

UNIVERSITY COLLEGE LONDON

DOCTORAL THESIS

**Essays in Consumption
Inequality and the Allocation of
Household Resources**

Alexandros Theloudis

A dissertation submitted in fulfillment
of the requirements for the degree of

Doctor of Philosophy

of

University College London

Department of Economics

September 7, 2016

I, Alexandros Theloudis, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Dedicated to my parents

Abstract

This thesis consists of three self-contained papers in household economics. Each uses an empirically tractable life-cycle model of consumption and family labor supply to study aspects of economic behavior of households, such as the allocation of expenditure among household members or the allocation of spousal time across paid and non-paid activities. Emphasis is put on modelling intra-household interactions.

The opening chapter examines how married people's allocation of time responds to wages and the gender wage gap. It develops a life-cycle collective model for spouses who allocate time across market work, home production, and leisure. The model features lack of commitment to lifetime marriage and the gender wage gap affects intra-family bargaining power. Results from the PSID suggest that the narrowing gender wage gap since 1980 improved women's bargaining power in the family resulting in a shift of household work from women to their husbands. The model is used to assess, counter factually, the implications of gender wage equality for family time allocations.

The second chapter studies how individual and aggregate consumption in the family respond to idiosyncratic wage shocks using a collective life-cycle model that features public and private consumption, endogenous family labor supply, asset accumulation, correlated wage shocks, and lack of spousal commitment to lifetime marriage. Preliminary results from the PSID suggest strong labor and consumption response to wage shocks and that hours and consumption are substitute goods at the intensive margin of labor supply. Wages have an economically significant effect on intra-family bargaining powers.

The last chapter studies theoretically the transmission of income shocks into consumption across households that exhibit unobserved preference heterogeneity. Heterogeneity is nonparametric and nonseparable from household preferences. I show how any moment of the distribution of consumption and labor supply elasticities can be identified with readily available household panel data. Identification does not rely on any specific parametrization of household preferences or their distribution.

Acknowledgements

This dissertation is the outcome of nearly six years of doctoral studies at UCL (including an initial year of coursework). These years were deeply fulfilling, educating, and transforming; despite the many hard moments the PhD inevitably had, I look back at them with appreciation and nostalgia.

I want to express my sincere gratitude to my supervisors Richard Blundell and Ian Preston. I thank them wholeheartedly for their support, guidance, endless time and patience throughout these years. I learned enormously from them. Their supervision goes far beyond the contents of this Thesis; they shaped deeply my economic thought and served as role models in their capacity as researchers, teachers, humans. It is an honor to call myself their student.

Special thanks go to Luigi Pistaferri who hosted me at the Department of Economics at Stanford University in 2014. I am grateful for his generous advice and time and for treating me like his own student. I consider my research visit at Stanford to be a turning point in my PhD, and Luigi Pistaferri is hugely responsible for this. Special thanks also go to Eric French who often acted as a third, unofficial, supervisor and provided guidance and advice. I am particularly grateful to him for his support during the job market months in 2015/2016. I am also grateful to Valérie Lechene for supporting me during the job market and for her insightful advising throughout the PhD.

I wish to thank Costas Meghir for playing a special role in my PhD. Costas Meghir was my advisor upon my admission to UCL. Although we did not work together as he left for Yale while I was still working on the first year's courses, he helped me shape and sharpen my research focus at a time when I was still

trying to understand where the inequality and labor literatures stood.

I owe thanks to many UCL faculty members and visitors for numerous discussions we had over the years. Among others, I thank Orazio Attanasio, Pierre-André Chiappori, Martin Cripps, Mariacristina De Nardi, Tim Lee, Jeremy Lise, Suphanit Piyapromdee, Itay Saporta Eksten, and Michela Tinicani. Magne Mogstad was particularly encouraging throughout the PhD; I thank him for this and for offering me numerous research opportunities, including at Statistics Norway. Finally, thanks go to faculty members at the University of Essex, where I studied Economics at the Master's level, for encouraging me to apply to UCL; João Santos Silva played a particular role at that time and he has been following my progress ever since.

This dissertation would not have been possible without the generous financial support I received initially from the Department of Economics at UCL, through a W.M. Gorman scholarship, and subsequently from the Economic and Social Research Council (ESRC). The ESRC funded most of my research travels to conferences and workshops as well as my research visit at Stanford (thanks to ESRC's Doctoral Training Centre at UCL and Shelley Channon). Finally, the Royal Economic Society provided financial support in the last year of the PhD through a Junior Fellowship.

I am indebted to many friends and coursemates at UCL, Alireza, Arun, Eeva, Eleni, Lukas, Mario, Ming, Priscilla, Valerio, for making these years the most enjoyable. I am mostly indebted to dear friends in London and Greece who closely supported me in my pursuits. Thank you guys.

Finally and most importantly, my most special thanks go to my family, my parents Giannis and Nopi, and my sister Dora Theloudis. I cannot express how grateful I am for all you have done and keep doing for me. I wish I could give back to you at least a little of your love and devotion to me.

Alexandros Theloudis
University College London
September 2016

Contents

Abstract	4
Acknowledgements	6
List of Figures	11
List of Tables	13
Introduction	14
1 Wages and Family Time Allocation	18
1.1 Introduction	18
1.2 An Empirical Overview	27
1.2.1 Data	27
1.2.2 Time Allocation Facts	29
1.2.3 The Gender Wage Gap	33
1.3 A Life-Cycle Collective Model without Commitment	38
1.3.1 Illustration of Key Features	39
1.3.2 Model	41
1.3.3 Parametrization	46
1.3.4 Retirement	53
1.4 Model Solution and Simulation	56
1.5 Identification and Estimation	59
1.5.1 External Processes	60
1.5.2 Model Parameters	64

1.5.3	Initialization of the Pareto Weight	69
1.6	Results	74
1.7	Implications of Model for Behavior	80
1.8	Implications of Gender Wage Equality	87
1.9	Discussion and Conclusions	94
2	Consumption Dynamics and Allocation in the Family	97
2.1	Introduction	97
2.2	A Life-Cycle Collective Model without Commitment	104
2.2.1	The Model	105
2.2.2	The Solution to the Family Problem	112
2.2.3	The Family Response to Wage Shocks	117
2.2.4	Identification of the Family Structure	122
2.3	A Life-Cycle Model for Singles	129
2.4	Empirical Implementation	133
2.4.1	Sample Selection and Definition of Variables	134
2.4.2	Pre-Estimated Parameters	142
2.4.3	Estimation Procedure	145
2.5	Results	148
2.5.1	Results with Separable Public Good and Full Commitment	151
2.5.2	Results with Nonseparable Public Good and Lack of Commitment	153
2.6	Discussion and Conclusions	156
3	Consumption Inequality across Heterogeneous Families	158
3.1	Introduction	158
3.2	A Life-Cycle Household Model for Consumption and Labor Supply	162
3.2.1	The Model	163
3.2.2	Closed-Form Expressions for Inequality	170
3.3	Identification	174
3.3.1	Income Shocks	175

3.3.2 Transmission Parameters of Income Shocks 176

3.3.3 Distribution of Preferences 177

3.4 Tests 180

3.4.1 Specification Test and Further Discussion 180

3.4.2 Unobserved Preference Heterogeneity 181

3.5 Conclusions 182

A Appendix to

Wages and Family Time Allocation 184

A.1 Data: Sample Selection and Variables 184

A.2 Model: Public and Private Consumption 187

A.3 Estimation: Exogenous Blocks 189

A.4 Estimation: Model Fit and Overidentification 191

B Appendix to

Consumption Dynamics and Allocation in the Family 199

B.1 Frisch Elasticities 199

B.2 Approximations to First Order Conditions 201

B.3 Approximation to Intertemporal Budget Constraint 205

B.4 Transmission Parameters 209

B.5 Measurement Error in Consumption 213

C Appendix to

Consumption Inequality across Heterogeneous Families 214

C.1 Approximations to First Order Conditions 214

C.2 Approximation to Intertemporal Budget Constraint 216

C.3 Derivation of $Cov(\Delta h_{2it}, \Delta h_{2it+1})$ 219

Bibliography 231

List of Figures

1.1	Average annual hours worked in the market	30
1.2	Employment trends: market participation, FT and PT work . .	31
1.3	Average weekly hours worked in the household	33
1.4	Non-participation in home production	34
1.5	Annual hours worked in the market and weekly hours worked in the household by parental status	34
1.6	Unconditional gender wage gap	36
1.7	Conditional gender wage gap	37
1.8	Simulated against actual wages (means)	61
1.9	Proportion of families in various fertility states	62
1.10	Graphical illustration of identification	68
1.11	Histogram of the initialized intra-family bargaining power of men in 1980-1989	73
1.12	Life-cycle model fit	79
1.13	Bargaining effects of a narrowing gender wage gap	83
1.14	Income/substitution effects of a narrowing gender wage gap . . .	85
1.15	Counterfactual life-cycle profiles of time allocations	92
2.1	The transmission of $\Delta u_{1it} > 0$ into public consumption and leisure	120
2.2	Pre-estimated parameters	144
A.1	Evolution of wages with mean age of household head	187
A.2	Actual and smoothed variances of permanent shocks	190
A.3	Non-targeted moments: transition probabilities	198

List of Tables

1.1	Major sample descriptive statistics	28
1.2	Time into market and household work	47
1.3	Fixed parameters	65
1.4	Minor sample descriptive statistics	70
1.5	Initialized intra-family bargaining power of men in 1980-1989 . .	72
1.6	Estimates of preference parameters	75
1.7	Estimates of the parameters in the Pareto weight	77
1.8	Proportions of people in different time allocations	78
1.9	Bargaining effects: changes in proportions of people in various time allocations	84
1.10	Income and substitution effects: changes in proportions of peo- ple in various time allocations	86
1.11	Equal pay counterfactuals	89
1.12	Counterfactual family time allocations	91
2.1	Frisch elasticities of spouse j	114
2.2	Demographics and labor market outcomes by year	135
2.3	Private consumption by year and marital status	139
2.4	Public consumption by year and marital status	140
2.5	Household wealth by year and marital status	141
2.6	Estimates of wage parameters	150
2.7	Estimates of Frisch elasticities	152
2.8	Bargaining effects	155

A.1	Second moments of wage shocks	192
A.2	Child care costs: price and probabilities	193
A.3	Earnings regressions: male and female divorcees	194
A.4	Married men's initialized bargaining power: correlation with spousal attributes	195
A.5	Targeted moments: proportions men	196
A.6	Targeted moments: proportions women	197
B.1	Estimates of consumption measurement error	213

Introduction

A household is a fundamental economic unit where some of the most important economic decisions are made, such as decisions pertaining to the level and type of household expenditure, the allocation of such expenditure among household members, the allocation of the members' time endowments across paid and non-paid activities, household investment in human capital, the level of investment in various risky and risk-free assets, child rearing, and many more. Analyzing such decisions involves, often, understanding the interactions among decision-making household members, their preferences, as well as factors affecting such interactions.

