also simulate $5,000$ singles of each gender. We simulate $1,500$ periods and discard the first 500 periods to avoid dependence on initial conditions. Initial Pareto weights are 0.5 since spouses are *ex-ante* identical.

4.2 Policy functions

We display policy functions as functions of the income of the wife. The husband's income is fixed at its mean value, $y_m = 1$, i.e., $\ln y_m = 0$. The relative Pareto weight of the wife from the previous period, $1/x$, is assumed to be at its long-run value 1 (there is no gender heterogeneity in the baseline parameterization). To highlight the impact of asset holdings on the strength of commitment issues, we consider three different asset levels in the figures, $a = 0$, $a = \bar{a}$, and $a = a_{\max}^{\text{simul}}$, where $\bar{a}$ denotes average asset holdings in the stationary economy and $a_{\max}^{\text{simul}}$ denotes the maximum asset level in the simulated economy.

Figure 13 shows how the relative Pareto weight of the wife depends on her own income. Under full commitment (solid line), the relative Pareto weight of the female is independent of the state of nature and household wealth. The dashed lines refers to the case of zero asset holdings under limited commitment, i.e., to a situation where the borrowing constraint tends to bind. In this case, low income realizations of the female cause a binding participation constraint of the male implying a decrease in the relative Pareto weight of the wife. In contrast, if the wife's income realization is high, her participation constraint binds such that her relative Pareto weight increases. Thus, the relative Pareto weight of the wife is weakly increasing in her income under limited commitment. The other two lines show a similar pattern in couples with higher asset holdings. Again, the relative Pareto weight of the female is weakly increasing in her income. Even in couples with very high asset holdings, see the dashed-dotted line, the relative Pareto weight of the female shows this pattern but the figure also shows that the relative Pareto weight reacts more modest when asset holdings become larger. Thus, Figure 13 suggests that commitment issues become less severe with higher asset holdings.

Figure 14 shows the savings decision of couples, for different levels of assets. Both, under full and limited commitment, household savings are increasing in female income. In all four panels, savings under limited commitment are higher than under full commitment. Even for high asset holdings, couples have an incentive to accumulate more assets than under full commitment. This reduces the likelihood of binding participation constraints in the future. However, the difference between full and limited commitment becomes smaller when households have more wealth. This is in line with the finding from Figure 13 that the strength of commitment issues decreases in asset holdings.

4.3 Simulation results

Commitment issues. The quantitative model was calibrated so that the variance of the Pareto weight of the male is 0.0061, see Section 4.1. As an untargeted statistic, we consider the rebargaining frequency as a second measure of the strength of commitment issues. The rebargaining frequency in the simulated model is 16.7%. This implies that couples rebargaing on average every 6 periods (years) over the relative Pareto weight.

Figure 14: Savings policy functions

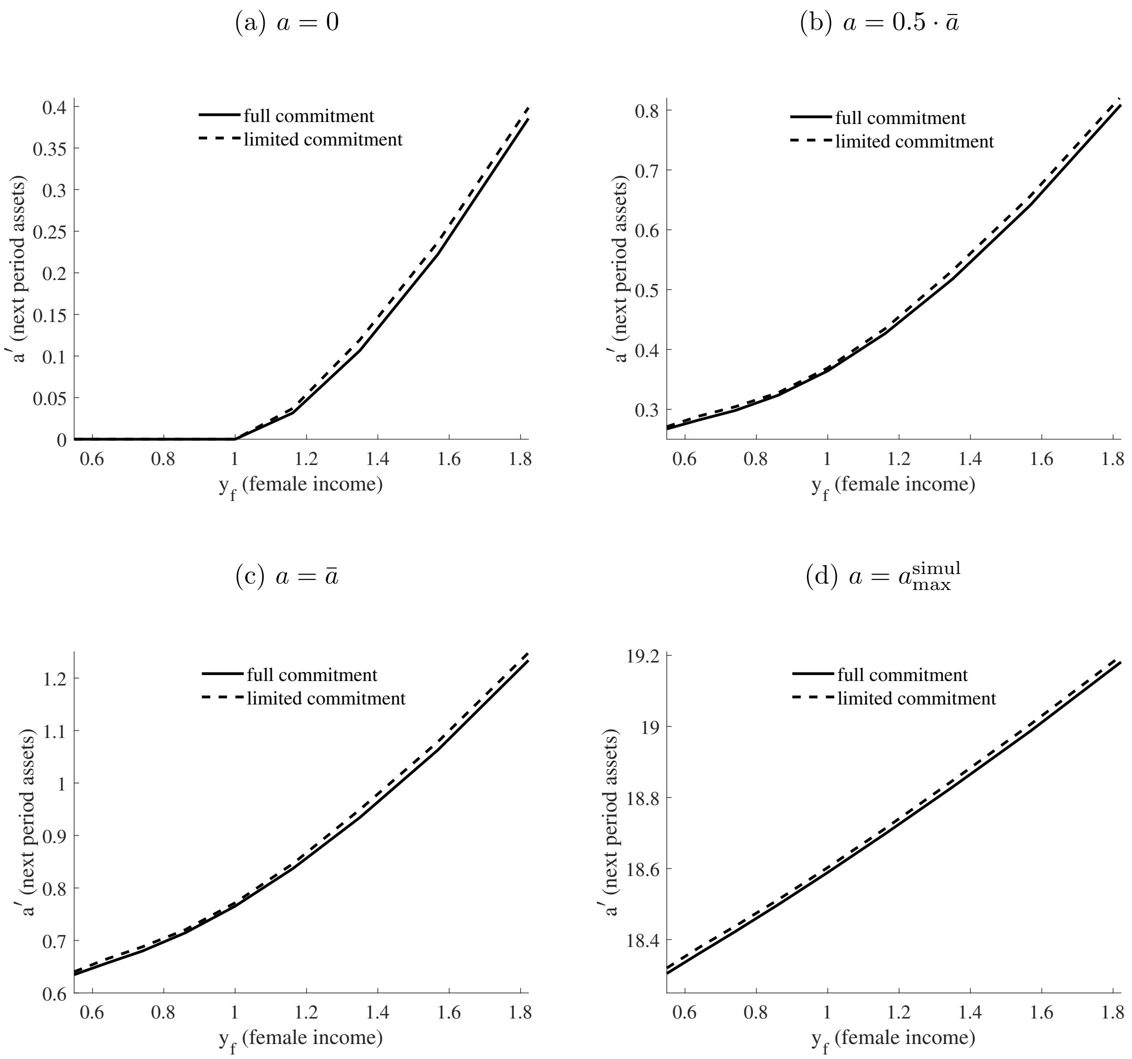

NOTE.–The figure shows savings policy functions as a function of the wife's income for different asset holdings. Income of the husband is hold constant its mean value. $\bar{a}$ denotes mean asset holdings of couples under limited commitment from the simulated model. $a^{\text{simul}}_{\text{max}}$ denotes the maximum asset level in the simulated economy.

Figure 15: Pareto weights and asset holdings

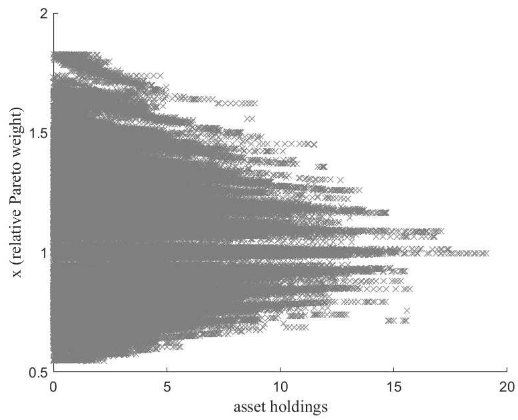

NOTE.–The figure shows a scatterplot of the relative Pareto weight of the male and asset holdings of the couple, simulated using our baseline model.

Commitment issues are smaller in wealthier households. This is illustrated in Figure 15 which shows a scatterplot of the relative Pareto weight of the male and asset holdings of the couple. The dispersion of Pareto weights becomes smaller when couples have more wealth. Wealth reduces the strength of commitment problems even when no wealth is lost to the couple upon divorce (as in our baseline calibration). The reason is that, from the perspective of the spouse who is tempted to file for divorce, the divorce is associated with a de-facto wealth loss. Remaining in the marriage, this spouse can expect to consume more than 50% of household consumption, including consumption out of household wealth, because his or her Pareto weight is high. By contrast, in case of a divorce, this spouse would have to cede 50% of household wealth to the other spouse.

Savings. Average asset holdings are 0.77 under full commitment and 0.87 under limited commitment. Thus, the savings of couples are on average about 12% higher under limited commitment than under full commitment. Figure 16 shows a histogram of the stationary distribution of asset holdings under full commitment (black bars) and under limited commitment (white bars). The shape of the stationary distribution is in both cases similar to a Pareto distribution which is usually used to specify the wealth distribution empirically. Under full commitment, the borrowing constraint is binding more often and very small asset holdings are observed more often than under limited commitment. By contrast, medium and larger values of asset holdings occur more often under limited commitment than under full commitment. Since the full commitment case is the Pareto efficient benchmark, couples under limited commitment over-save inefficiently.

We now investigate asset flows instead of stocks. For this, we perform a counterfactual

Figure 16: Stationary distribution of assets

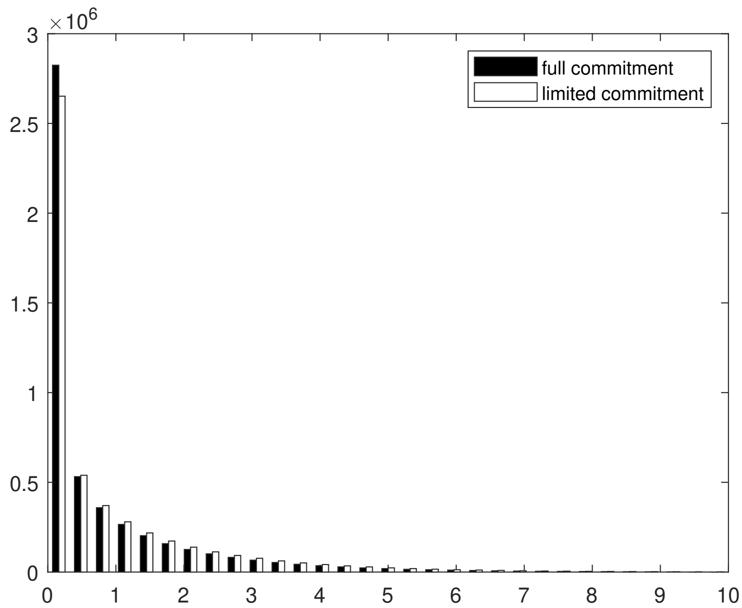

NOTE.–The figure shows the stationary distribution of asset holdings under full commitment (black bars) and under limited commitment (white bars). Maximum asset holdings in the simulation are about 20.

simulation. To discriminate between potentially different savings decisions caused by different asset holdings and different relative Pareto weights, we simulate a counterfactual where the Pareto weight is held constant at its previous period value but the asset holdings are taken from the original limited-commitment couple. This experiment allows to assess which amount couples would save when the couple could commit to the previous-period relative Pareto weight from now on, given their actual current wealth. We find that, on average, the couple under limited commitment saves 4% more than the counterfactual couple. This deviation from the Pareto efficient benchmark causes a welfare loss under limited commitment compared to a situation of full commitment. In summary, couples have on average 11% higher asset holdings (stock dimension) and save on average 4% more in every period (flow dimension) when commitment among spouses is limited.

The flow dimension helps to illustrate the efficiency loss associated with the increased saving under limited commitment. While being 11% wealthier gives couples a larger buffer stock against income risk, their utility is reduced by the 4% additional saving each year and the associated drop in yearly consumption. Under full commitment, couples efficiently balance the trade-off between consumption and wealth. The limited-commitment friction induces them to deviate from this efficient behavior and save inefficiently much and hence consume inefficiently little.