In this thesis I develop methods to study some of these issues, both theoretically and empirically, and I apply such methods to rich survey data from the United States. The emphasis of the first chapter is on the allocation of spousal time across different activities, whereas the second chapter puts emphasis on the allocation of expenditure among household members and, therefore, on intra-household consumption inequality. The third chapter looks theoretically into how observed household behavior can convey information about spousal preferences for consumption and time, even if such preferences exhibit unobserved cross-sectional heterogeneity.

In more detail, the first chapter studies the allocation of married people's time across various activities and how such allocation responds to wages and the gender wage gap. In the US, the gender wage gap has narrowed down by as much as 25% over the last three and a half decades. At the same time, women's

labor supply has increased and, while couples spend less time on household work, men's share of household work has risen. In this chapter I develop a life-cycle collective model for individuals in a household (spouses) who differ in preferences and bargaining power but share a common budget constraint. The spouses allocate their time across market work, home production, and leisure, and they also decide about public consumption and savings. The model features lack of commitment to lifetime marriage and the gender wage gap can affect the spouses' bargaining power in the household. I estimate gender-specific preferences and how intra-family bargaining power responds to the gender gap using data on married and divorced individuals from the Panel Study of Income Dynamics in the US. The results suggest that the narrowing gender wage gap improved women's bargaining power in the family resulting in a shift of household work from women to their husbands. The model is used to assess, counter factually, the implications of gender wage equality for family time allocations. If the gender gap was eliminated altogether, the proportion of women in full-time market work would increase by up to 32% whereas couples' time into home production would decrease by up to 21% as women would reduce their household work by up to 7 hours per week.

The second chapter studies how individual and aggregate consumption in the family respond to idiosyncratic wage changes using a collective life-cycle model for a household of two decision-making spouses. The model incorporates endogenous family labor supply, public but also private consumption, asset accumulation, correlated wage shocks, and general nonseparable, spouse-specific preferences. Contrary to the first chapter, the model in the second chapter abstracts from home production. Wages enter the household budget constraint, but also the spouses' intra-family bargaining powers implying lack of spousal commitment to future allocations. I derive analytical expressions for the dynamics of earnings and consumption using a novel approach developed by Blundell and Preston (1998) and other papers; this approach relies on Taylor approximations to the lifetime budget constraint and the first-order conditions

of the household problem. I show how such analytical expressions can be used to identify the household structure (spouse-specific preferences, allocation of consumption between spouses, a rich set of bargaining effects) with panel data on hours, earnings, assets, and household-level consumption only. The identifying assumption is that spouses have the same preferences with their single counterparts. Preliminary evidence from the Panel Study of Income Dynamics suggests strong labor and consumption response to wage shocks and that hours and consumption are substitute goods at the intensive margin of labor supply. Wages have an economically significant effect on intra-family bargaining power, but not statistically so.

The last chapter studies the transmission of income shocks into consumption across two-earner unitary households with unobserved preference heterogeneity. The treatment of heterogeneity is general as heterogeneity is non-parametric and nonseparable from household preferences. I derive analytical expressions for cross-sectional earnings and consumption inequality following the approximations approach developed in Blundell and Preston (1998) and other papers. I show how these expressions can be used to identify any moment of the distribution of policy-relevant parameters, such as the consumption and hours elasticities with respect to wages, if panel data on consumption, hours, and earnings are available. Identification does not rely on any specific parametrization of household preferences or their distribution. Finally, I propose a test for unobserved preference heterogeneity that is straightforward to implement using panel data.

The first chapter contributes to our knowledge of collective household decisions, in this case time allocation decisions, in a dynamic setting without commitment. It builds on the initial static implementations of the collective model by Chiappori (1988, 1992) and Blundell et al. (2005), as well as the subsequent extensions to the dynamics case by Mazzocco (2007), Lise and Yamada (2014), and Voena (2015). The second chapter complements the first one by focusing

on consumption rather than time allocation decisions. It too builds on the aforementioned literature on the collective model, as well as on the literature that studies the transmission of wage/income shocks into consumption (for example, Blundell et al., 2008; Heathcote et al., 2014; Blundell et al., 2016). This is the literature to which also the third chapter contributes, advancing our understanding of identification of, effectively random, spousal preferences.

I hope the methods developed and findings presented in this thesis put future research in a better position to understand household decisions and serve as a basis to analyze the effects a wide variety of policies may have on household outcomes.

Chapter 1

Wages and Family Time Allocation

1.1 Introduction

How do wages affect married couples' allocation of time? What does a narrowing gender wage gap imply for the bargaining power spouses have in their households? How does their bargaining power affect their allocation of time? How would gender wage equality impact such allocation? To address these questions, I develop a rich life-cycle collective model of family time allocation, consumption and savings. Decision makers in the household (the spouses) choose jointly how to allocate their time across market work, work in the household (home production) and leisure in the presence of uncertainty in their wages and family composition.^{1,2} The model features lack of commitment to lifetime marriage meaning that the spouses do not commit to staying together for life. Changes in wages, and specifically changes in the gender wage gap, can induce shifts in the bargaining power that one or another spouse has in the household decision process. Such shifts reflect better or worse options that a spouse may have outside the household (for example in case of divorce) as a result of a changing gender wage gap. I estimate the model using data from

¹I use the terms 'household' and 'family' interchangeably throughout this chapter. The same applies to the terms 'decision makers', 'spouses', 'partners', or 'individuals'.

²The importance of distinguishing between leisure and non-market work is stressed in Becker (1965).

the PSID. I exploit cross-sectional variation in wages and family composition as well as the sharp decline in the gender wage gap that favored women after 1980. Focusing on one cohort whose life-cycle spans years 1980-2009, I find that the narrowing gender wage gap improved women's intra-family bargaining power resulting, primarily, in a shift of household work from women to their husbands. Such change in intra-family bargaining power is not consistent with full commitment between spouses. Finally, I use the model to investigate the likely implications the elimination of the gender wage gap has for family time allocations. In such counterfactual environment, the rate of female full-time market work would increase strongly by up to 32% even during the child-bearing years and women would enter the labor market when they previously would not participate. Moreover, the allocation of time into home production would become more equal between spouses (primarily with women reducing their much higher hours) but their joint total time into home production would decrease by as much as 21%.

Since 1980 the gender wage gap in the US, as measured by the ratio of male to female hourly wages, has fallen sharply by as much as 25%.³ This decline has occurred systematically over most of the last three and a half decades even if one accounts for cohort effects, spousal education, fertility and other factors. It is the result of growth in male and female real wages, with the latter outperforming the former. Over the same period of time, the proportion of women in full-time market work has increased strongly with women switching away from part-time work as well as entering the labor market when they previously did not participate. Women also halved the time they devote to home production whereas men's household work has remained flat. The chapter revolves currently around a single cohort whose life-cycle spans the period when the biggest changes in the gender wage gap and family time allocations occurred, namely the 3 decades after 1980. The chapter, therefore, focuses currently on how the narrowing gender gap affected married people's

³This figure is based on raw data from the PSID described in section 1.2.

time allocations *over their life-cycle* and serves as a *first* only step towards understanding how wages affect family time allocations *over time*.

I hypothesize that the narrowing of the gender wage gap has a direct effect on family planning (by increasing women’s monetary reward for market work) as well as an indirect one through impacting on the way decisions between spouses are made in the household. The channels through which wages likely affect family time allocations are the following. First, an increase in one’s hourly wage may render work in the labor market more attractive, along both the extensive (participation) and the intensive (hours) margins. Second, keeping labor supply fixed, a wage rise implies higher income and, in turn, higher expenditures and savings. If purchased goods (expenditures) are the material inputs to home production, higher expenditures may reduce or boost the time inputs to home production depending on the nature of complementarity between material and time inputs. Third, shifts in relative wages within a family may alter the task specialization spouses engage in; for example, a spouse with a relatively higher wage may engage fully in the labor market whereas the other one in home production. Fourth, shifts in relative wages may make a spouse’s outside option more or less attractive. To deter a person from exercising their outside option, their partner may consent to increase that person’s weight (bargaining power) in the family decision process which, in turn, is likely to affect a number of time allocation and other household outcomes. These channels are all interrelated reinforcing or mitigating each other making it harder to analyze the relationship between wages and married people’s allocation of time outside a structural model.⁴

The model allows for all the aforementioned channels. Two spouses⁵ have

⁴Another potential effect of wages and the gender wage gap is on the selection of individuals into marriage and, in general, on marital patterns. This chapter abstracts from this feature taking marriage as given. To some extent, Chiappori et al. (2015) address this question developing an equilibrium model of education, marriage, and labor supply. Expected returns in the labor market affect education and marital choices people make early on in their life-cycles; however, the paper shuts down many of the aforementioned channels through which wages (returns) affect choices, such as shifts in intra-family bargaining powers resulting from lack of commitment.

⁵I use the terms ‘spouses’ to refer to two decision making individuals in the model;

their own, gender-specific, preferences over private leisure (in the spirit of Chiappori, 1988, 1992) and a public consumption good (in the spirit of Blundell et al., 2005).⁶ The public good is produced in the household with inputs raw materials purchased in the goods market (‘public expenditures’) and time devoted to home production by each individual (‘household work’). The spouses are separately endowed with a fixed amount of time which they allocate across work in the labor market, work in the household, and leisure. An hour of work in the labor market is compensated by a gender-specific stochastic wage which individuals take as exogenous; earnings are used to purchase raw materials in the market or save for the future.

The spouses choose public expenditures/savings and their allocation of time to maximize the (expected, discounted, and inter-temporally separable) weighted sum of their respective, gender-specific, utility functions over their lifetime. The weights on their utilities can be seen as the bargaining power the spouses have in the household decision process; such bargaining power is not necessarily constant over time due to lack of commitment in the spirit of Mazzocco (2007) and Lise and Yamada (2014). Lack of commitment restricts choices by a set of marriage participation constraints, one per partner and time period, that ensure spouses receive at least as much utility from inside their joint household as they can possibly get from their outside option. To help fix ideas I take divorce to be the relevant outside option available to them, even though, strictly speaking, I do not have to specify this explicitly. I make the value of divorce for each spouse depend on their own wage offers, which reflect the value of their skills in the labor market, and on family composition regarding the presence and age of children. Wages and family composition are exogenous and subject to uncertainty.

the model applies equally to traditional nuclear families as well as more modern forms of cohabitation.

⁶The model in this chapter belongs to the family of ‘collective’ models as introduced by Chiappori (1988, 1992) and Apps and Rees (1988). These models treat the family as a group of individuals who act together under common constraints and, therefore, respect the fundamental principle of methodological individualism (Manser and Brown, 1980; McElroy and Horney, 1981).

Using cross-sectional and inter-temporal variation in wages and cross-sectional variation in family composition I identify time-use preferences for married men and women as well as how intra-family bargaining power changes with the gender wage gap over time. A major difficulty arises because wages affect the budget set *and* bargaining powers simultaneously. I distinguish between the two channels by fixing, essentially normalizing, intra-family bargaining power at the start of the life-cycle. Specifically, I form a proxy for the initial intra-family bargaining power by comparing married spouses' lifetime earnings in the hypothetical scenario of divorce. I estimate their hypothetical earnings should they get divorce using reduced-form information on divorcees in the PSID (*divorcees* because I treat divorce as the relevant outside option) rather than solving the divorcees' problem explicitly.

I estimate the model by the method of simulated moments using data from the PSID after 1980. Currently focusing on one cohort, the model reproduces life-cycle patterns of time allocations of married men and women. I find that, especially for families with young children, women's disutility from full-time market work is greater than the disutility from part-time work, which, in turn, is greater than the disutility from work in the household. Consumption and leisure are substitute goods for the majority of women; however for approximately 1/4 of them they are complements. Finally, men suffer greater disutility from work in the household than women do if the two supply the same amount of household hours.

The data reveal that both mechanisms (monetary reward and changes in bargaining power) are important when the gender wage gap narrows down. A 10% closing of the gap favoring women decreases women's household work by 14% and increases their rate of full time market work by approximately 4%. Half of the decrease in women's household work is due to the higher monetary reward of market work, thus to women switching to some form of market work. The other half is due to women becoming relatively stronger in the household decision process, and therefore able to extract more leisure. The

rise in women's market work is the result of two opposite forces: the higher monetary reward pushes women's market work up (dominating force) whereas their improved bargaining power pushes work down replacing it with leisure.

As the income women bring in the household rises, the spouses are in a better financial position to replace household chores such as child care or cleaning with similar services purchased from the market. In principle this could benefit men too by cutting down their household work (an income effect). In reality, however, men keep their household work unchanged as their weakening bargaining position counterbalances the income effect. As women become more powerful thanks to a narrower gender wage gap, they shift household work from themselves to their husbands.

These findings are suggestive of the likely implications gender wage equality would have for family time allocations. If women were paid on average their husbands' wage, female market participation would increase strongly throughout the life-cycle. The most striking effects would occur in the childbearing years when the rate of female full-time market work would rise to approximately 75% compared to 57% in the data (this change corresponds to an increase by approximately 32%). Only 1/8 of this increase comes from women switching from part- to full-time work; the rest comes from women entering the labor market when they previously did not participate. Gender wage equality would render the allocation of spousal time into home production more equal between spouses but it would also decrease the total time into household work by as much as 7 hours per week during the childbearing years (a decrease of 21% compared to the data). The timing of establishing gender wage equality in the life-cycle matters for the severity of the effects especially in the childbearing years. Perhaps not unexpectedly, the largest effects occur when wage equality is established early on in the partners' lives.

Relation to the literature This chapter builds on two strands of literature. On one side is the literature on models of household decision making with Chiappori (1988)'s and Apps and Rees (1988)'s collective concept being

the most prominent representation. As I illustrate below, there is a number of recent papers that extend the original static collective model to allow for life-cycle dynamics.⁷ On the other side is the literature that, from a unitary standpoint, studies the evolution of male or female life-cycle labor supply usually alongside a number of other choices such as consumption or retirement.

The papers in the first strand of literature that this article is mostly related to are Lise and Yamada (2014), Knowles (2013), and Mazzocco et al. (2014). Lise and Yamada (2014) use a dynamic collective model of the household with which the model in my article shares common features. They study how intra-family bargaining power varies across as well as within households when wage shocks hit. They estimate the parameters of the model using the first-order conditions and a unique panel dataset from Japan with expenditure information for each spouse. They find that relative wages affect intra-family allocations in the cross-section and large wage shocks induce changes in those allocations from one period to the next; the latter, they argue, serves as evidence against full intra-household commitment. Unlike Lise and Yamada (2014), I allow for extensive- as well as intensive-margin labor supply and I solve explicitly for spousal choices over the entire life-cycle. This is likely to be important if one expects the impact of wages -or of wage-related policies- on time allocations to vary with age and family composition.

Knowles (2013) asks how important bargaining between spouses is for labor supply since the 1970s. He develops a stylized two-period model in which the bargaining power depends on a marriage market equilibrium; he abstracts from life-cycle features such as savings or fertility. He finds intra-family bargaining is critical for explaining trends in gender-specific labor supply but he does not separate changes occurring *across* cohorts from changes occurring *within* cohorts. The present chapter, by contrast, uses a life-cycle model with savings, fertility, extensive as well intensive labor choices, and, finally, does not specify a form for intra-family bargaining (other than assuming ex-post

⁷Early empirical implementations of the *static* collective model include Browning et al. (1994) and Fortin and Lacroix (1997).

efficiency).

Mazzocco et al. (2014) investigate the interconnectedness of family labor supply, savings, and marital decisions using PSID data between 1984-1996. They develop a dynamic collective model with intensive and extensive labor supply and home production but abstract from public expenditures. Moreover, their paper does not allow changes in the gender wage gap to affect intra-family bargaining powers.

Fernández and Wong (2014) use a life-cycle collective model to study the increase in female labor supply in the second half of the 20th century. They abstract from home production and men's time allocations, and impose full commitment between spouses. Voena (2015) explores how divorce and property division laws impact married people's intertemporal choices using policy reforms in the US in the 1970s and 1980s. She specifies a life-cycle collective model for female market participation without home production; although laws controlling divorce and property division affect spouses' outside options in her model, serving as distribution factors, wages or the gender wage gap do not. Her findings support lack of intra-household commitment as in Mazzocco (2007), one of the first implementations of a dynamic collective model.⁸

There are several papers in the second strand of literature that this article relates to. French (2005) studies the labor supply and retirement behavior of men using a life-cycle model with wage and health uncertainty. He focuses particularly on the behavioral effects of social security benefits. Attanasio et al. (2008) study the strong increase in American women's labor force participation after the 1970s using a life-cycle model of labor supply, savings, and human capital. They argue that a narrowing gender wage gap and a declining cost of child care can explain the aforementioned increase. Eckstein and Lifshitz

⁸In addition, Gemici (2011) uses a dynamic collective model with Nash bargaining to study household migration decisions. A recent review of this literature, including static and dynamic collective models, is provided by Browning et al. (2014) and Chiappori and Mazzocco (2014). Lise and Seitz (2011) and Chiappori and Meghir (2014) argue why the intra-household allocation of resources should not be ignored.

(2011) also study women's employment rates, paying particular attention to the differential patterns that married and singles have experienced. Blundell et al. (2013) study the implications that welfare programs have in the short (labor supply) and the long run (human capital accumulation) using a life-cycle model of female labor supply, education, human capital, and savings. Finally, in an earlier paper Francesconi (2002) estimates a dynamic model of female labor supply allowing for endogenous fertility decisions but not savings.⁹

The papers in the second strand of literature, with their various specifications and assumptions, have three features in common: they focus on male *or* female labor supply, they abstract from home production, and they ignore intra-family allocations. By contrast, this chapter reserves, in principle, an explicit role for all these features. However, I abstract from endogenous human capital (that several of those papers model explicitly) for reasons pertaining to identification of the collective household structure (and which are discussed in section 1.3.2).

In relation to the literature, this chapter is the first one to (i) study female labor supply, on the intensive and extensive margins, using a collective life-cycle model with lack of commitment, home production, *and* household-level public expenditure; (ii) investigate the relationship between the gender wage gap and intra-family bargaining power, (iii) assess the implications of equal pay between men and women through eliminating counterfactually the gender wage gap.