Table 27: Measures of consumption risk and risk sharing

Measure	Bachelor	Full commitment	Limited commitment
Variance of consumption growth			
Household, Ξ	–	0.27	0.25
Men, Ξ_m	0.40	0.27	0.28
Women, Ξ_f	0.40	0.27	0.28
Kaplan-Violante coefficients			
Men, ϕ_m	0.40	0.65	0.60
Women, ϕ_f	0.40	0.65	0.60

NOTE.– Ξ denotes the variance of consumption growth at the couple level and Ξ_i denotes the variance of individual consumption growth for spouse i, $\Xi_i = \text{var}(\ln c_i' - \ln c_i)$. ϕ_i is the Kaplan and Violante (2010) insurance coefficient of spouse i, $\phi_i = 1 - \text{cov}\,(\ln c_i' - \ln c_i, \varepsilon_i)/\text{var}\,(\varepsilon_i)$.

Risk sharing. Couples have two sources of insurance against idiosyncratic income shocks, intra-household risk sharing and precautionary savings. The first three rows in Table 27 show the variance of consumption growth, $\Xi_i = \text{var}(\ln c_i' - \ln c_i)$, as a measure of uninsured consumption risk. The variance of individual consumption growth is highest for bachelor households as bachelors can insure themselves only through precautionary savings. Individuals in full-commitment couple households have the lowest individual consumption growth variance. Limited commitment couples achieve more stable consumption than bachelor households but their consumption is more volatile than under full commitment. Under limited commitment, the aggregate couple consumption growth variance is not equal to the individual ones which shows that intra-household risk sharing is imperfect under limited commitment. The higher wealth which is accumulated under limited commitment leads to a lower volatility of household consumption compared to the full-commitment case. However, time-varying individual consumption shares which reflect changes in relative bargaining positions induce individual consumption to vary more strongly than under full commitment. Hence, individuals are effectively insured worse under limited commitment despite the fact that they live in wealthier households.

An alternative measure of insurance against idiosyncratic income shocks is the Kaplan and Violante (2010) insurance coefficient which measures how strongly an individual's consumption reacts to shocks to the individual's income. Formally, the Kaplan and Violante (2010) coefficient is defined as $\phi_i = 1 - \text{cov}\,(\ln c_i' - \ln c_i, \varepsilon_i)/\text{var}\,(\varepsilon_i)$. The Kaplan and Violante (2010) insurance coefficient document the same pattern as the consumption variances, see the fourth and fifth row of Table 27. The insurance coefficients are smallest for bachelor households and

Table 28: Welfare losses from limited commitment (consumption equivalents)

Group	total effect	inefficient saving	imperfect risk sharing
Male, ζ_m	0.31	0.14	0.17
Female, ζ_f	0.32	0.12	0.20
Couple, ζ	0.32	0.13	0.19

NOTE.–ζ_i denotes the welfare loss from limited commitment in terms of consumption equivalents (percentage consumption compensation). The effect due to increased savings is determined in a counterfactual analysis where limited-commitment couples are allowed to fix their Pareto weights but savings policy functions are evaluated at their actual wealth levels from the limited-commitment solution. The effect due to imperfect risk sharing is calculated as the difference between the total effect and the effect due to inefficient savings. Differences between genders are caused by numerical imprecision.

highest in the case of full commitment. Insurance coefficients for couples under limited commitment lie in-between. Quantitatively, the insurance coefficient of married individuals is 25 percentage points higher than the one of single individuals. Under limited commitment, this number is reduced to 20 percentage points. Put differently, limited commitment reduces the insurance value of marriage by one fifth.

Welfare. Table 28 shows the welfare losses from limited commitment due to imperfect risk sharing and inefficient savings, expressed as consumption equivalents. We determine welfare under limited commitment and under full commitment, respectively, in simulations of the respective model versions that start at the same distribution of wealth. Specifically, we start both simulations at the ergodic wealth distribution of the limited commitment economy. We then determine the consumption equivalent for gender i, ζ_i as the percentage increase in simulated consumption under limited commitment that would yield the same average level of lifetime utility for individuals of this gender as under full commitment.

On average, one needs to give couples about 0.32% more consumption to compensate them for the welfare losses from limited commitment. Welfare losses are equal for both genders up to numerical precision since spouses are homogeneous. To put this welfare loss due to limited commitment among spouses into perspective, the welfare losses of the business cycle are measured between 0.08%, see Imrohoroğlu (1989), and 1% of consumption, see Krusell et al. (2009). The general set-up considered in the latter study is particularly comparable to ours as their economy is also an incomplete-markets economy. Further, the welfare effects of the limited-commitment friction in household decision making seems to affect married couples' well-being to an extent that is similar to the effect of their special tax treatment. Guner et al. (2011) estimate a consumption equivalent of 0.2% for the welfare effect of abolishing opportunities for married couples to file income taxes jointly.

To disentangle the two channels through which limited commitment affects welfare, we apply a counterfactual simulation of the full-commitment model where we draw for each household a set of relative Pareto weights from the ergodic distribution of these weights in the limited-commitment model. This simulation differs from the baseline simulation of the limited-commitment model only in long-run household wealth but agents face the same risk over their individual shares in household consumption. Therefore, comparing average welfare in these two simulations measures the welfare effect of inefficient savings behavior under limited commitment. Quantitatively, the welfare effect of inefficient saving is equivalent to 0.13% of lifetime consumption. Thus, oversaving is responsible for about 40% of the welfare loss due to limited commitment.

5 Property division upon divorce

When commitment in the marriage is limited, the threat of divorce affects behavior in the marriage. Consequently, the rules for divorces impact on married couples' decisions as they determine when and for whom it can be tempting to leave the household. These rules can be set by policy makers in the form of divorce laws or by the spouses themselves in the form of, e.g., prenuptial agreements. In this section, we consider rules regarding property division upon divorce, i.e., which share of spouses' common wealth is owned by whom after a divorce.

The impact of property division upon divorce on behavior within the marriage has been studied by a number of papers in the literature. Our analysis is closely related to Ábrahám and Laczó (2016) who investigate optimal property division upon divorce in a limited-commitment framework. In their model, there is a trade-off between risk sharing within the marriage and insurance across marital states. Rules that improve the former usually worsen the latter and vice versa. We complement their analysis quantitatively by studying the effects of property division rules on risk sharing within the marriage in a less restrictive setting, most importantly we abstain from the simplifying assumption of constant household income which shuts off precautionary savings under full commitment. Rainer (2007) and Abrahám and Gottardi (2017) also study the effects of property division rules on savings incentives. Rainer (2007) does not provide quantitative results and Abrahám and Gottardi (2017) focus on a different friction in household decision making, asymmetric information about outside options. Further, there is a number of papers that study the impact of property division rules on other decisions of married couples. For example, Voena (2015) focuses on labor supply and Stevenson (2007) on education choices.

While we assumed in the baseline model that spouses do not face any costs when filing for divorce and that they share their wealth equally after divorce, we now consider alternatives to these assumptions. First, we introduce pecuniary costs of divorce in terms of wealth losses. Thereafter, we consider different regimes of the distribution of joint asset holdings after divorce. As before, we focus on couples' saving and risk-sharing behavior under limited commitment.

5.1 Resource costs upon divorce

Costs of divorce occur because the divorcees have to pay legal fees or payments to lawyers. Divorces can also be costly when divorcees have to sell illiquid assets such as housing property in a short amount of time which is usually associated with a loss. Hence, resource costs upon divorce can be set by policy makers in divorce laws or by the spouses themselves through their portfolio choices.

The literature has mostly considered one of two polar cases, either that divorcees do not lose any assets, see for example Mazzocco et al. (2013) or Voena (2015), or that divorcees lose all assets, see for example Mazzocco (2007) or Ábrahám and Laczó (2018). Ábrahám and Laczó (2016) search for the optimal asset losses and the optimal distribution of assets among divorcees but their analysis is not quantitative in nature. We use our quantitative model to investigate how different degrees of wealth losses upon divorce impact on saving and risk sharing in the marriage. For this, we solve the model for different values for the share of assets that get lost upon divorce, φ, using an equidistant grid between 0% and 100%.

Panel (a) of Figure 17 shows that the extent of commitment issues, measured by the rebargaining frequency, decreases when asset losses upon divorce increase. The mechanism is that higher asset losses upon divorce make outside options relatively less attractive which reduces commitment problems. The quantitative effects are quite substantial. While the rebargaining frequency is about 17% in the baseline parameterization where there are no ressource losses when couples file for divorce, the rebargaining frequency drops to about 11% when spouses lose all assets upon divorce.

Panel (b) of Figure 17 shows the effects on average asset holdings. Theoretically, the direction of these effects is not clear. Recall that the saving motive due to limited commitment aims at making divorce costly. When the share of assets that is lost upon divorce is higher, every unit of additional saving has a stronger effect on divorce costs. But also any given level of targeted divorce costs can be obtained by lower levels of savings. In the restricted set-up of Ábrahám and Laczó (2018), the first effect dominates and couples save more when higher shares of assets are lost upon divorce. Our quantitative results confirm that higher asset losses upon divorce induce larger average asset holdings so that the difference between average asset holdings under full and limited commitment becomes larger (the solid line shows the full commitment reference case). Quantitatively, while couples save about 20 percent more when commitment is limited and if no assets get lost upon divorce, this number almost triples when all assets are lost upon divorce. The effect is highly non-linear and already small asset losses upon divorce impact substantially on savings. For example, a 10% wealth loss upon

Figure 17: Ressource costs of filing for divorce

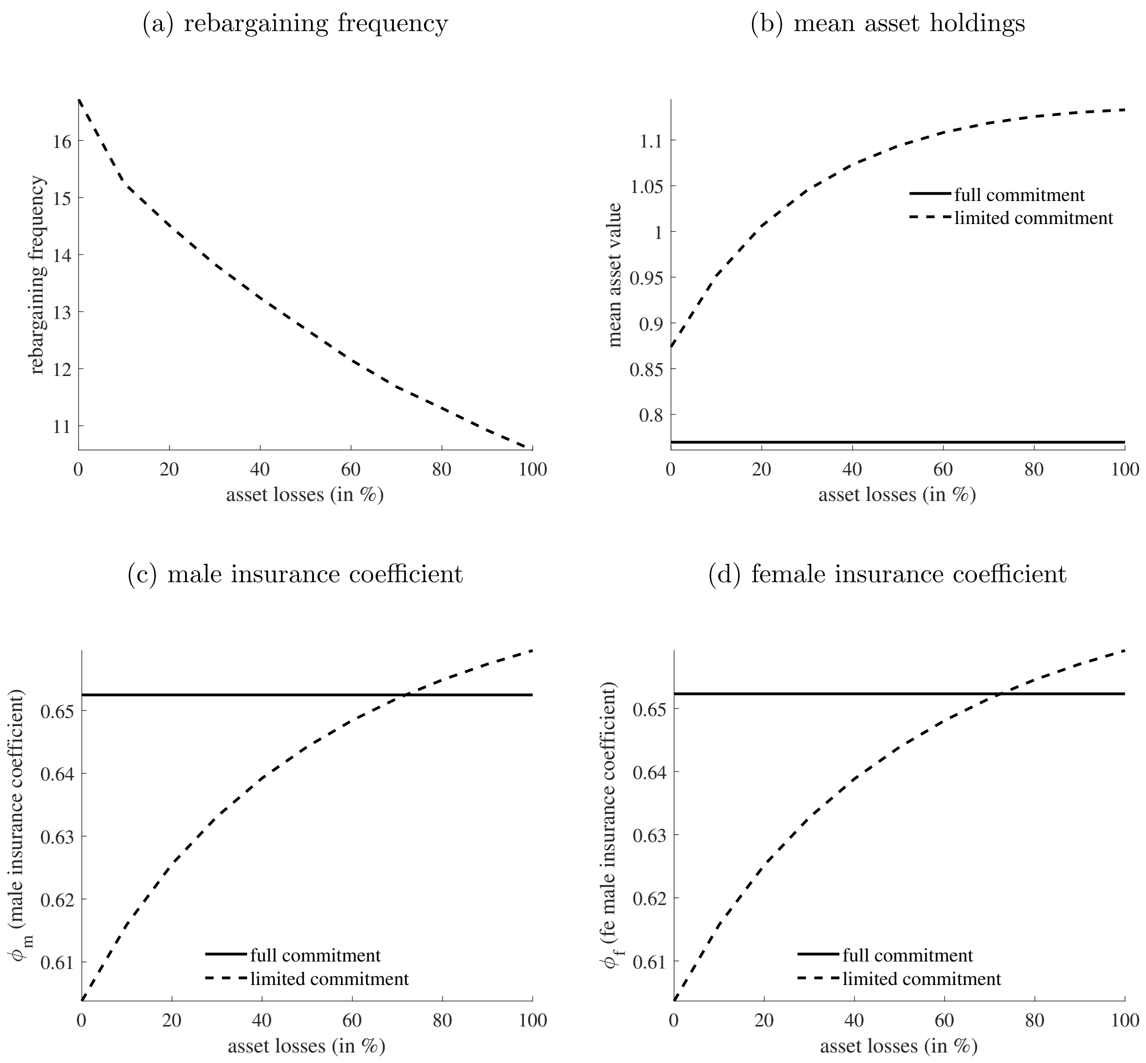

NOTE.–The figure shows the impact of resource costs of filing for divorce on the extent of commitment issues, mean asset holdings, and the Kaplan and Violante (2010) insurance coefficients. Resource costs are paid by both spouses. Solid lines indicate the case where spouses can fully commit to an allocation rule. Dashed lines indicate the case of limited commitment between spouses.

divorce leads to a 25% increase in savings due to limited commitment.