The chapter is arranged as follows. Section 1.2 describes the data and the empirical facts that motivate this research. Section 2.2 develops the model of household decision making. Section 1.4 discusses technical aspects of the model and section 1.5 discusses identification and estimation. Section 2.5 presents the results. Section 1.7 discusses the implications of the model for behavior and section 1.8 describes the policy experiment. Section 1.9 concludes.

⁹Important earlier papers in this strand of literature also include Eckstein and Wolpin (1989), who model women's labor force participation and fertility choices when current participation affects future earnings, and van der Klaauw (1996), who models women's labor force participation jointly with their marital choices.

1.2 An Empirical Overview

This section overviews the data used in this chapter, lays out the family time allocation facts this study aims to reproduce, and discusses the evolution of the gender wage gap over time. Section 1.2.1 presents the data and some baseline summary statistics, section 1.2.2 illustrates the time allocation facts, and section 1.2.3 is devoted to the gender wage gap.

1.2.1 Data

This chapter uses data from the Panel Study of Income Dynamics (PSID). This provides rich income and employment data for households and their members since 1968 as well as information on time devoted to home production.

The PSID¹⁰ started in 1968 tracking a -then- nationally representative sample of households; repeated annually until 1997 the survey collected detailed information on incomes, market work, food consumption, and demographics of adult household members and their linear descendants should they split off and establish their own households. Over time the scope of the PSID widened allowing the collection of even richer information such as amounts of time devoted to household work (from late 1970s onwards). After 1997 the survey became biennial but also added information on a variety of household expenditures and wealth. I make no use of the expenditure or wealth information as this spans a relatively short period of time only.

I select men and women aged 25 to 65 from the core sample ('Survey Research Center') between years 1980 and 2009. I impose the aforementioned age restriction because the model in this chapter does not deal with early-life (education) or late-life (retirement) decisions. I split this into two distinct and non-overlapping samples: (i) a major sample of households of continuously married men and women throughout the years they are observed, and (ii) a minor sample of singles of both genders. I use the former for the main part of my analysis and I describe it in more detail below. I postpone a discussion of

¹⁰Detailed information on the PSID, as well as access to the data, is available at psidonline.isr.umich.edu.

Table 1.1: Major sample descriptive statistics

	Men	Women
% some college	0.63	0.55
% working	0.94	0.80
Annual earnings	65846	31352
Annual work hours	2255	1610
Hourly wage rate	28.86	18.73
Num. of kids	1.25	
Observations (household-year)	15917	

Notes: ‘some college’ is defined as any education above the 12th grade. ‘% working’ is defined as the proportion of those working in a given year. Earnings and working hours are presented for those working. Hourly wages are for those working using the central 96% of the relevant distribution. All monetary amounts are expressed in 2010 dollars and all descriptive statistics for stable households are calculated across all stable household-year observations.

Source: Major PSID sample.

the latter sample until section 1.5.3.

In the major sample I follow households headed by a married opposite-sex couple.¹¹ The analysis in this chapter revolves currently around one only widely-defined cohort. Additional cohorts will be added in future work. I define this cohort as those households whose male spouse is born between years 1943 and 1955. The average age of the male spouse is 30 in 1980 and 59 in 2009. A narrower definition of a cohort would be desirable but this is not possible without running into small sample sizes. Given that the age difference between spouses in approximately two thirds of households in this cohort does not exceed ± 3 years, I do not explicitly condition on similar years of birth for the female spouse. I remove inflation from all monetary values¹² and, to account partly for measurement error, I drop households for which earnings of a *working* spouse fall below 1% or above 99% of the (gender- and time-specific) distribution. Finally, I require that households are stable in that they do not experience compositional changes in the head couple. The resulting dataset is an unbalanced panel of 1279 households observed over at least two consecutive

¹¹I also consider couples that are permanently cohabiting (a tiny proportion in the data).

¹²I express all monetary amounts in 2010 dollars. To deflate I use the All-Urban-Consumers CPI available by the BLS at www.bls.gov/cpi.

years. More than 55% of households are observed for at least 10 years and more than 30% for at least 20. Some key descriptive statistics are presented in table 1.1; appendix A.1 provides further details.


I concentrate on continuously married (stable) couples because I do not solve explicitly for the divorcees' problem (see sections 2.2 and 1.5). The main caveat is whether excluding unstable couples (*i.e.* couples that separate or divorce) biases my results. I discuss the direction of this potential bias in section 2.5.

1.2.2 Time Allocation Facts

In this section I illustrate the main facts about married men's and women's life-cycle time allocations. Specifically I focus on the time they spend working in the labor market and in the household.

Figure 1.1 plots average annual hours of market work for workers and non-workers. Three features stand out. First, women work much less in the market than men. Second, over the first two thirds of their life-cycle, men's labor supply is flat at approximately 2,250 hours annually; women's labor supply on the other hand increases steadily from less than 1,000 hours annually to a peak of 1,550 hours at mean age 50. Third, both men and women decrease their hours of market work in the last third of their life-cycle, possibly due to retirement.

To understand these trends better, figure 1.2a plots the proportion of people who participate in the labor market over the life-cycle. A person is classified modestly as participating if he/she reports working at least 10 hours and earning at least \$10 in any given year. For women the picture is clear. There is a big increase in labor market participation over the first two thirds of their life-cycle and a subsequent decrease in the last third; this extensive margin shift is the main force behind the strong increase in women's working hours reported in figure 1.1. For men things are different. A nearly full participation in the first years is followed by a sudden downwards jump around mean age 43. Participation then flattens out again (at around 90% now) until

Figure 1.1: Average annual hours worked in the market

Notes: This figure plots average annual hours of market work for workers and non-workers. A 95% confidence interval appears in gray shade.


Source: Major PSID sample.

it starts declining in the last few years.

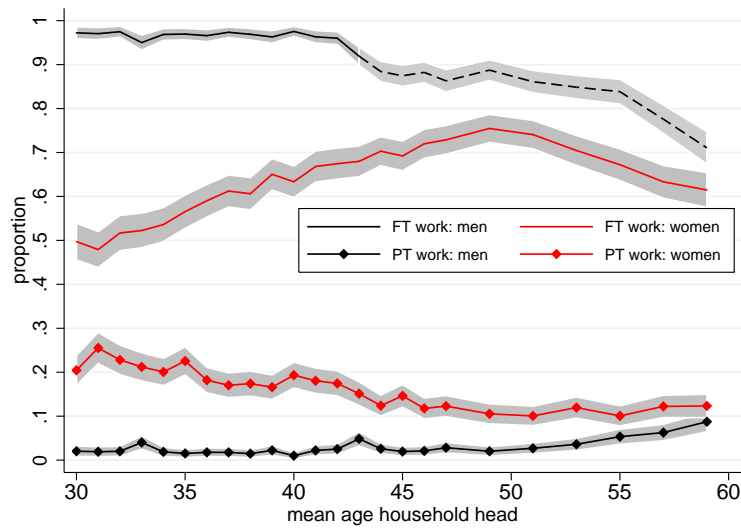
A careful look at the data flags up an inconsistency in the measure of male earnings that occurs in 1993 and affects men in the main sample at mean age 43 onwards. The definition of earnings changes slightly after 1993 and the available measure excludes some previously included earnings components such as the labor part of business income (see appendix A.1 for further information). This seems to be the reason behind the downwards jump in male employment at mean age 43. Indeed, until 1993 around 10 men in the major sample report 0 earnings every year and the majority of them also reports 0 working hours. After 1993, however, the number of men reporting 0 earnings jumps to around 70 every year with only 20% of them also reporting 0 hours. Among those reporting 0 earnings after 1993, mean annual working hours are around 1,800, *i.e.* sufficiently close to the unconditional mean of figure 1.1. I conclude that men's employment jump at mean age 43 is the result of a data design flaw and it does not reflect a true incident in men's labor supply.

Figure 1.2b delves deeper into the employment trends and plots the proportions of people working full- or part-time in the labor market (among all

Figure 1.2: Employment trends: market participation, FT and PT work



(a) Labor market participation




(b) Full-time and part-time market work

Notes: This figure plots the proportion of people who participate in the labor market, and the proportions of people working full- and part-time. A 95% confidence interval appears in gray shade. Data for men after mean age 43 suffer from a data design flaw (see main text for details).
Source: Major PSID sample.

workers and non-workers). A person is classified as working full-time (part-time) if he/she participates in the market and reports working more than 1,000 (up to 1,000) hours annually. The figure paints an opposite picture for married men versus women: men work full-time for most of their life-cycle (with the same caveat about employment around mean age 43) and they only start reducing slowly their full-time work in the last third of the life-cycle. Even then, a noticeable proportion seems to revert to part-time work rather than quit the market totally. Women, on the other hand, increase their full-time work by more than the overall increase in their participation, partly because they move gradually away from part-time work. Hence, the increase in female working hours in figure 1.1 is a combination of a strong increase in the extensive margin of labor supply (figure 1.2a) and a smaller increase in the intensive margin (figure 1.2b).

Turning to household work (time devoted to home production), figure 1.3 plots *weekly* hours of household work for married men and women including those who report 0 such hours. Household work refers to any work in and around the household, such as cooking or cleaning, but excludes time spent with children. Two features stand out. First, men supply much fewer hours than women. Second, women's hours drop dramatically over the first two thirds of their life-cycle and they level off in the last third. Men's hours, on the other hand, remain flat at approximately 7 weekly hours throughout the life-cycle.

To investigate these patterns further, figure 1.4 plots the proportion of people over time who report supplying 0 weekly hours to home production. To improve legibility, I plot the actual proportions (squares and circles) as well as separate smoothing curves that pass through the scatters. Around 13% of men do not participate in household chores whereas for women this proportion is effectively 0. As there are no obvious trends in the extensive margin of household work, one infers that the big drop in women's household work reported in figure 1.3 is almost exclusively due to a decrease in its

Figure 1.3: Average weekly hours worked in the household

Notes: This figure plots average weekly hours of household work for household workers and non-workers. A 95% confidence interval appears in gray shade.

Source: Major PSID sample.


intensive rather than extensive margin.

As one would expect, family time allocations vary across different subgroups of the population. The presence of kids in the household is likely to be one of the most important factors impacting their parents' time use. Indeed figure 1.5 redraws the initial graphs for market and household hours splitting the sample by the parental status of the household (parents versus non-parents). Two facts emerge. First, men's time allocation is not affected by the presence of children. Second, women's time allocation is affected severely by children, with childless women experiencing trends very similar to men (albeit at different magnitudes). These facts are true for work in the market and work in the household.

1.2.3 The Gender Wage Gap

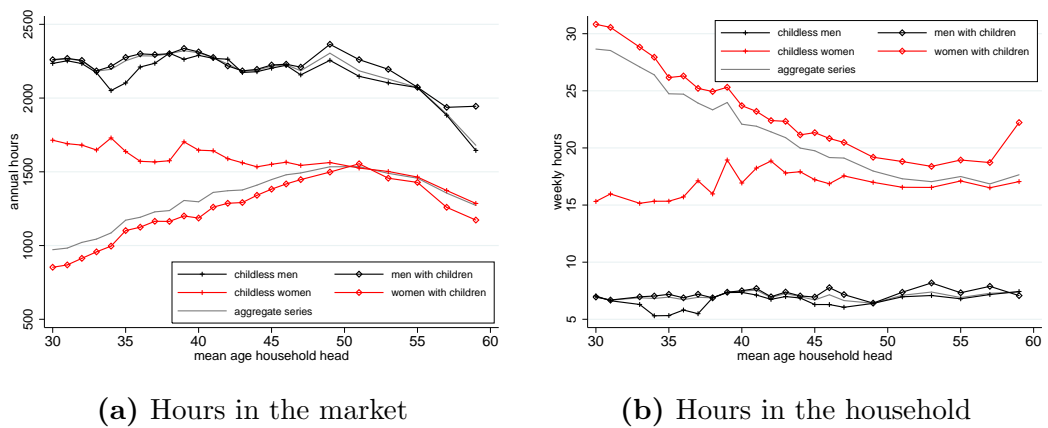
As the chapter focuses currently on one cohort, the gender wage gap statistics I report below refer to that cohort only (see section 1.2.1 for sample selection). I study the evolution of the wage gap over calendar years 1980-2009; these correspond to the lifespan of the aforementioned cohort. Note that the re-

Figure 1.4: Non-participation in home production



Notes: This figure plots the proportion of spouses who report 0 weekly hours of household work. A 95% confidence interval around the original (non-smoothed) proportions appears in gray shade. *Source:* Major PSID sample.

Figure 1.5: Annual hours worked in the market and weekly hours worked in the household by parental status



(a) Hours in the market

(b) Hours in the household

Notes: This figure plots average annual hours of market work for market workers and non-workers as well as average weekly hours of household work for household workers and non-workers. Confidence intervals have been suppressed to ease legibility of the graphs. *Source:* Major PSID sample.

restriction to one cohort implies that calendar time coincides with the mean age of the household head. I calculate the unconditional gender wage gap in two alternative ways: (a) as the ratio of median male-female hourly wages in the economy (‘economy-wide’ gender gap); (b) as the median ratio of male-female hourly wages in the family (‘within-household’ gender gap). Figure 1.6 plots these measures of the gender wage gap against calendar time (and, therefore, mean age). I plot the actual estimates of the gap (circles) as well as separate smoothing curves that pass through the scatters.


The gender wage gap narrowed down steadily in favor of women throughout their lifespan: in the start of the 1980s the ‘median’ man commands an hourly wage around 1.7-1.8 times higher than the ‘median’ woman; in 2009 the gender gap is around 1.3 or 25% lower. Within the family, the median ratio of spousal wages was approximately 1.55 in 1980 but only 1.35 in 2009 (13% lower). For completeness, figure A.1 in the appendix reports the levels of wages (medians and means) by gender. The narrowing of the gender gap is not specific to this particular cohort only. An earlier cohort¹³ also experiences an improvement in women’s relative wages, at least in the second half of their lifetime, even though the gap between genders had been everywhere wider than in the cohort of focus.

The narrowing of the gender wage gap is robust to a number of richer specifications. Figure 1.7, panel (a), plots the evolution of the gender (log) wage gap after controlling for spousal education and number of children, and after correcting women’s wages for selection into the labor market. In this graph I define the gender wage gap as

$$GWG_t = \text{median}(\tilde{w}_{1it}) - \text{median}(\tilde{w}_{2it})$$

where \tilde{w}_{jit} is the log hourly wage of a married person of gender j ($j = 1$ for men, $j = 2$ for women) after removing the effects of the aforementioned

¹³The earlier cohort consists of stable households whose male head is born between 1933 and 1945; his mean age is 30 in year 1970. The same exactly selection criteria apply to the earlier cohort as to the cohort of focus.

Figure 1.6: Unconditional gender wage gap**(a)** ‘Economy-wide’ gender wage gap **(b)** ‘Within-household’ gender wage gap

Notes: This figure plots two alternative definitions of the gender wage gap over time. Only the central 96% of the wage distribution by gender and time is used.

Source: Major PSID sample.

observable characteristics and correcting women’s wages for self-selection in the labor market.¹⁴ Figure 1.7, panel (b), plots the gender (log) wage gap *within* the family after controlling for spousal education and number of children, and after correcting it for women’s selection in the labor market. In this graph I define the gender wage gap as

$$GWG_t = \text{median}(\widetilde{\Delta w}_{it})$$

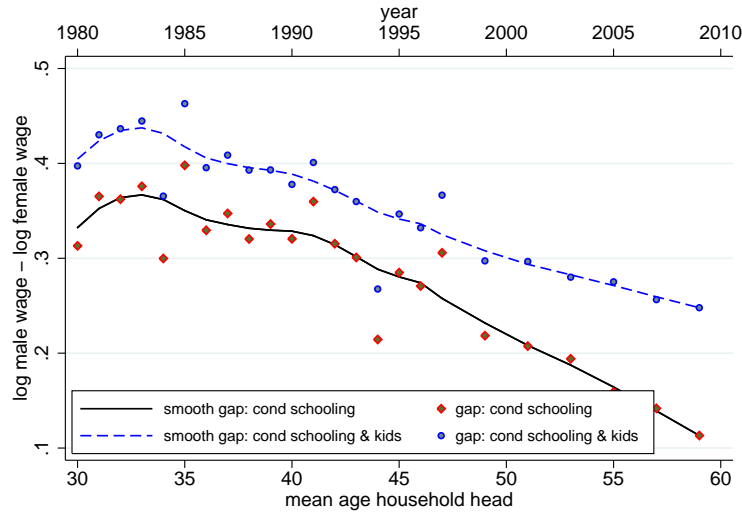
where $\widetilde{\Delta w}_{it}$ is the within-household gap in log hourly wages after removing the effect of observable characteristics and correcting for women’s participation selection. Appendix A.1 provides the details of these calculations, including the correction for women’s selection in the labor market.

Across all figures the picture that emerges points to an improvement of the economic status of women relative to that of men (at least as reflected upon their wages). Such improvement is robust to a number of factors that potentially affect the gender wage gap, such as women’s education, labor market participation, or number of children.

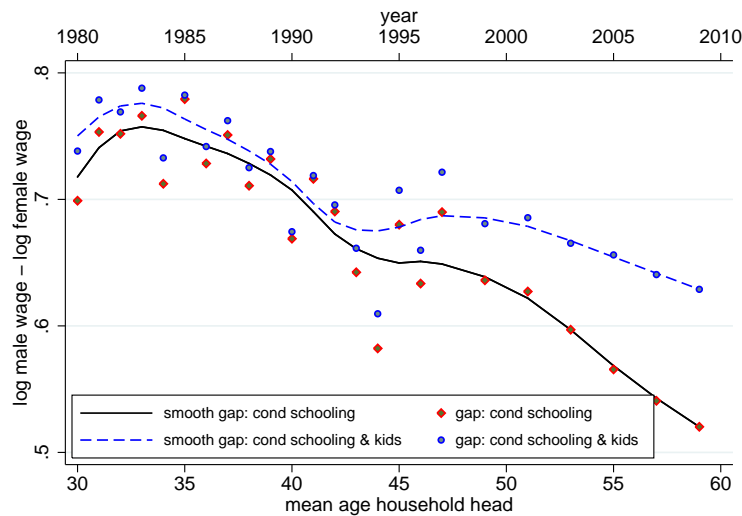
In a series of papers, Blau and Kahn (1997, 2006) investigate the reasons

¹⁴I do not correct male wages for selection into the labor market due to men’s very high, almost full, participation rate throughout of this study (see figure 1.2a).

Figure 1.7: Conditional gender wage gap



(a) 'Economy-wide' gender wage gap



(b) 'Within-household' gender wage gap

Notes: This figure plots the evolution of the gender wage gap over time in a number of different specifications. In graph (a), the gender wage gap is defined as $\text{median}(\tilde{w}_{1it}) - \text{median}(\tilde{w}_{2it})$ where \tilde{w}_{jit} is the log hourly wage of a married person of gender j ($j = 1$ for men, $j = 2$ for women) conditional on observable characteristics and after correcting women's wages for market participation selection. In graph (b), the gender wage gap is defined as $\text{median}(\widehat{\Delta}w_{it})$ where $\widehat{\Delta}w_{it}$ is the within-household gap in log hourly wages conditional on observable characteristics and after correcting for women's participation selection. Only the central 96% of the wage distribution by gender and calendar time is used.