Panels (c) and (d) of Figure 17 show the Kaplan and Violante (2010) insurance coefficients.[1] Under full commitment, only about 20% of an individual's income shock translates directly to his or her consumption (solid line). Under limited commitment, insurance coefficients increase with asset losses upon divorce (dashed line). This is a consequence of the improved risk sharing in the household due to less frequent and smaller adjustments of Pareto weights and the larger cushioning of household consumption against shocks due to increased wealth.

Interestingly, the insurance coefficient under limited commitment becomes even higher than under full commitment when more than 70% of assets get lost upon divorce. In general, couples with limited commitment save more than under full commitment leading to a strong buffer stock against shocks. With relatively stark wealth losses upon divorce, their Pareto weights are adjusted rather rarely such that intra-household risk sharing is relatively strong. In combination, this leads to better insurance against income shocks than under full commitment. Nevertheless, this constitutes an inefficiently high degree of insurance against income shocks, i.e., married individuals are overinsured in this situation. Insurance through precautionary savings comes at the price of reduced consumption in order to achieve the necessary wealth level. Without the limited-commitment friction, spouses would balance the trade-off between consumption and insurance at a lower level of insurance and more consumption.

Importantly, the effect of property losses upon divorce on the insurance coefficients is concave with small asset losses exerting the largest marginal effects and large property losses having only small marginal effects. A wealth loss in the range of 20% already closes half of the gap in consumption insurance due to limited commitment. This concavity has strong implications for the trade-off between risk sharing within the marriage and insurance across marital states, where one would want to avoid large asset losses upon divorce because they induce high consumption volatility across marital states. The concavity in the benefits from asset losses in terms of improved insurance within marriage implies that from an efficiency perspective, a limited asset loss upon divorce is rather desirable while for large asset losses the marginal disadvantages probably dominate.

5.2 Distribution of assets upon divorce

We now study how the distribution of assets upon divorce affects intra-household decision making. We start with assigning partner's uneven shares of assets which are fixed before

[1]The insurance coefficient is the same for husbands and wives since there are no gender differences in the model.

the marriage. That is, in our first experiment, the wealth share which one of the spouses receives upon divorce can differ from 50% but is independent from what happens in the marriage or from who files for the divorce. The relevant application of this analysis concerns the treatment of wealth which was accumulated by one of the spouses before the formation of marriage (and the capital income obtained from this wealth) in case of divorce. A spouse can also obtain a disproportionate share of wealth accumulated within marriage, e.g., when spouses save by enlarging or maintaining a house brought into the marriage by this spouse or when they invest into an enterprise which was owned by this spouse before the marriage. This treatment can be determined in divorce laws or in prenuptial agreements of the spouses.

In the baseline parameterization, it was assumed that spouses share assets equally upon divorce, $\kappa = 0.5$. To analyze the impact of different asset shares of spouses upon divorce, we now consider asset shares between 35% and 65%.[2]

Panel (a) in Figure 18 shows that commitment problems are smallest when assets upon divorce are distributed evenly among spouses ($\kappa = 0.5$). The rebargaining frequency increases for smaller and higher asset shares of the husband.[3] Because divorce is unilateral, an uneven division of assets upon divorce make participation constraints binding more often as the partner potentially receiving the higher share of resources is more often tempted to leave the household. However, the effects are modest quantitatively. The rebargaining frequency rises by less than half a percentage points even for very asymmetric property divisions.

Panel (b) shows that average asset holdings under limited commitment are hump-shaped in the asset share of the husband upon divorce (dashed line). Couple save most when both spouses receive equal shares of assets upon divorce but, also here, the effects are small quantitatively. Even when property division is relatively asymmetric, average wealth is affected by only about one percent.

The insurance coefficients shown in panels (c) and (d) of Figure 18 become smaller when the asset share assigned to the respective individual in case of divorce increases. When a spouse receives a larger share of assets upon divorce, his or her Pareto weight in the marriage tends to increase. For example, consider the case where the wife would receive a larger share of assets upon divorce. In this case, her participation constraint tends to bind already for moderate positive shocks to her income as she can credibly threaten to leave the household and receive a large share of household wealth. On the other hand, negative shocks to the

[2] If asset shares of spouses upon divorce are too high, divorces become attractive even in the presence of (large) direct utility gains from marriage. Since such values of κ would create incentives to terminate efficient marriages, we restrict our attention to cases where this does not occur.

[3] The effects to the left and to the right of $\kappa = 0.5$ are symmetric up to numerical precision.

Figure 18: Distribution of assets upon divorce

(a) Rebargaining frequency

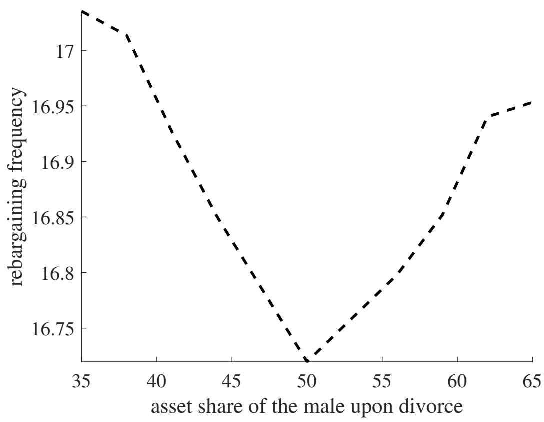

(b) mean asset holding

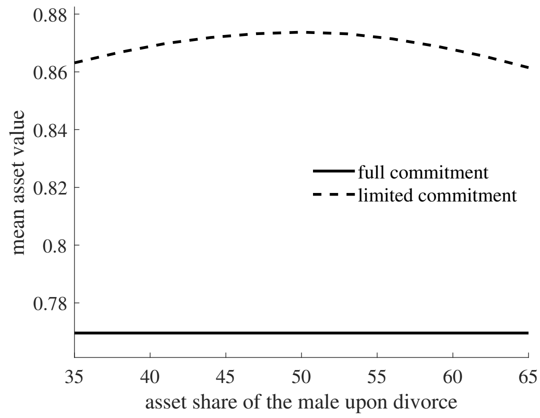

(c) male insurance coefficient

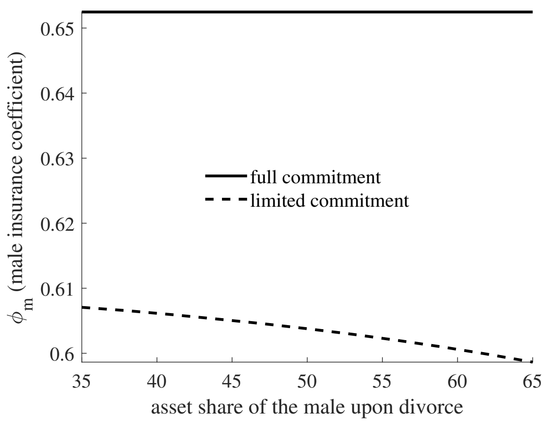

(d) female insurance coefficient

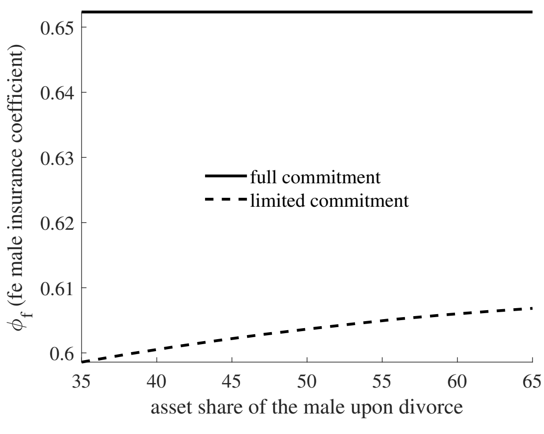

NOTE.–The figure shows the impact of the asset share distributed to the husband upon divorce on the extent of commitment issues, mean asset holdings, and the Kaplan and Violante (2010) insurance coefficients. Solid lines indicate the case where spouses can fully commit to an allocation rule. Dashed lines indicate the case of limited commitment between spouses.

Figure 19: Policy function for consumption of the wife when the wife receives a large share of wealth upon divorce.

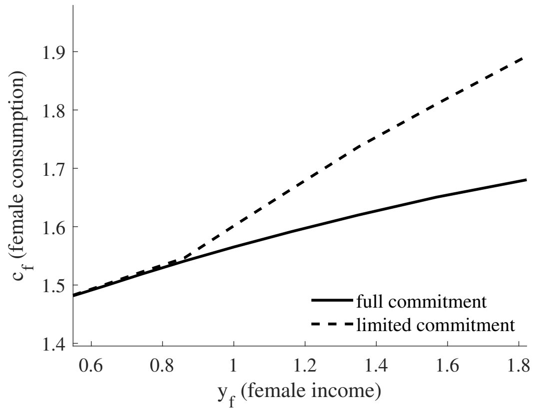

NOTE.–The figure shows consumption of the wife as a function of her income. Income of the husband is hold constant at its mean value. The wife receives the share $1 - \kappa = 0.65$ of all household assets upon divorce.

wife's income do not tend to make the husband's participation constraint binding because, for him, the outside option is relatively unattractive due to the low share of wealth he would receive. For the wife, who is treated preferentially in a divorce, this leads to a good insurance against downward income risk while she can keep the upward risk to her income mostly to herself. This is illustrated in Figure 19 which plots, for the case where the wife's share in household wealth upon divorce is 65%, the wife's consumption against her income holding constant the husband's income. The figure shows that, compared to the full-commitment benchmark, the wife's consumption responds more strongly to shocks to her own income. However, the wife is actually better off than under full commitment as her consumption is always larger under limited commitment.

While uneven property division upon divorce induce some interesting dynamics in intra-household decision making, we find that the quantitative effects on savings and risk sharing are moderate. This implies that, while uneven property division strengthen problems related to limited commitment to some extent, couples may be willing to accept this when other arguments–such as making it attractive for an individual with higher pre-marital wealth to join the marriage in the first place– are in favor of such an uneven property division rule.

5.3 Penalty for filing for divorce

We now consider the case where the asset share of the individual who files for divorce may be different from the share of his or her spouse. An example for divorce rules including such penalties for filing for divorce are alimony obligations. In our set-up, a spouse is tempted to leave the household when his or her income exceeds the income of the other spouse substantially. If alimony obligations are connected to relative incomes, this constitutes a liability for the spouse who files for divorce and this spouse's net wealth is reduced by the expected present value of the alimony liability. In turn, the other spouse's net wealth is increased by the expected present value of the alimony stream. In prenuptial agreements, spouses are free to determine property division rules that may depend on who files for divorce. Ábrahám and Laczó (2016) show that divorce rules that include such penalties can be very effective in ameliorating commitment problems within the household as they make divorce unattractive to exactly the spouse who would otherwise be tempted to file for one.

Penalties for filing for divorce affect the participation constraints as, in contrast to the baseline model, it is now important *who* files for divorce. The participation constraints are now given by

$$\begin{aligned} u_m\left(c_m\right)+\Psi+\beta\cdot\mathbb{E}\left[V_m\left(\Omega',\mu_m',\mu_f'\right)\right] &\geq S_m\left(y_m,(1-\varphi)\cdot\kappa\cdot(1-\varrho)\cdot a\right)\\ u_f\left(c_f\right)+\Psi+\beta\cdot\mathbb{E}\left[V_f\left(\Omega',\mu_m',\mu_f'\right)\right] &\geq S_f\left(y_f,(1-\varphi)\cdot(1-\kappa)\cdot(1-\varrho)\cdot a\right). \end{aligned}$$

As in the baseline model, spouses lose a constant share, φ, of all assets if they file for divorce. The husband receives the share κ and the wife receives the share $1-\kappa$ of the remaining assets. The difference to the baseline model is that the person who files for divorce faces a wealth penalty, ϱ. Thus, a when the husband files for divorce, he receives the share $(1-\varphi)\cdot\kappa\cdot(1-\varrho)$ of the household's assets while the wife faces no penalty and receives share $(1-\varphi)\cdot(1-\kappa)$. Reversely, when the wife files for divorce, she would face the penalty and receive share $(1-\varphi)\cdot(1-\kappa)\cdot(1-\varrho)$ and the husband would face no penalty and receive share $(1-\varphi)\cdot\kappa$.

Figure 20 shows the rebargaining frequency, average household wealth, and the Kaplan and Violante (2010) insurance coefficients plotted against the penalty for the spouse filing for divorce. The figure is identical to Figure 17. Under limited commitment, only the asset share of the person whose participation constraint is binding is relevant which coincides with the person who would file for divorce if his or her Pareto weight were not adjusted. If this person received an asset share of, say, 40%, he or she does not care about whether the household as a whole looses 20% of assets and share the remainder evenly or whether there is a 20% penalty for filing for divorce. Hence, the strength of commitment issues and the resulting inefficiencies within the marriage are identical in both scenarios. However, the rules have starkly different implications regarding transitions between marital states because penalties for filing for divorce do not include actual asset losses for the household as a whole. As a consequence, spouses should prefer penalties for filing for divorce over asset losses upon divorce when setting up a prenuptial agreement.

Figure 20: Penalty for filing for divorce

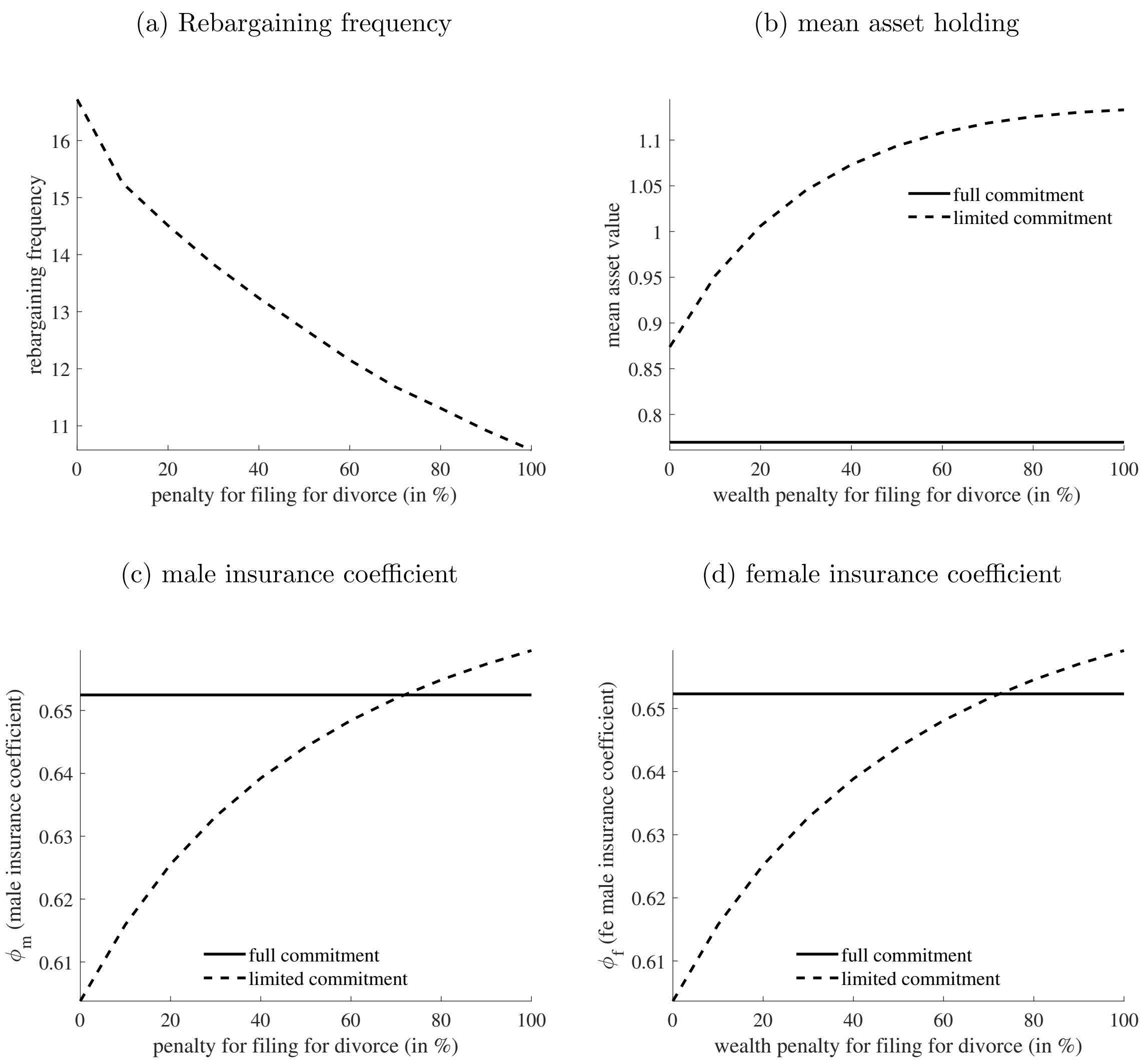

NOTE.–The figure shows the impact of individual resource costs of filing for divorce on the extent of commitment issues, mean asset holdings, and the Kaplan and Violante (2010) insurance coefficients. Resource costs are paid by the spouse who files for divorce. Thus, when the husband files for divorce, he receives the share $(1-\varphi)\cdot\kappa\cdot(1-\varrho)$ of the assets of the couple, where φ denotes total asset losses upon divorce, κ denotes the share of assets distributed to the husband, and φ denotes the penalty for filing for divorce. Analogously, the wife would receive the share $(1-\varphi)\cdot(1-\kappa)\cdot(1-\varrho)$ when she files for divorce. Solid lines refer to the case where spouses can fully commit to an allocation rule. Dashed lines refer to the case of limited commitment between spouses.

6 Conclusion

Limited commitment between household members induces them to save inefficiently much and to share inefficiently little risk. In this paper, we have introduced limited commitment in a dual-earner heterogeneous-agents incomplete-markets economy and have investigated these inefficiencies quantitatively. To do so, we have developed a practically feasible and fast numerical approach for solving intertemporal decision problems of multi-person households in presence of limited commitment. We have calibrated our model to generate realistic strengths of commitment problems.

Our results show that the effects of limited commitment on savings behavior are substantial. On average, married couples' yearly saving is increased by 4% and their long-run wealth is increased by 11%. Despite living in wealthier households, married individuals share risk less efficiently with their spouses when commitment between them is limited. Quantitatively, limited commitment reduces the insurance value of marriage by about 20% and the wellbeing of married individuals by an amount equivalent to 0.3% of lifetime consumption. About 40% of this welfare loss is due to oversaving.

We have also analyzed how property division upon divorce affects savings and risk-sharing behavior in the marriage. Wealth losses upon divorce and penalties for the spouse filing for divorce can reduce commitment problems in the marriage substantially. Losses or penalties in the range of 20% are found to close about half of the efficiency gap in consumption insurance due to limited commitment.

A Appendix

A.1 The Marcet and Marimon (2019) approach in a model with savings

In the couple household model with savings, the period budget constraint in period t is

$$c_m\left(s^t\right)+c_f\left(s^t\right)+a\left(s^t\right)\leq y_m\left(s_t\right)+y_f\left(s_t\right)+(1+r)\,a\left(s^{t-1}\right).$$

Both couple household types, full and limited commitment, face initial participation constraints

$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot\left[u_m\left(c_m\left(s^t\right)\right)+\Psi\right]\geq\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot u_m\left(c_m^{bachelor}\left(s^t\right)\right),$$
$$\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot\left[u_f\left(c_f\left(s^t\right)\right)+\Psi\right]\geq\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\middle|\,s^0\right)\cdot u_f\left(c_f^{bachelor}\left(s^t\right)\right),$$

with $\pi\left(s^0\middle|\,s^0\right)=1$. These initial participation constraints also depend on initial asset holdings and the recursive formulation is

$$V_m\left(\mathbf{y}\left(s_0\right),a_m\left(s_{-1}\right)+a_f\left(s_{-1}\right),\mu_{m0},\mu_{f0}\right)\geq S_m\left(y_m\left(s_0\right),a_m\left(s_{-1}\right)\right),$$
$$V_f\left(\mathbf{y}\left(s_0\right),a_m\left(s_{-1}\right)+a_f\left(s_{-1}\right),\mu_{m0},\mu_{f0}\right)\geq S_f\left(y_f\left(s_0\right),a_f\left(s_{-1}\right)\right).$$

Thus, initial expected lifetime utility of both spouses must be at least as high within the household than as a single.

In the limited commitment case, expected lifetime utility of both spouses must be at least as high within the couple as in the outside option in any state of nature, after any history of states and in very period of time:

$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_m\left(c_m\left(s^r\right)\right)+\Psi\right]\geq\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_m\left(c_m^{bachelor}\left(s^r\right)\right),\ \forall s^t,\ \forall t,$$
$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_f\left(c_f\left(s^r\right)\right)+\Psi\right]\geq\sum_{t=0}^{\infty}\sum_{s^r}\beta^{t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_f\left(c_f^{bachelor}\left(s^r\right)\right),\ \forall s^t,\ \forall t,$$

with $\pi\left(s^t\middle|\,s^t\right)=1$. These participation constraints also depend on asset holdings. Additionally, the rule for asset division upon divorce matters since it determines how many assets each divorcee receives. The participation constraints in recursive formulation are

$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_m\left(c_m\left(s^r\right)\right)+\Psi\right]\geq S_m\left(y_m\left(s_r\right),\left(1-\varphi\right)\cdot\kappa\cdot a\left(s^{r-1}\right)\right),$$
$$\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_f\left(c_f\left(s^r\right)\right)+\Psi\right]\geq S_f\left(y_f\left(s_r\right),\left(1-\varphi\right)\cdot\left(1-\kappa\right)\cdot a\left(s^{r-1}\right)\right).$$

Couples lose a constant share, φ, of all assets if they file for divorce. The remaining joint asset holdings are distributed with shares κ and $1-\kappa$ to the former husband and the former wife, respectively.