Source: Major PSID sample.

behind the narrowing of the gender wage gap in the 1980s and 1990s (the years most of my data also come from). Using the PSID, they provide evidence of sex-biased institutional and technical change contributing to a faster growth in women's wages relative to men's. Such factors include improvements in the relative treatment of women in the labor market (possibly in response to the federal government's anti-discrimination policies in the 1970s) or demand-driven increased rents in industries where women had a comparative advantage (for example, services). In the light of this evidence, the present chapter aims to investigate the extent to which an *exogenous* narrowing of the gender wage gap can help explain household time allocations.

A caveat is due here. As the chapter currently focuses on one cohort only, one cannot separate calendar-time from life-cycle effects on the evolution of the gender wage gap or time allocations. The patterns in the figures above are likely to embody both types of effects. Therefore, the present chapter should be seen as investigating the relation between wages and family time allocations within one cohort's life-cycle, and this is only a first step towards understanding how the gender wage gap may affect family time allocations *over time* and *across cohorts*.

1.3 A Life-Cycle Collective Model without Commitment

This section develops the life-cycle collective model of family time allocation, public consumption and savings, which features lack of commitment to lifetime marriage. Two spouses are characterized by their own, gender-specific, preferences; each of them is fit to work in the labor market and earn a gender-specific hourly wage that is subject to idiosyncratic productivity shocks.

The life-cycle consists of two distinct periods: the working period, when the couple may also have children, and the retirement period. In section 1.3.1 I summarize the key features of the model during the working period. The details are given in section 1.3.2, where I lay out the model's building blocks

including its recursive formulation, and in section 1.3.3, where I detail the model's specification. Section 1.3.4 describes the retirement period.

1.3.1 Illustration of Key Features

The decision making spouses, subscripted by $j = \{1, 2\}$, consume a public (non-rival) good and allocate their time to leisure, market work, and home production. There may be children in the household but children are not decision makers.¹⁵ Spouse j has preferences U_j given by

$$U_j(Q, l_j; \mathbf{z}_j).$$

Here Q is the public consumption good and l_j is j 's private leisure. \mathbf{z}_j is a vector of observable taste shifters affecting j 's preferences; possible taste shifters are j 's education or the number and age of his/her children. An extension to preferences over private consumption goods is considered in appendix A.2.

The public good Q is produced domestically via a home production function given by

$$f(K, \tau_1, \tau_2; \mathbf{Z})$$

with inputs public expenditures K and time τ_j devoted to home production by each partner. The public good comprises items such as food at home or a clean house. In the former case K can be viewed as the amounts paid in grocery shopping whereas τ_j as the time each partner spends cooking. Here \mathbf{Z} is a vector of production shifters for which the obvious candidates are again the number and age of children in the household.

The partners stay together as members of the same household from period $t = 0$ (age 30) until the deterministic end of their working (T ; 30 years later) and retirement lives (T^R ; 40 years later). For simplicity I assume that both individuals are of the same age and post the schooling periods of their lives. I do not model marriage decisions; instead the focus of this chapter is on

¹⁵See Dauphin et al. (2011) and Dunbar et al. (2013) for static collective models where children act as decision makers.

the partners' choices after they have formed a household (i.e. conditional on marriage). The model accounts for initial conditions that arise from assortative patterns in the marriage market (see the wage process in section 1.3.3).

The spouses do not commit *ex ante* to one another for life. In each period that they stay together, they do so because each of them satisfies, among other things, their participation constraints in the household. Such constraints take the form of lower bounds that the utility each partner enjoys from inside the household must respect in each time period. The participation constraints essentially ensure that both partners enjoy at least as much utility from inside their joint household as they could possibly enjoy from their best outside option, which I take to be divorce.¹⁶ The outside options (the lower bounds) are not constant over time or across different states of the world; this changing nature of theirs imposes limits to commitment and risk sharing between spouses and affects household behavior. In this chapter, I make the outside options depend on the wages the spouses can command in the labor market to reflect the possibility that higher paid individuals may be able to attract better outside options.¹⁷

During the working period of life, I model annual choices over public consumption/savings and the allocation of time across leisure, market work, and work in the household. Market work generates income to fund public expenditures in the goods market or save for the future; work in the household contributes to the home production of the public consumption good. Publicness of consumption is an important element in the model as it permits economies of scale and complementarities between partners' preferences regardless of the specific functional forms that will represent them.

The value of each individual's time in the labor market is captured by the

¹⁶Consistent with most of the literature (Chiappori et al., 2002; Knowles, 2013; Voena, 2015) I choose divorce as the spouses' most relevant outside option. Other papers, however, consider non-cooperative cohabitation as a more realistic outside option (see, for example, Lechene and Preston, 2011).

¹⁷By contrast, I assume that savings or labor supply during marriage do not affect spouses' outside options. This simplification ensures the model's tractability and permits identification of the household structure.

hourly wage they can earn. Inside the model individuals cannot affect their wage and the model abstracts from human capital accumulation and similar features. The wage is seen as the exogenous gender-specific price of one's skills in the labor market and the individuals take it as given in each period.¹⁸ Wages can affect the trade-off among the different activities one can engage in and, therefore, the extent to which one or another individual specializes in market versus household work.

Family composition regarding children is an important determinant of individual choices during the working period of life. To capture the impact of children on behavior I model an exogenous stochastic 'fertility' process that reproduces the dynamics observed in the data over the life-cycle. Individuals make choices conditional on their family composition rather than choosing 'fertility' explicitly (something that would complicate the model considerably).¹⁹

1.3.2 Model

Given the previous points, the household in the working period of life can be seen as solving

$$\max_{\{Q_t, A_{t+1}, l_{jt}, \tau_{jt}\}_{t=0, j=\{1,2\}}^{t=T}} \mathbb{E}_0 \sum_{t=0}^T \beta_t U_1(Q_t, l_{1t}; \mathbf{z}_{1t}) \quad (1.1)$$

subject to the following constraints

$$\mathbb{E}_0 \sum_{t=0}^T \beta_t U_2(Q_t, l_{2t}; \mathbf{z}_{2t}) \geq \mathcal{U}_2(\mathbf{x}_1, \mathbf{x}_2) \quad (1.2)$$

$$A_t + \sum_{j=1}^2 w_{jt} h_{jt} = K_t + CC_t(h_{2t}, N_t) + \frac{A_{t+1}}{1+r} \quad A_{t+1} \geq \underline{A}_{t+1} \quad (1.3)$$

$$U_1(Q_t, l_{1t}; \mathbf{z}_{1t}) \geq \bar{U}_1(w_{1t}; \mathbf{d}_{1t}; \mathbf{z}_{1t}) \quad (1.4)$$

¹⁸The wage may be a function of prior educational choices but these are outside the control of individuals in the time frame of this model. See Blundell et al. (2013) or Chiappori et al. (2015) for a treatment of schooling choices in the context of a dynamic unitary or collective model respectively.

¹⁹Francesconi (2002) and Keane and Wolpin (2010) are examples of studies that endogenize fertility, both in a unitary context.

$$U_2(Q_t, l_{2t}; \mathbf{z}_{2t}) \geq \bar{U}_2(w_{2t}; \mathbf{d}_{2t}; \mathbf{z}_{2t}) \quad (1.5)$$

$$Q_t = f(K_t, \tau_{1t}, \tau_{2t}; \mathbf{Z}_t) \quad (1.6)$$

$$l_{jt} + h_{jt} + \tau_{jt} = \mathcal{T} \quad j = \{1, 2\}. \quad (1.7)$$

Constraints (1.3)-(1.7) must be satisfied in every period t . Expression (1.1) involves the maximization of the first individual's time-0-expected discounted lifetime utility; β_t is the common discount factor at t . Expression (1.2) is a promise keeping constraint, essentially an agreement set out at $t = 0$ that individual 2's expected discounted lifetime utility will not fall below a minimum level \mathcal{U}_2 (more on this to follow). Equation (1.3) is the sequential budget constraint linking available resources to expenditure and savings in each period of working life, (1.4)-(1.5) are the participation constraints in the household, (1.6) is the household production function, and (1.7) is the time budget per individual for a total time endowment \mathcal{T} . Much of the notation is already introduced; the remaining notation is as follows: (i) in the budget constraint A_t is household common assets, w_{jt} is spouse j 's hourly wage at t , h_{jt} is his/her hours of market work, $CC_t(h_{2t}, N_t)$ is child care costs that families with young children may have to meet (N_t summarizes the family composition; more on this to follow), r is the deterministic and known market interest rate, and \underline{A} is a borrowing limit; (ii) in the participation constraints $\bar{U}_j(\cdot)$ is the utility individual j can get from his/her outside option at t . The above program is written *as if* household member 1 makes all the choices in the household which obviously goes against the collective spirit. Decentralization is feasible but requires a combination of Lindahl (personal) and shadow prices for Q because this is a good that is both public and domestic (see Chiappori and Meghir, 2014).

In writing the outside options I have assumed that only exogenous variables enter \bar{U}_j , mainly the wage, the observable taste shifter \mathbf{z}_{jt} , and a vector of distribution factors \mathbf{d}_{jt} . By distribution factors I refer to any exogenous variables that affect choices through shifting partners' outside options but not

their preferences or the budget set.²⁰ Allowing the outside option to depend on individual choices while married would lead to inefficient allocations of time and would jeopardize the model's tractability. To see why, suppose \bar{U}_j is an increasing function of one's market work (say, through the dependence of wages on past labor supply). In this case the individual supplies labor for two reasons: first, labor generates income that can be used to buy current and future goods; second, labor improves one's outside option and boosts, consequently, his/her bargaining power in the household. As a result labor is over-supplied in this family beyond what is Pareto optimal and both partners can be better off if they agree to supply less. For a detailed illustration of this point see section 6.2.3 in Browning et al. (2014).