The Lagrangian of the household maximization problem is

$$\begin{aligned}
\mathcal{L}=\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\cdot\pi\left(s^t\right)\cdot\Big\{&\mu_0\cdot u_m\left(c_m\left(s^t\right)\right)+\left(1-\mu_0\right)\cdot u_f\left(c_f\left(s^t\right)\right)+\Psi\\
&+\nu_m\left(s^t\right)\cdot\begin{pmatrix}\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_m\left(c_m\left(s^r\right)\right)+\Psi\right]-\\ \sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_m\left(c_m^{bachelor}\left(s^r\right)\right)\end{pmatrix}\\
&+\nu_f\left(s^t\right)\cdot\begin{pmatrix}\sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot\left[u_f\left(c_f\left(s^r\right)\right)+\Psi\right]-\\ \sum_{r=t}^{\infty}\sum_{s^r}\beta^{r-t}\cdot\pi\left(s^r\middle|\,s^t\right)\cdot u_f\left(c_f^{bachelor}\left(s^r\right)\right)\end{pmatrix}\\
&+\lambda\left(s^t\right)\cdot\left(y_m\left(s_t\right)+y_f\left(s_t\right)+\left(1+r\right)a\left(s^{t-1}\right)-c_m\left(s^t\right)-c_f\left(s^t\right)-a\left(s^t\right)\right)\\
&+\xi\left(s^t\right)\cdot a\left(s^t\right)\Big\},
\end{aligned}$$

where $\nu_i\left(s^t\right)$ denotes the Kuhn-Tucker multiplier on person i's participation constraint, and $\lambda\left(s^t\right)$ denotes the Lagrangian multiplier on the budget constraint. This Lagrangian is non-stationary since it includes forward looking constraints. Marcet and Marimon (2019) provide a method that transforms this non-stationary Lagrangian to a stationary one.

The idea of Marcet and Marimon (2019) is to combine elements of the former Lagrangian approach in a way such that the maximization problem becomes stationary. The left-hand sides of the participation constraints (the inside options), which are the first term in the brackets in the second and third line of the Lagrangian, can be combined with the weighted utility of the particular spouse in the household target function. Marcet and Marimon (2019) show that the introduction of additional state variables transforms the problem to a stationary one. In particular, Marcet and Marimon (2019) show that introducing a "new" Pareto weight defined according to $\mu_i\left(s^t\right)=\mu_i\left(s^{t-1}\right)+\nu_i\left(s^t\right)$ with $\mu_m\left(s^{-1}\right)=\mu_0$ and $\mu_f\left(s^{-1}\right)=1-\mu_0$, makes the problem stationary. Thus, whenever the participation constraint of spouse i is binding, i.e., $\nu_i\left(s^t\right)>0$, his or her weight in the household target function is increased as little as possible such that the participation constraint holds with equality. Marcet and

Marimon (2019) show that the transformed maximization problem has the same solution as the original one.

Using the Marcet and Marimon (2019) approach, the Lagrangian can be rewritten as

$$\begin{aligned}
\mathcal{L} =& \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \cdot \pi\left(s^t\right) \cdot \left\{\mu_m\left(s^t\right) \cdot u_m\left(c_m\left(s^t\right)\right) + \mu_f\left(s^t\right) \cdot u_f\left(c_f\left(s^t\right)\right)\right. \\
& + \left(\mu_m\left(s^t\right) + \mu_f\left(s^t\right)\right) \cdot \Psi \\
& - \nu_m\left(s^t\right) \cdot S_m\left(y_m\left(s_t\right), (1-\varphi) \cdot \kappa \cdot a\left(s^{t-1}\right)\right) \\
& - \nu_f\left(s^t\right) \cdot S_f\left(y_f\left(s_t\right), (1-\varphi) \cdot (1-\kappa) \cdot a\left(s^{t-1}\right)\right) \\
& + \lambda\left(s^t\right) \cdot \left(y_m\left(s_t\right) + y_f\left(s_t\right) + (1+r)\, a\left(s^{t-1}\right) - c_m\left(s^t\right) - c_f\left(s^t\right) - a\left(s^t\right)\right) \\
& \left. + \xi\left(s^t\right) \cdot a\left(s^t\right)\right\}.
\end{aligned}$$

The recursive formulation of this problem is

$$\begin{aligned}
& V\left(\mathbf{y}\left(s^t\right), a\left(s^{t-1}\right), \mu_m\left(s^{t-1}\right), \mu_f\left(s^{t-1}\right)\right) = \\
& \mu_m\left(s^t\right) \cdot u_m\left(c_m\left(s^t\right)\right) + \mu_f\left(s^t\right) \cdot u_f\left(c_f\left(s^t\right)\right) \\
& \quad + \left(\mu_m\left(s^t\right) + \mu_f\left(s^t\right)\right) \cdot \Psi \\
& \quad + \beta \cdot \sum_{s^{t+1}} \pi\left(s^{t+1} \middle| s^t\right) \cdot V\left(\mathbf{y}\left(s^{t+1}\right), a\left(s^t\right), \mu_m\left(s^t\right), \mu_f\left(s^t\right)\right) \\
& \quad - \nu_m\left(s^t\right) \cdot S_m\left(y_m\left(s_t\right), (1-\varphi) \cdot \kappa \cdot a\left(s^{t-1}\right)\right) \\
& \quad - \nu_f\left(s^t\right) \cdot S_f\left(y_f\left(s_t\right), (1-\varphi) \cdot (1-\kappa) \cdot a\left(s^{t-1}\right)\right) \\
& \quad + \lambda\left(s^t\right) \cdot \left(y_m\left(s_t\right) + y_f\left(s_t\right) - c_m\left(s^t\right) - c_f\left(s^t\right)\right) \\
& \quad + \xi\left(s^t\right) \cdot a\left(s^t\right).
\end{aligned}$$

To obtain the Bellman equation, we skip the budget constraint and the borrowing constraint:

$$\begin{aligned}
& V\left(\mathbf{y}\left(s^t\right), a\left(s^{t-1}\right), \mu_m\left(s^{t-1}\right), \mu_f\left(s^{t-1}\right)\right) = \\
& \mu_m\left(s^t\right) \cdot u_m\left(c_m\left(s^t\right)\right) + \mu_f\left(s^t\right) \cdot u_f\left(c_f\left(s^t\right)\right) \\
& \quad + \left(\mu_m\left(s^t\right) + \mu_f\left(s^t\right)\right) \cdot \Psi \\
& \quad + \beta \cdot \sum_{s^{t+1}} \pi\left(s^{t+1} \middle| s^t\right) \cdot V\left(\mathbf{y}\left(s^{t+1}\right), a\left(s^t\right), \mu_m\left(s^t\right), \mu_f\left(s^t\right)\right) \\
& \quad - \nu_m\left(s^t\right) \cdot S_m\left(y_m\left(s_t\right), (1-\varphi) \cdot \kappa \cdot a\left(s^{t-1}\right)\right) \\
& \quad - \nu_f\left(s^t\right) \cdot S_f\left(y_f\left(s_t\right), (1-\varphi) \cdot (1-\kappa) \cdot a\left(s^{t-1}\right)\right),
\end{aligned}$$

and, without the state of nature notation, we obtain equation (132) from the main text:[1]

$$V\left(\mathbf{y}, a, \mu_m, \mu_f\right) = \mu_m' \cdot u_m\left(c_m\right) + \mu_f' \cdot u_f\left(c_f\right) + \left(\mu_m' + \mu_f'\right) \cdot \Psi + \beta \cdot \mathbb{E}\left[V\left(\mathbf{y}, a', \mu_m', \mu_f'\right)\right] -$$
$$\nu_m \cdot S_m\left(y_m, (1-\varphi) \cdot \kappa \cdot a\right) - \nu_f \cdot S_f\left(y_f, (1-\varphi) \cdot (1-\kappa) \cdot a\right),$$

[1] $\mu_m\left(s^t\right) = \mu_m'$ is notational convention.

A.2 Derivation of first-order conditions

A.2.1 Full commitment couple

The full commitment couple household maximization problem is

$$V\left(\Omega^{fc},\mu_0^{fc}\right)=\max_{c_m,c_f,a'}\ v\left(c_m,c_f,\mu_0^{fc}\right)+\beta\mathbb{E}\left[V\left(\Omega^{fc\prime},\mu_0^{fc}\right)\right],$$

subject to the period budget constraint

$$c_m+c_f+a'\leq y_m+y_f+(1+r)\,a,$$

and the joint borrowing constraint

$$a'\geq 0,$$

given household preferences

$$v\left(c_m,c_f,\mu_0^{fc}\right)=\mu_0^{fc}\cdot u_m\left(c_m\right)+\left(1-\mu_0^{fc}\right)\cdot u_f\left(c_f\right)+\Psi.$$

For deriving the first-order conditions, we introduce the time index t. Then, the lifetime utility of the household is given by

$$V\left(\Omega^{fc},\mu_0^{fc}\right)=\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\cdot v\left(c_{mt},c_{ft},\mu_0^{fc}\right).$$

The Lagrangian of the household maximization problem under full commitment is

$$\mathcal{L}=\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left[v\left(c_{mt},c_{ft},\mu_0^{fc}\right)+\lambda_t\cdot\left(y_{mt}+y_{ft}+(1+r)\,a_t-c_{mt}-c_{ft}-a_{t+1}\right)+\xi_t a_{t+1}\right].$$

The first-order conditions are

$$\begin{aligned}
\frac{\partial\mathcal{L}}{\partial c_{mt}}&=\beta^t\left(\frac{\partial v\left(c_{mt},c_{ft},\mu_0^{fc}\right)}{\partial c_{mt}}-\lambda_t\right)=0,\\
\frac{\partial\mathcal{L}}{\partial c_{ft}}&=\beta^t\left(\frac{\partial v\left(c_{mt},c_{ft},\mu_0^{fc}\right)}{\partial c_{ft}}-\lambda_t\right)=0,\\
\frac{\partial\mathcal{L}}{\partial a_{t+1}}&=-\beta^t\lambda_t+\beta^{t+1}\left(1+r\right)\mathbb{E}_t\lambda_{t+1}+\beta^t\xi_t=0,\\
\frac{\partial\mathcal{L}}{\partial\lambda_t}&=y_{mt}+y_{ft}+(1+r)\,a_t-c_{mt}-c_{ft}-a_{t+1}=0,\\
\xi_t&\geq 0,\ a_{t+1}\geq 0,\ \xi_t a_{t+1}=0.
\end{aligned}$$

The first-oder conditions can be re-arranged

$$\frac{\partial v\left(c_{mt}, c_{ft}, \mu_0^{fc}\right)}{\partial c_{mt}} = \lambda_t,$$

$$\frac{\partial v\left(c_{mt}, c_{ft}, \mu_0^{fc}\right)}{\partial c_{ft}} = \lambda_t,$$

$$\lambda_t = \beta\left(1+r\right)\mathbb{E}_t\lambda_{t+1} + \xi_t,$$

$$y_{mt} + y_{ft} + \left(1+r\right)a_t = c_{mt} + c_{ft} + a_{t+1},$$

$$\xi_t \geq 0,\ a_{t+1} \geq 0,\ \xi_t a_{t+1} = 0.$$

The first two conditions imply the risk-sharing condition

$$\lambda_t = \frac{\partial v\left(c_{mt}, c_{ft}, \mu_0^{fc}\right)}{\partial c_{mt}} = \mu_0^{fc}\cdot\frac{\partial u_m\left(c_{mt}\right)}{\partial c_{mt}} = \frac{\partial v\left(c_{mt}, c_{ft}, \mu_0^{fc}\right)}{\partial c_{ft}} = \left(1-\mu_0^{fc}\right)\cdot\frac{\partial u_f\left(c_{ft}\right)}{\partial c_{ft}}.$$