The assumption that only exogenous variables enter \bar{U}_j serves also another purpose, that of simplifying the representation of the model (1.1)-(1.7). Consider representing the problem by its Lagrangian formulation. Let ν_2 be the Lagrange multiplier on (1.2); also let $\tilde{\nu}_{1t}$ be the Lagrange multiplier on participation constraint (1.4) and $\tilde{\nu}_{2t}$ on (1.5). Then the above problem is equivalent to

$$\max_{\{Q_t, A_{t+1}, l_{jt}, \tau_{jt}\}_{t=0, j=\{1,2\}}^{t=T}} \mathbb{E}_0 \sum_{t=0}^T \beta_t \left[\left(1 + \frac{\tilde{\nu}_{1t}}{\beta_t}\right) U_1(Q_t, l_{1t}; \mathbf{z}_{1t}) + \left(\nu_2 + \frac{\tilde{\nu}_{2t}}{\beta_t}\right) U_2(Q_t, l_{2t}; \mathbf{z}_{2t}) \right]$$

or, written more compactly, to

$$\max_{\{Q_t, A_{t+1}, l_{jt}, \tau_{jt}\}_{t=0, j=\{1,2\}}^{t=T}} \mathbb{E}_0 \sum_{t=0}^T \beta_t \left[\mu_{1t} U_1(Q_t, l_{1t}; \mathbf{z}_{1t}) + \mu_{2t} U_2(Q_t, l_{2t}; \mathbf{z}_{2t}) \right] \quad (1.1')$$

subject to constraints (1.3), (1.6) and (1.7) only (section 3.1, Chiappori and Mazzocco, 2014). Then, $\mu_{jt} = \nu_j + \frac{\tilde{\nu}_{jt}}{\beta_t}$ is individual j 's bargaining power in the household at time t ($\nu_1 = 1$) or, equivalently, the weight his/her preferences carry in the household decision process at that time. Moreover, if one imposes

²⁰Chiappori et al. (2002) and Voena (2015) provide examples of distribution factors such as the sex ratio in the local marriage market or laws governing divorce and property sharing. See also Bourguignon et al. (2009).

the normalization $\mu_{1t} + \mu_{2t} = 1$ then μ_{jt} can be seen also as the Pareto weight a social planner attaches to member j 's preferences at t .²¹

What determines the weights μ_{jt} is given by the nature of the constraints that their underlying elements serve as Lagrange multipliers to. ν_j is the weight attached to individual j 's expected *lifetime* utility at the beginning of time, hence the lack of a time subscript. This may be a function of the individual's predetermined characteristics, some economy-wide attributes, as well as beginning-of-time expectations about possible changes in these characteristics/attributes in the future. I denote such variables by vector \mathbf{x}_j ; candidate variables may include spousal education, occupation, or parental income. These individual characteristics at $t = 0$ determine \mathcal{U}_2 in (1.2) and, as a consequence, contribute to determining the initial weight ν_j each spouse's expected lifetime utility carries. $\tilde{\nu}_{jt}$ is the multiplier on j 's participation constraint in period t . Whatever affects the outside option \bar{U}_j at t will affect $\tilde{\nu}_{jt}$ too. Pooling all the components of μ_{jt} together and normalizing the weights to add up to 1 implies

$$\mu_{jt} = \mu_j(\mathbf{x}_1, \mathbf{x}_2, w_{1t}, w_{2t}, \mathbf{d}_{1t}, \mathbf{d}_{2t}, \mathbf{z}_{1t}, \mathbf{z}_{2t}).$$

The aforementioned normalization of the sum of the weights to 1 is an additional reason why *both* partners' wages, distribution factors and pre-determined attributes enter μ_{jt} .

The Pareto weights μ_{1t} and μ_{2t} summarize the allocation of bargaining power in the household. Starting, hypothetically, from $\mu_{1t} = \mu_{2t} = \frac{1}{2}$, partner 1 becomes relatively more (less) powerful when $\mu_{1t} > \mu_{2t}$ ($\mu_{1t} < \mu_{2t}$). If the partners commit fully to never exploit their outside options, which is equivalent to removing the participation constraints, $\tilde{\nu}_{jt} = 0$ and $\mu_{jt} = \nu_j$ in each period ('full commitment' benchmark). If such commitment is not possible, j 's bargaining power will shift when any of the time-dependent factors affecting the outside options, therefore affecting $\tilde{\nu}_{jt}$ too, shifts ('no commitment'

²¹This normalization is possible if β_t is appropriately redefined in each period to avoid distorting the intertemporal incentives when $\mu_{1t} + \mu_{2t} \neq \mu_{1s} + \mu_{2s}$ for $t \neq s$.

benchmark). As an example, an increase in j 's wage will improve her outside option and result in $\tilde{v}_{jt} > 0$; this will raise her bargaining power by $\frac{\tilde{v}_{jt}}{\beta_t}$ and decrease her partner's power by the same amount.

Note that the model in this chapter cannot distinguish between lack of commitment ('no commitment') and limited commitment. The dependence of intra-family bargaining power on contemporaneous wages and distribution factors is, strictly speaking, consistent with the 'no commitment' framework of Mazzocco (2007) because the spouses adjust their bargaining power after *any* of the factors affecting their outside options changes (for example, after *any* change in their wages). Formally this implies that one spouse has their participation constraint always bind. By contrast, limited commitment, as developed in Ligon et al. (2002), requires that intra-family bargaining power shift only after a person's participation constraint binds, which will not necessarily be true after *any* wage change. In this chapter I model lack of commitment and I test it against full commitment that is nested within it. Rejection of full commitment, however, may be due to lack of commitment being a reasonable representation of the data or, alternatively, due to lack of commitment serving as a proxy for limited commitment.

Pareto efficiency The participation constraints prohibit the spouses from reaching the first-best or *ex-ante* efficient allocation of their resources. The solution to the above problem is *ex-post* efficient as the household still maximizes the weighted sum of their static utilities in each period. Ex-post efficiency implies that no better allocation of resources can take place once information at time t is revealed without violating the prevailing participation constraints; for details see Chiappori and Mazzocco (2014) or section 6.2.2 in Browning et al. (2014).

Recursive formulation Let $\mathcal{S}_t = \{\mathbf{z}_{1t}, \mathbf{z}_{2t}, \mathbf{Z}_t, \mathbf{x}_1, \mathbf{x}_2, w_{1t}, w_{2t}, \mathbf{d}_{1t}, \mathbf{d}_{2t}\}$ be the set of exogenous state variables at time t ; current period assets A_t is the endogenous state variable. Moreover, let $\mathcal{C}_t = \{K_t, l_{1t}, l_{2t}, \tau_{1t}, \tau_{2t}\}$ be the set of choice variables alongside next period's assets A_{t+1} (a total of 6 variables).

Finally, let \mathcal{U} denote the weighted sum of the partners' intra-temporal utility functions given by

$$\mathcal{U}(\mathcal{C}_t, \mathcal{S}_t) = \sum_{j=1}^2 \mu_j(\mathbf{x}_1, \mathbf{x}_2, w_{1t}, w_{2t}, \mathbf{d}_{1t}, \mathbf{d}_{2t}, \mathbf{z}_{1t}, \mathbf{z}_{2t}) U_j(f(K_t, \tau_{1t}, \tau_{2t}; \mathbf{Z}_t), l_{jt}; \mathbf{z}_{jt})$$

Program (1.1') can be written recursively as

$$V_t(A_t, \mathcal{S}_t) = \max_{\mathcal{C}_t, A_{t+1}} \{ \mathcal{U}(\mathcal{C}_t, \mathcal{S}_t) + \beta \mathbb{E}_{\mathcal{S}_{t+1} | \mathcal{S}_t} V_{t+1}(A_{t+1}, \mathcal{S}_{t+1}) \}$$

subject to constraints (1.3) and (1.7). V_t is the value function of the married household at t ; $\mathbb{E}_{\mathcal{S}_{t+1} | \mathcal{S}_t}$ denotes expectations over the exogenous state space at $t + 1$ conditional on its realization at t . Discounting is assumed geometric.

I do not write the value function for if the spouses divorce and I do not solve for the value of divorce numerically. The gender-specific value of divorce, which I only approximate using a reduced-form approach, enters the problem through dictating a normalization of the Pareto weight in the first few years of the family life-cycle (this point is discussed in section 1.5.3). Once this initial normalization takes place, I allow intra-family bargaining power to change subsequently with the gender wage gap and no further information on divorcees is used. This shortcut eases the computational needs of the model (saving the burden of solving the divorced man's and woman's life-cycle problems) but comes with the cost of restricting the estimation sample to stable households only (as shown in section 1.2.1). I discuss the implications of this restriction in section 2.5.

1.3.3 Parametrization

In each period of their life in the household, which I take to be one year, the partners maximize expected lifetime utility (1.1') taking as given their individual characteristics and their economic circumstances. Their individual characteristics are described by their education ($educ_1, educ_2$) and the woman's

Table 1.2: Time into market and household work

Activity	Intensity	Abbrev.	Daily hours
<i>Hours per day</i>			<i>24</i>
<i>Sleep & personal care</i>			<i>8</i>
<i>Remaining productive hours</i>			<i>16</i>
<i>Market</i>	no work	<i>NW</i>	0
	part time	<i>PT</i>	4
	full time	<i>FT</i>	8
<i>Household</i>	low	<i>L</i>	0.4
	low middle	<i>LM</i>	1.6
	high middle	<i>HM</i>	3
	maximum	<i>MAX</i>	6

utility costs of work (θ_2);²² their economic circumstances are described by age (t), their common assets (A), the presence of kids in the household and the age of the youngest among them (N), and their respective idiosyncratic productivity in the labor market (v_1, v_2).