Thus, the first-order conditions are

$$\lambda_t = \mu_0^{fc}\cdot\frac{\partial u_m\left(c_{mt}\right)}{\partial c_{mt}} = \left(1-\mu_0^{fc}\right)\cdot\frac{\partial u_f\left(c_{ft}\right)}{\partial c_{ft}},$$

$$\lambda_t = \beta\left(1+r\right)\mathbb{E}_t\lambda_{t+1} + \xi_t,$$

$$y_{mt} + y_{ft} + \left(1+r\right)a_t = c_{mt} + c_{ft} + a_{t+1},$$

$$\xi_t \geq 0,\ a_{t+1} \geq 0,\ \xi_t a_{t+1} = 0,$$

or, in standard dynamic programming notation

$$\lambda = \mu_0^{fc}\cdot\frac{\partial u_m\left(c_m\right)}{\partial c_m} = \left(1-\mu_0^{fc}\right)\cdot\frac{\partial u_f\left(c_f\right)}{\partial c_f},$$

$$\lambda - \xi = \beta\left(1+r\right)\mathbb{E}\lambda',$$

$$y_m + y_f + \left(1+r\right)a = c_m + c_f + a',$$

$$\xi \geq 0,\ a' \geq 0,\ \xi a' = 0.$$

A.2.2 Limited commitment couple

The limited commitment couple household maximization problem in recursive formulation is

$$V\left(\Omega, \mu_m, \mu_f\right) = \max_{c_m, c_f, a} v\left(c_m, c_f, \mu_m', \mu_f'\right) - \nu_m\cdot S_m\left(y_m, \left(1-\varphi\right)\cdot\kappa\cdot a\right) -$$
$$\nu_f\cdot S_f\left(y_f, \left(1-\varphi\right)\cdot\left(1-\kappa\right)\cdot a\right) + \beta\cdot\mathbb{E}\left[V\left(\Omega', \mu_m', \mu_f'\right)\right],$$

subject to the period budget constraints

$$c_m + c_f + a' \leq y_m + y_f + (1 + r)\, a,$$

and the joint borrowing constraint

$$a' \geq 0,$$

given the household preferences

$$v\left(c_m, c_f, \mu'_m, \mu'_f\right) = \mu'_m \cdot u_m\left(c_m\right) + \mu'_f \cdot u_f\left(c_f\right) + \left(\mu'_m + \mu'_f\right) \cdot \Psi.$$

For deriving the first-order conditions, we introduce the time index t. The current state of nature is denoted by s_t and implies an income pair $\mathbf{y}\left(s_t\right) = \left(y_m\left(s_t\right), y_f\left(s_t\right)\right)$. The unconditional probability of a particular state is given by $\pi\left(s_t\right)$. The history of events until time t is $s^t = \left(s_0, ..., s_t\right)$. Thus, the unconditional probability of a particular history of states is $\pi\left(s^t\right)$. The conditional probability of a particular history of states in period r, given the history of states in period t is $\pi\left(s^r \middle| s^t\right)$ with $\pi\left(s^t \middle| s^t\right) = 1$. The Lagrangian of the maximization problem is given by

$$\begin{aligned}
\mathcal{L} = & \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \cdot \pi\left(s^t\right) \cdot \left\{\mu_m\left(s^t\right) \cdot u_m\left(c_m\left(s^t\right)\right) + \mu_f\left(s^t\right) \cdot u_f\left(c_f\left(s^t\right)\right) + \Psi\right. \\
& - \nu_m\left(s^t\right) \cdot S_m\left(y_m\left(s_t\right), (1-\varphi) \cdot \kappa \cdot a\left(s^{t-1}\right)\right) \\
& - \nu_f\left(s^t\right) \cdot S_f\left(y_f\left(s_t\right), (1-\varphi) \cdot (1-\kappa) \cdot a\left(s^{t-1}\right)\right) \\
& + \lambda\left(s^t\right) \cdot \left(y_m\left(s_t\right) + y_f\left(s_t\right) + (1+r)\, a\left(s^{t-1}\right) - c_m\left(s^t\right) - c_f\left(s^t\right) - a\left(s^t\right)\right) \\
& \left. + \xi\left(s^t\right) \cdot a\left(s^t\right)\right\},
\end{aligned}$$

where $\mu_i\left(s^t\right) = \mu'_i$ is notational convention.

The resulting first-order conditions are

$$\frac{\partial\mathcal{L}}{\partial c_m\left(s^t\right)}=\mu_m\left(s^t\right)\cdot\frac{\partial u_m\left(c_m\left(s^t\right)\right)}{\partial c_m\left(s^t\right)}-\lambda\left(s^t\right)=0,$$

$$\frac{\partial\mathcal{L}}{\partial c_f\left(s^t\right)}=\mu_f\left(s^t\right)\cdot\frac{\partial u_f\left(c_f\left(s^t\right)\right)}{\partial c_f\left(s^t\right)}-\lambda\left(s^t\right)=0,$$

$$\frac{\partial\mathcal{L}}{\partial a\left(s^t\right)}=-\pi\left(s^t\right)\cdot\left(\lambda\left(s^t\right)-\xi\left(s^t\right)\right)+\beta\cdot\sum_{s^{t+1}}\pi\left(s^{t+1}\right)\left[\begin{array}{c}(1+r)\cdot\lambda\left(s^{t+1}\right)\\-\nu_m\left(s^{t+1}\right)\cdot\\\frac{\partial S_m\left(y_m(s_{t+1}),(1-\varphi)\cdot\kappa\cdot a\left(s^t\right)\right)}{\partial a(s^t)}\\-\nu_f\left(s^{t+1}\right)\cdot\\\frac{\partial S_f\left(y_f(s_{t+1}),(1-\varphi)\cdot(1-\kappa)\cdot a\left(s^t\right)\right)}{\partial a(s^t)}\end{array}\right]=0,$$

$$\xi\left(s^t\right)\geq 0\ ,\ a\left(s^t\right)\geq 0\ ,\ \text{and}\ \xi\left(s^t\right)\cdot a\left(s^t\right)\geq 0.$$

Using $\pi\left(s^{t+1}\right)/\pi\left(s^t\right)=\pi\left(s^{t+1}\middle|\,s^t\right)$, the first order-conditions are

$$\mu_m\left(s^t\right)\cdot\frac{\partial u_m\left(c_m\left(s^t\right)\right)}{\partial c_m\left(s^t\right)}=\lambda\left(s^t\right),$$

$$\mu_f\left(s^t\right)\cdot\frac{\partial u_f\left(c_f\left(s^t\right)\right)}{\partial c_f\left(s^t\right)}=\lambda\left(s^t\right),$$

$$\lambda\left(s^t\right)-\xi\left(s^t\right)=\beta\cdot\sum_{s^{t+1}}\pi\left(s^{t+1}\middle|\,s^t\right)\left[\begin{array}{c}(1+r)\cdot\lambda\left(s^{t+1}\right)-\nu_m\left(s^{t+1}\right)\cdot\\\frac{\partial S_m\left(y_m(s_{t+1}),(1-\varphi)\cdot\kappa\cdot a\left(s^t\right)\right)}{\partial a(s^t)}-\nu_f\left(s^{t+1}\right)\cdot\\\frac{\partial S_f\left(y_f(s_{t+1}),(1-\varphi)\cdot(1-\kappa)\cdot a\left(s^t\right)\right)}{\partial a(s^t)}\end{array}\right],$$

$$\xi\left(s^t\right)\geq 0\ ,\ a\left(s^t\right)\geq 0\ ,\ \text{and}\ \xi\left(s^t\right)\cdot a\left(s^t\right)\geq 0.$$

Skipping the state of nature notation, the first-order conditions are

$$\lambda=\mu_m'\cdot\frac{\partial u_m\left(c_m\right)}{\partial c_m}=\mu_f'\cdot\frac{\partial u_f\left(c_f\right)}{\partial c_f},$$

$$\lambda-\xi=\beta\mathbb{E}\left[(1+r)\cdot\lambda'-\nu_m'\cdot\frac{\partial S_m(y_m',(1-\varphi)\cdot\kappa\cdot a')}{\partial a'}-\nu_f'\cdot\frac{\partial S_f\left(y_f',(1-\varphi)\cdot(1-\kappa)\cdot a'\right)}{\partial a'}\right],$$

$$\xi\geq 0\ ,\ a'\geq 0\ ,\ \text{and}\ \xi\cdot a'\geq 0,$$

which are the first-order conditions of the limited commitment model in Section 2 with $\mu_i'=\mu_i\left(s^t\right)$.

A.3 Numerical solution bachelor household

Solving the limited commitment model requires in Step 2 of Algorithm 8 the solution of the bachelor household model. The solution is needed for checking the participation constraints in Step 6. Algorithm 10 summarizes the numerical solution of the bachelor household maximization problem.

Algorithm 10 (Bachelor households) *The algorithm consists of the following four steps:*

1: Set convergence criteria ζ_1 and ζ_2. Discretize the exogenous wage rate processes using the approach of Tauchen (1986). Build a grid for current period assets a, $\{a_1, ..., a_{N_a}\}$. Build a grid for next period assets a', $\{\tilde{a}_1, ..., \tilde{a}_{N_a}\}$.

2: EGM step.

2a: Choose an initial guess for the asset policy functions $\mathcal{A}^S_{m,0}(\Omega_m)$ and $\mathcal{A}^S_{f,0}(\Omega_f)$.

2b: Given next period assets a' by the exogenous grid $\{\tilde{a}_1, ..., \tilde{a}_{N_a}\}$ and the guess for the asset policy function $\mathcal{A}^S_{i,j}(\Omega_i)$, find next period consumption of bachelor i as a function of next period income and next period assets period using the budget constraint (119).

2c: Use the Euler equation (122) to calculate current period consumption as a function of the current period income and next period assets.

2d: Calculate the endogenous grid of current period assets $\hat{a}_i$ using the current period budget constraint (119). This yields a mapping from the endogenous grid to the exogenous grid of a': $a'(\hat{a}_i)$.

2e: Use $a'(\hat{a}_i)$ to interpolate a' on the exogenous grid $\{a_1, ..., a_{N_a}\}$ to find the asset policy function $\mathcal{A}^S_{i,j+1}(\Omega_i)$.

2f: Check whether the borrowing constraint (120) is binding. Replace all values of $\mathcal{A}^S_{i,j+1}(\Omega_i)$ smaller than the borrowing limit by the borrowing limit.

2g: If sup $\left|\mathcal{A}^S_{m,j+1}(\Omega_m) - \mathcal{A}^S_{m,j}(\Omega_m)\right| < \zeta_1$, and sup $\left|\mathcal{A}^S_{f,j+1}(\Omega_f) - \mathcal{A}^S_{f,j}(\Omega_f)\right| < \zeta_1$ stop, set $\mathcal{A}^S_m(\Omega_m) = \mathcal{A}^S_{m,j+1}(\Omega_m)$ and $\mathcal{A}^S_f(\Omega_f) = \mathcal{A}^S_{f,j+1}(\Omega_f)$. Otherwise set $j = j+1$ and go back to step 2b. Iterate until convergence.

3: Find policy functions for consumption $\mathcal{C}^S_i(\Omega_i)$, $i \in \{m, f\}$, using the period budget constraint (119).

4: Choose initial guesses for the value functions $V^S_{m,0}(\Omega_m)$, and $V^S_{f,0}(\Omega_f)$.

4a: Given that you are in iteration step j, use the Bellman equation (118) to get individual value functions:

$$V^S_{m,j+1}(\Omega_m) = u_m\left(\mathcal{C}^S_m(\Omega_m)\right) + \beta \cdot \mathbb{E}\left[V^S_{m,j}(\Omega'_m)\right], \text{ and}$$
$$V^S_{f,j+1}(\Omega_f) = u_f\left(\mathcal{C}^S_f(\Omega_f)\right) + \beta \cdot \mathbb{E}\left[V^S_{f,j}(\Omega'_f)\right].$$

4b: If sup $\left|V^S_{m,j+1}(\Omega_m) - V^S_{m,j}(\Omega_m)\right| < \zeta_2$, and sup $\left|V^S_{f,j+1}(\Omega_f) - V^S_{f,j}(\Omega_f)\right| < \zeta_2$ stop. Otherwise set $j = j+1$ and go back to Step 4a. Iterate until convergence.

A.4 Deriving the re-arranged Euler equation under limited commitment

Using male marginal utility of consumption, $\partial u_m(c_m)/\partial c_m = c_m^{-\sigma}$, the Euler equation (137) can be written as

$$\mu_m' \cdot c_m^{-\sigma} = \beta \cdot \mathbb{E}\left[(1+r) \cdot \mu_m'' \cdot \left(c_m'\right)^{-\sigma} - \nu_m' \cdot \frac{\partial S_m}{\partial a'} - \nu_f' \cdot \frac{\partial S_f}{\partial a'}\right].$$

Dividing by μ_m' and using $\mu_m'' = \mu_m' + \nu_m'$ yields

$$c_m^{-\sigma} = \beta \cdot \mathbb{E}\left[(1+r) \cdot \frac{\mu_m''}{\mu_m'' - \nu_m'} \cdot \left(c_m'\right)^{-\sigma} - \frac{\nu_m'}{\mu_m'} \cdot \frac{\partial S_m}{\partial a'} - \frac{\nu_f'}{\mu_m'} \cdot \frac{\partial S_f}{\partial a'}\right].$$

Using $\upsilon_m' = \nu_m'/\mu_m''$, the equation becomes

$$c_m^{-\sigma} = \beta \cdot \mathbb{E}\left[(1+r) \cdot \frac{1}{1-\upsilon_m'} \cdot \left(c_m'\right)^{-\sigma} - \frac{\upsilon_m'}{1-\upsilon_m'} \cdot \frac{\partial S_m}{\partial a'} - \frac{\nu_f'}{\mu_m'} \cdot \frac{\partial S_f}{\partial a'}\right].$$

Finally, expansion of the fraction in the last term on the right-hand side yields

$$\begin{aligned}
c_m^{-\sigma} &= \beta \cdot \mathbb{E}\left[(1+r) \cdot \frac{1}{1-\upsilon_m'} \cdot \left(c_m'\right)^{-\sigma} - \frac{\upsilon_m'}{1-\upsilon_m'} \cdot \frac{\partial S_m}{\partial a'} - \frac{\nu_f'}{\mu_m'} \cdot \frac{\partial S_f}{\partial a'}\right] \\
&= \beta \cdot \mathbb{E}\left[(1+r) \cdot \frac{1}{1-\upsilon_m'} \cdot \left(c_m'\right)^{-\sigma} - \frac{\upsilon_m'}{1-\upsilon_m'} \cdot \frac{\partial S_m}{\partial a'} - \frac{\nu_f'}{\mu_f''} \cdot \frac{\mu_f''}{\mu_m''} \cdot \frac{\mu_m''}{\mu_m'} \cdot \frac{\partial S_f}{\partial a'}\right] \\
&= \beta \cdot \mathbb{E}\left[(1+r) \cdot \frac{1}{1-\upsilon_m'} \cdot \left(c_m'\right)^{-\sigma} - \frac{\upsilon_m'}{1-\upsilon_m'} \cdot \frac{\partial S_m}{\partial a'} - \frac{\upsilon_f'}{(1-\upsilon_m') \cdot x''} \cdot \frac{\partial S_f}{\partial a'}\right].
\end{aligned}$$

Chapter VI

Summary

In this Thesis, I have analyzed intra-household behavior and the resulting economic outcomes in quantitative models of the family. The particular focus was the estimation of Frisch labor-supply elasticities in couple households and inefficiencies in economic outcomes in presence of limited commitment between spouses.

Conventional estimates of Frisch labor-supply elasticities are biased in presence of borrowing constraints. In **Chapter II** (joint work with Christian Bredemeier and Falko Jüßen), we have developed an incomplete-markets model with two-earner households and full commitment. We have derived a new estimation approach for the Frisch elasticity that yields unbiased estimates even in samples that include borrowing-constrained households. Our approach exploits that the strength of the estimation bias depends on individuals' relative contribution to household earnings. It takes the form of a simple interaction-term model with minimum data requirements. Using PSID data, we have estimated Frisch elasticities of about 0.7 for men and rather homogeneous Frisch elasticities across the population.

In **Chapter III**, we have developed and calibrated an incomplete-markets model with limited commitment in dual-earner households. We have used the model to quantify how strongly empirically reasonable degrees of limited commitment reduce individuals' possibilities to share risk within the household. We also have identified for which types of agents intra-household risk sharing is particularly aggravated by limited commitment. Our results show that limited commitment reduces the insurance value of marriage by one fourth to one third. This is associated with a welfare loss equivalent to 0.6% of lifetime consumption. In the cross-section, risk-sharing possibilities are particularly reduced for individuals whose patience or risk aversion is below average while individuals with strong gains from marriage or high degrees of altruism are less affected. Our results imply an increased demand for other forms of insurance such as public insurance and provide a rationale for means-testing that targets individuals rather than households.

In **Chapter IV** (joint work with Christian Bredemeier and Falko Jüßen), we have investigated how Frisch labor-supply elasticities can be estimated when commitment between spouses is limited. In principle, the Frisch elasticity can be estimated in a regression of hours worked on the hourly wage rate when controlling for consumption. Theory suggests to control for consumption of the individual worker but most household panel surveys contain consumption information only at the household level (if at all). We have shown that proxying individual consumption by household consumption causes substantial biases in estimated Frisch elasticities when commitment between household members is limited. We have developed two improved estimation approaches that reduce or fully eliminate this bias. Applying these approaches to household panel data from the U.S., we have found Frisch elasticities of about 0.7.

In **Chapter V** (joint work with Christian Bredemeier and Falko Jüßen), we have extended our analysis of limited commitment by introducing intertemporal decisions in the form of savings. We have introduced limited commitment into a quantitative dual-earner heterogeneous-agents incomplete-markets economy and provide a practically feasible and fast numerical approach for solving such a model. We use this framework to study the inefficiencies in savings and risk-sharing behavior due to limited commitment. The effects of limited commitment on savings are substantial, raising married couples' yearly saving by 4% and their long-run wealth by 11%. The insurance value of marriage is reduced by 20%. Property division rules upon divorce can mitigate these inefficiencies. For example, a 20% wealth loss upon divorce or a 20% wealth penalty for the spouse filing for divorce close about 50% of the insurance gap due to limited commitment.

Bibliography

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.

Abrahám, A. and P. Gottardi (2017). Optimal asset division rules for dissolving partnerships.

Ábrahám, Á. and S. Laczó (2016). Household behavior and optimal property division upon divorce. Working paper.

Ábrahám, Á. and S. Laczó (2018). Efficient risk sharing with limited commitment and storage. *The Review of Economic Studies 85*(3), 1389–1424.

Acemoglu, D., M. Golosov, and A. Tsyvinski (2011). Power fluctuations and political economy. *Journal of Economic Theory 146*(3), 1009–1041.

Addoum, J. M., H. Kung, and G. Morales (2016). Limited Marital Commitment and Household Portfolios. Working Paper.

Aguiar, M. and M. Amador (2011). Growth in the shadow of expropriation. *The Quarterly Journal of Economics 126*(2), 651–697.

Aguiar, M. and E. Hurst (2007). Measuring trends in leisure: the allocation of time over five decades. *The Quarterly Journal of Economics 122*(3), 969–1006.

Aguiar, M. and E. Hurst (2016). The macroeconomics of time allocation. In *Handbook of Macroeconomics*, Volume 2, pp. 203–253. Elsevier.

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

Alesina, A., A. Ichino, and L. Karabarbounis (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy 3*(2), 1–40.

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

Altonji, J. G. (1986). Intertemporal substitution in labor supply: Evidence from micro data. *Journal of Political Economy 94*(3, Part 2), S176–S215.

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

Anderson, S. and J.-M. Baland (2002). The economics of roscas and intrahousehold resource allocation. *The Quarterly Journal of Economics 117*(3), 963–995.

Andreski, P., G. Li, M. Z. Samancioglu, and R. Schoeni (2014). Estimates of annual consumption expenditures and its major components in the psid in comparison to the ce. *American Economic Review 104*(5), 132–35.

Angelini, V., M. Bertoni, L. Stella, and C. T. Weiss (2019). The ant or the grasshopper? the long-term consequences of unilateral divorce laws on savings of european households. *European Economic Review 119*, 97–113.

Attanasio, O. and L. Pistaferri (2014). Consumption inequality over the last half century: some evidence using the new psid consumption measure. *American Economic Review 104*(5), 122–26.

Attanasio, O. and J.-V. Rıos-Rull (2000). Consumption smoothing in island economies: Can public insurance reduce welfare? *European Economic Review 44*(7), 1225–1258.

Barsky, R. B., F. T. Juster, M. S. Kimball, and M. D. Shapiro (1997). Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. *The Quarterly Journal of Economics 112*(2), 537–579.

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

Bayot, D. and A. Voena (2015). Prenuptial contracts, labor supply and household investments. Working paper.

Becker, G. S. (1960). An economic analysis of fertility. In *Demographic and economic change in developed countries*, pp. 209–240. Columbia University Press.

Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal 75*(299), 493–517.

Becker, G. S. (1974a). A theory of marriage: Part II. *Journal of Political Economy 82*(2), S11–S26.

Becker, G. S. (1974b). A theory of social interactions. *Journal of Political Economy 82*(6), 1063–1093.

Becker, G. S. (1991). *A Treatise on the Family* (Revised edition ed.). Harvard University Press.

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

Blau, D. M. and R. M. Goodstein (2016). Commitment in the household: Evidence from the effect of inheritances on the labor supply of older married couples. *Labour Economics 42*, 123–137.

Blau, F. D. and L. M. Kahn (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature 55*(3), 789–865.

Blomquist, N. S. (1985). Labour supply in a two-period model: The effect of a nonlinear progressive income tax. *The Review of Economic Studies 52*(3), 515–524.

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

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

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

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy 126*(S1), S73–S115.

Blundell, R. and I. Walker (1986). A life-cycle consistent empirical model of family labour supply using cross-section data. *The Review of Economic Studies 53*(4), 539–558.

Borjas, G. J. (1980). The relationship between wages and weekly hours of work: The role of division bias. *The Journal of Human Resources 15*(3), 409–423.

Bredemeier, C., J. Gravert, and F. Juessen (2019). Estimating labor supply elasticities with joint borrowing constraints of couples. *Journal of Labor Economics 37*(4), 1215–1265.

Bronson, M. A. (2015). Degrees are forever: Marriage, educational investment, and lifecycle labor decisions of men and women. Working Paper.

Bronson, M. A. and M. Mazzocco (2018). Taxation and household decisions: an intertemporal analysis. Working Paper.

Brown, S., L. Farrell, M. N. Harris, and J. G. Sessions (2006). Risk preference and employment contract type. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 169*(4), 849–863.

Browning, M. and P.-A. Chiappori (1998). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica 66*(6), 1241–1278.

Browning, M., P.-A. Chiappori, and Y. Weiss (2014). *Economics of the Family.* Cambridge University Press.

Browning, M., A. Deaton, and M. Irish (1985). A profitable approach to labor supply and commodity demands over the life-cycle. *Econometrica 53*(3), 503–543.

Brunnermeier, M. K. and S. Nagel (2008). Do wealth fluctuations generate time-varying risk aversion? micro-evidence on individuals. *American Economic Review 98*(3), 713–36.

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

Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2017). The effect of wealth on individual and household labor supply: evidence from swedish lotteries. *American Economic Review 107*(12), 3917–46.

Charles, K. K., S. Danziger, G. Li, and R. Schoeni (2014). The intergenerational correlation of consumption expenditures. *American Economic Review 104*(5), 136–40.

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica 56*(1), 63–90.

Chiappori, P.-A. and M. Mazzocco (2017). Static and intertemporal household decisions. *Journal of Economic Literature 55*(3), 985–1045.

Chiappori, P.-A., J. A. Molina, J. I. Giménez, and J. Velilla (2019). Intertemporal labor supply and intra-household commitment. IZA Discussion Paper.

Chiappori, P.-A. and M. Paiella (2011). Relative risk aversion is constant: Evidence from panel data. *Journal of the European Economic Association 9*(6), 1021–1052.

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

Coate, S. and M. Ravallion (1993). Reciprocity without commitment: Characterization and performance of informal insurance arrangements. *Journal of Development Economics 40*(1), 1–24.

Conesa, J. C. and D. Krueger (2006). On the optimal progressivity of the income tax code. *Journal of Monetary Economics 53*(7), 1425–1450.

Croson, R. and U. Gneezy (2009). Gender differences in preferences. *Journal of Economic Literature 47*(2), 448–74.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica 59*(5), 1221–1248.

Del Boca, D. and C. Flinn (2012). Endogenous household interaction. *Journal of Econometrics 166*(1), 49–65.

Dercon, S. and P. Krishnan (2000). In sickness and in health: Risk sharing within households in rural ethiopia. *Journal of Political Economy 108*(4), 688–727.

Devereux, K. and L. Turner (2016). Divorce, Remarriage, and Fertility with On-the-Marriage Search. Working Paper.

Díaz-Jiménez, J., A. Glover, and J.-V. Ríos-Rull (2011). Vupdated facts on the us distributions of earnings, income and wealth in the united states: 2007 update. *Federal Reserve Bank of Minneapolis Quarterly Review, Summer 34*(1), 2–31.

Dixit, A., G. M. Grossman, and F. Gul (2000). The dynamics of political compromise. *Journal of Political Economy 108*(3), 531–568.

Doepke, M. and M. Tertilt (2009). Women's liberation: What's in it for men? *The Quarterly Journal of Economics 124*(4), 1541–1591.

Doepke, M. and M. Tertilt (2016). Families in macroeconomics. In *Handbook of Macroeconomics*, Volume 2, pp. 1789–1891. Elsevier.

Doepke, M. and M. Tertilt (2019). Does female empowerment promote economic development? *Journal of Economic Growth 24*(4), 309–343.

Domeij, D. and M. Floden (2006). The labor-supply elasticity and borrowing constraints: Why estimates are biased. *The Review of Economic Dynamics 9*(2), 242–262.

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

Druedahl, J. and T. H. Jørgensen (2017). A general endogenous grid method for multi-dimensional models with non-convexities and constraints. *Journal of Economic Dynamics and Control 74*, 87–107.

Dubois, P., B. Jullien, and T. Magnac (2008). Formal and informal risk sharing in ldcs: Theory and empirical evidence. *Econometrica 76*(4), 679–725.

Duflo, E. and C. Udry (2004). Intrahousehold resource allocation in cote d'ivoire: Social norms, separate accounts and consumption choices. Technical report, National Bureau of Economic Research.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics 133*(4), 1645–1692.

Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics 17*(2), 329–344.

Fernández, R. (2014). Women's rights and development. *Journal of Economic Growth 19*(1), 37–80.

Finke, M. S. and S. J. Huston (2013). Time preference and the importance of saving for retirement. *Journal of Economic Behavior & Organization 89*, 23–34.

Floden, M. and J. Lindé (2001). Idiosyncratic risk in the united states and sweden: Is there a role for government insurance? *Review of Economic Dynamics 4*(2), 406–437.

Foerster, H. (2019). The impact of post-marital maintenance on dynamic decisions and welfare of couples. Working paper.

Fortin, B. and G. Lacroix (1997). A test of the unitary and collective models of household labour supply. *The Economic Journal 107*(443), 933–955.

Foster, A. D. and M. R. Rosenzweig (2001). Imperfect commitment, altruism, and the family: Evidence from transfer behavior in low-income rural areas. *Review of Economics and Statistics 83*(3), 389–407.

Gemici, A. and S. Laufer (2014). Marriage and Cohabitation. Working Paper.

Gierlinger, J. and S. Laczó (2017). Matching to share risk without commitment. *The Economic Journal 128*(613), 2003–2031.

Golosov, M., A. Tsyvinski, and N. Werquin (2016). Recursive contracts and endogenously incomplete markets. In *Handbook of Macroeconomics*, Volume 2, pp. 725–841. Elsevier.

Greenwood, J., N. Guner, and G. Vandenbroucke (2017). Family economics writ large. *Journal of Economic Literature 55*(4), 1346–1434.

Guner, N., R. Kaygusuz, and G. Ventura (2011). Taxation and household labour supply. *The Review of Economic Studies 79*(3), 1113–1149.

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

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

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

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

Heathcote, J., K. Storesletten, and G. L. Violante (2009). Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics 1*(1), 319–354.

Heathcote, J., K. Storesletten, and G. L. Violante (2010). The Macroeconomic Implications of Rising Wage Inequality in the United States. *Journal of Political Economy 118*(4), 681–722.

Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review 104*(7), 2075–2126.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal Tax Progressivity: An Analytical Framework. *The Quarterly Journal of Economics 132*(4), 1693–1754.

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

Heer, B. and A. Maussner (2009). *Dynamic General Equilibrium Modelling* (2 ed.). Springer.

Hertzberg, A. (2019). Time-consistent individuals, time-inconsistent households. *Journal of Finance (forthcoming).*

Hintermaier, T. and W. Koeniger (2010). The method of endogenous gridpoints with occasionally binding constraints among endogenous variables. *Journal of Economic Dynamics and Control 34*(10), 2074–2088.

Hokayem, C. and J. P. Ziliak (2014). Health, human capital, and life cycle labor supply. *American Economic Review 104*(5), 127–31.

Holden, K. C. and P. J. Smock (1991). The economic costs of marital dissolution: Why do women bear a disproportionate cost? *Annual review of sociology 17*(1), 51–78.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control 17*(5-6), 953–969.

Hyslop, D. R. (2001). Rising us earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *American Economic Review 91*(4), 755–777.

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

Imrohoroğlu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy 97*(6), 1364–1383.

Iskhakov, F., T. H. Jørgensen, J. Rust, and B. Schjerning (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics 8*(2), 317–365.

Kabukçuoğlu, A. and E. Martínez-García (2016). The market resources method for solving dynamic optimization problems. Globalization and Monetary Policy Institute. Working Paper No. 274, Federal Reserve Bank of Dallas.

Kalai, E. (1977). Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica 45*(7), 1623–1630.

Kalai, E. and M. Smorodinsky (1975). Other solutions to nash's bargaining problem. *Econometrica 43*(3), 513–518.

Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics 2*(4), 53–87.

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.

Kehoe, P. J. and F. Perri (2002). International business cycles with endogenous incomplete markets. *Econometrica 70*(3), 907–928.

Kimball, M. S. (1988). Farmers' cooperatives as behavior toward risk. *American Economic Review 78*(1), 224–232.

Kimmel, J. and T. J. Kniesner (1998). New evidence on labor supply:: Employment versus hours elasticities by sex and marital status. *Journal of Monetary Economics 42*(2), 289–301.

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

Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *The Review of Economic Studies 63*(4), 595–609.

Konrad, K. A. and K. E. Lommerud (1995). Family policy with non-cooperative families. *The Scandinavian Journal of Economics 97*(4), 581–601.

Krebs, T. (2007). Job displacement risk and the cost of business cycles. *American Economic Review 97*(3), 664–686.

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

Krusell, P., T. Mukoyama, A. Şahin, and A. A. Smith Jr (2009). Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics 12*(3), 393–404.

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy 106*(5), 867–896.

Kuhn, M. and J.-V. Ríos-Rull (2016). 2013 update on the us earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review 37*(1), 2–73.

Laczó, S. (2014). Does risk sharing increase with risk aversion and risk when commitment is limited? *Journal of Economic Dynamics and Control 46*, 237–251.

Laczó, S. (2015). Risk sharing with limited commitment and preference heterogeneity: Structural estimation and testing. *Journal of the European Economic Association 13*(2), 265–292.

Lawrance, E. C. (1991). Poverty and the rate of time preference: evidence from panel data. *Journal of Political Economy 99*(1), 54–77.

Ligon, E. (2002). Dynamic bargaining in households. Manuscript, University of California, Berkeley.

Ligon, E., J. P. Thomas, and T. Worrall (2000). Mutual insurance, individual savings, and limited commitment. *Review of Economic Dynamics 3*(2), 216–246.

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

Lise, J. and K. Yamada (2019). Household sharing and commitment: Evidence from panel data on individual expenditures and time use. *The Review of Economic Studies 86*(5), 2184–2219.

Low, H., C. Meghir, L. Pistaferri, and A. Voena (2018). Marriage, labor supply and the dynamics of the social safety net. Technical report, National Bureau of Economic Research.

Ludwig, A. and M. Schön (2018). Endogenous Grids in Higher Dimensions: Delaunay Interpolation and Hybrid Methods. *Computational Economics 51*(3), 463–492.

Lundberg, S. and R. A. Pollak (1993). Separate spheres bargaining and the marriage market. *Journal of Political Economy 101*(6), 988–1010.

Lundberg, S. and R. A. Pollak (1994). Noncooperative bargaining models of marriage. *American Economic Review 84*(2), 132–137.

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.

Manser, M. and M. Brown (1980). Marriage and household decision-making: A bargaining analysis. *International Economic Review 21*(1), 31–44.

Marcet, A. and R. Marimon (2019). Recursive contracts. *Econometrica 87*(5), 1589–1631.

Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies 74*(3), 857–895.

Mazzocco, M., C. Ruiz, and S. Yamaguchi (2013). Labor supply, wealth dynamics and marriage decisions. Working Paper.

Mazzocco, M., C. Ruiz, and S. Yamaguchi (2014). Labor supply and household dynamics. *American Economic Review 104*(5), 354–59.

McElroy, M. B. and M. J. Horney (1981). Nash-bargained household decisions: Toward a generalization of the theory of demand. *International Economic Review 22*(2), 333–349.

McManus, P. A. and T. A. DiPrete (2001). Losers and winners: The financial consequences of separation and divorce for men. *American Sociological Review 66*(2), 246.

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

Meier, V. and H. Rainer (2015). Pigou meets Ramsey: gender-based taxation with non-cooperative couples. *European Economic Review 77*, 28–46.

Mommaerts, C. (2016). Long-term care insurance and the family. Working paper.

Morten, M. (2019). Temporary migration and endogenous risk sharing in village india. *Journal of Political Economy 127*(1), 1–46.

Nash, J. F. (1950). The bargaining problem. *Econometrica 18*(2), 155–162.

Obermeier, T. (2019). Intra-household inequality and the optimal progressivity of the income tax. Working paper.

OECD (2019). OECD Family Database.

Oikonomou, R. (2018). Unemployment insurance with limited commitment wage contracts and savings. *The BE Journal of Macroeconomics 18*(1), 1–20.

Oikonomou, R. and C. Siegel (2015). Capital taxes, labor taxes and the household. *Journal of Demographic Economics 81*(3), 217–260.

Ortigueira, S. and N. Siassi (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics 60*(6), 650–666.

Park, Y. (2014). Optimal taxation in a limited commitment economy. *The Review of Economic Studies 81*(2), 884–918.

Pencavel, J. (1986). Labor supply of men: a survey. In *Handbook of Labor Economics*, Volume 1, pp. 3–102. Elsevier.

Peterson, R. R. (1996). A re-evaluation of the economic consequences of divorce. *American Sociological Review 61*(3), 528–536.

Potoms, T. (2019). Intra-household insurance, leverage constraints and housing. Working Paper.

Rainer, H. (2007). Should we write prenuptial contracts? *European Economic Review 51*(2), 337–363.

Reynoso, A. (2019). The impact of divorce laws on the equilibrium in the marriage market. Working paper.

Robinson, J. (2012). Limited insurance within the household: Evidence from a field experiment in kenya. *American Economic Journal: Applied Economics 4*(4), 140–64.

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

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *The Review of Economic Studies 68*(1), 205–229.

Shephard, A. (2019). Marriage market dynamics, gender, and the age gap. Working Paper.

Stevenson, B. (2007). The impact of divorce laws on marriage-specific capital. *Journal of Labor Economics 25*(1), 75–94.

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.

Stokey, N., R. Lucas, and E. Prescott (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Storesletten, K., C. I. Telmer, and A. Yaron (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. *European Economic Review 45*(7), 1311–1339.

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

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

Theloudis, A. (2018). Wages and family time allocation. Working Paper.

Thomas, J. and T. Worrall (1988). Self-enforcing wage contracts. *The Review of Economic Studies 55*(4), 541–554.

Turon, H. (2019). Home production of childcare and labour supply decisions in a collective household model. *IZA Discussion Paper No. 12148*.

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review 105*(8), 2295–2332.

Wu, C. and D. Krueger (2015). How much consumption insurance in bewley models with endogenous family labor supply? Technical report, Working paper, University of Pennsylvania.

Ziliak, J. P. and T. J. Kniesner (1999). Extimating life cycle labor supply tax effects. *Journal of Political Economy 107*(2), 326–359.