Time allocations I assume time can take on discrete values across three activities: market work, work in the household (home production), and leisure. On a daily basis the time put into these activities by each spouse must add up to 24 hours net of 8 hours that people need for sleep and personal care (Biddle and Hamermesh, 1990). Table 1.2 summarizes the discrete values market- and household work can take. Market work can take on three values (for ‘no work’, ‘part time’, and ‘full time’) whereas work in the household four values (for ‘low’, ‘low middle’, ‘high middle’, ‘maximum’). The specific numerical values attached to these labels are not arbitrary; instead they correspond to the values most frequently reported in the PSID and the distribution implied by table 1.2 serves as a discrete approximation of the empirical distribution of time observed in the data.

For computational reasons, not all choices of market and household work

²²The model is written with education $educ_j$ as a state variable. However, current estimation results have education suppressed; results are reported without differentiating individuals by education.

are applicable to both men *and* women. I restrict men's household work to 'low' or 'low middle' (consistent with men in the PSID supplying very few hours to home production); I also restrict women's household work to 'high middle' or 'maximum' (consistent with women's observed household hours after 1980). Finally, I restrict men's market work to 'full time' only as there are very few men in my PSID sample not working full-time.²³ These restrictions imply that men's daily leisure is restricted between 6.4 and 7.6 hours whereas women's daily leisure between 2 and 13.

Preferences and home production I parameterize preferences U_j of spouse j by a non-separable function

$$U_j(Q_t, l_{jt}; \mathbf{z}_{jt}) = \frac{1}{1-\gamma} (Q_t/s(N_t))^{1-\gamma} \times \exp(g_j(l_{jt}; \mathbf{z}_{jt})) \quad (1.8)$$

where $\gamma > 1$ is the common coefficient of relative risk aversion and $s(N_t)$ is an equivalence scale that depends on the presence and age of the youngest child N_t .²⁴ Function $g_j(\cdot)$ reflects how the marginal utility of consumption changes with leisure (and thus with market and household work) and depends on N_t and j 's education $educ_j$. I specify

$$g_j(l_{jt}; \mathbf{z}_{jt}) = \begin{cases} g_1(l_{1t}; \mathbf{z}_{1t}) & = \sum_n \kappa_1^n \times \mathbf{1}[N_t = n] + \sum_e \kappa_1^e \times \mathbf{1}[educ_1 = e] \\ g_2(l_{2t}; \mathbf{z}_{2t}, \theta_2) & = \sum_n \kappa_2^n \times \mathbf{1}[N_t = n] + \sum_e \kappa_2^e \times \mathbf{1}[educ_2 = e] + \theta_2 \end{cases}$$

where the first row corresponds to $j = 1$ and the second row to $j = 2$. The sum \sum_n is over the different values $\{n\}$ of the age of the youngest child (if any) and the sum \sum_e is over the different levels $\{e\}$ of j 's education. For all possible

²³Part-time work for men does increase slightly towards the end of the life-cycle (figure 1.2b). However, I shut down men's market work choice for reasons pertaining to the feasibility of the computations herein.

²⁴The role of this equivalence scale is to account for the different needs that families with children of different ages have. It is not a means of comparison between a multi-member family and singles. I specify $s(N_t) = 1$ if the family has no children, $s(N_t) = 1.17$ if the youngest child is at most 5 years old, $s(N_t) = 1.23$ if it is between 5 and 10 years, and $s(N_t) = 1.32$ if it is between 10 and 18. These numbers come from the McClements equivalence scale after normalizing the scale to 1 in the case of a childless 2-adult-member family.

values $\{n\}$, $\{\mathbf{1}[N_t = n]\}_n$ constitutes a set of mutually exclusive dummies, each of which becomes active whenever $N_t = n$. Similarly, for all possible values $\{e\}$, $\{\mathbf{1}[educ_j = e]\}_e$ constitutes a set of mutually exclusive dummies, each of which becomes active whenever $educ_j = e$. The parameters κ_j^n and κ_j^e depend on the amount of leisure individual j enjoys and thus on the amount of market and household work he/she supplies.²⁵ Finally, θ_2 is a permanent individual-specific random cost of work that depends on the amount of work the female spouse puts into the market and household sectors.²⁶ In practice, θ_2 is drawn from a two point discrete distribution whose support and probability mass depend on the amount of work in the labor market and in the household. The distribution of θ_2 is estimated inside the model. From this specification it follows that there are two relevant preference shifters affecting U_j , these are $\mathbf{z}_{jt} = (N_t, educ_j)'$.

I parameterize the household production function f by the constant returns to scale specification

$$f(K_t, \tau_{1t}, \tau_{2t}; \mathbf{Z}_t) = K_t^\phi (\pi_1 \tau_{1t}^\varphi + \pi_2 \tau_{2t}^\varphi)^{\frac{1-\phi}{\varphi}} \quad (1.9)$$

with the additional restriction that $\pi_1 + \pi_2 = 1$. In the current specification the vector of production shifters \mathbf{Z}_t is left empty.

Budget constraint The budget constraint is given by the assets evolution equation (1.3). The borrowing limit \underline{A}_t is set at 10% of the family's minimum discounted lifetime income at t including pension income (more on

²⁵For a given age n of the family's youngest child I specify for males ($j = 1$):

$$\kappa_1^n(l_{1t}) = \kappa_{1,0}^n + \kappa_{1,1}^n \mathbf{1}[\tau_{1t} = LM]$$

and for females ($j = 2$):

$$\kappa_2^n(l_{2t}) = \kappa_{2,0}^n + \kappa_{2,1}^n \mathbf{1}[h_{2t} = FT] + \kappa_{2,2}^n \mathbf{1}[h_{2t} = PT] + \kappa_{2,3}^n \mathbf{1}[\tau_{2t} = MAX].$$

I specify $\kappa_j^e(l_{jt})$ similarly. I normalize $\kappa_{j,0}^n = 0, \forall j, n$, and $\kappa_{j,0}^e = 0, \forall j, e$; these are the dummies that correspond to j supplying the fewest possible hours in the labor market and the household. As a result, a positive g_j implies that work lowers the utility of consumption (given that $1 - \gamma < 0$) and that consumption and leisure are substitutes.

²⁶I do not model random costs of work for men because these cannot be identified when men effectively have a binary time use choice only, $\tau_{1t} = LM$ or $\tau_{1t} = L$.

pension income in section 1.3.4). This is not a generous borrowing limit as lifetime earnings are hindered by the possibility that the female spouse abstains from market. Hourly wages w_{jt} and child care costs $CC_t(h_{2t}, N_t)$ are described below.

Wages Each household member is fit to work in the market and earn an hourly wage that evolves according to the following permanent/transitory process

$$\begin{aligned}\ln w_{jt} &= \bar{W}_{jt} + v_{jt} + \xi_{jt} \\ v_{jt} &= v_{jt-1} + \zeta_{jt}.\end{aligned}\tag{1.10}$$

This process has been shown to fit the PSID data well (Blundell et al., 2016). The hourly wage is assumed exogenous and the individuals are viewed as price-takers in the labor market. The process is education specific but those subscripts are removed to simplify the notation. \bar{W}_{jt} is the mean of j 's log wage at t which is common across people of the same gender j and education.

The sum $v_{jt} + \xi_{jt}$ represents the stochastic idiosyncratic productivity which consists of a permanent and a transitory component, v_{jt} and ξ_{jt} respectively. The permanent component is the only economically relevant component and follows a unit root subject to a permanent shock ζ_{jt} . I allow shocks to be correlated across family members; specifically I assume ζ_{1t} and ζ_{2t} are jointly normally distributed according to

$$\begin{pmatrix} \zeta_{1t} \\ \zeta_{2t} \end{pmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{\zeta_1,t}^2 & \sigma_{\zeta_1\zeta_2,t} \\ \sigma_{\zeta_1\zeta_2,t} & \sigma_{\zeta_2,t}^2 \end{bmatrix} \right).$$

This process is estimated directly from the data and details are provided in section 1.5.1. The beginning-of-life permanent components, $v_{1t=0}$ and $v_{2t=0}$, are also correlated to reflect initial conditions that arise from the marriage market (assortative patterns in marriage); this correlation is also estimated directly in the data.

The transitory shock is viewed as measurement error that does not affect

choices; a similar approach is taken by French (2005) using PSID data or Blundell et al. (2013). It follows that the ‘within-household’ gender wage gap in period t is given by $\exp(\overline{W}_{1t} + v_{1t}) / \exp(\overline{W}_{2t} + v_{2t})$.

Pareto weight Let $m_t = \{\mathbf{x}_1, \mathbf{x}_2, w_{1t}, w_{2t}, \mathbf{d}_{1t}, \mathbf{d}_{2t}, \mathbf{z}_{1t}, \mathbf{z}_{2t}\}$ be the set of variables that enter the intra-family bargaining power (the Pareto weight). As this must be bounded in the unit interval, I employ the logistic function to represent it. Let partner 1’s weight be given by

$$\mu_{1t} = \frac{\exp(\eta(m_t))}{1 + \exp(\eta(m_t))} \quad (1.11)$$

whereas partner 2’s weight by $\mu_{2t} = 1 - \mu_{1t}$. For $\eta(m_t)$ I specify

$$\eta(m_t) = \eta^{(0)} + \sum_n \eta^{(n)} \times \frac{w_{1t}}{w_{2t}} \times \mathbf{1}[N_t = n].$$

Although not explicitly shown to economize on notation, $\eta^{(n)}$ is a function of spouses’ education; I specify $\eta^{(n)} = \sum_e \sum_{e'} \eta^{n,e,e'} \mathbf{1}[educ_1 = e, educ_2 = e']$ where the sum is over a set of mutually exclusive education dummies. $\eta^{(n)}$ reflects how intra-family bargaining power changes with the gender wage gap *within* a particular family and it varies with family composition and spousal education. As education $educ_j$ likely affects the initial allocation of bargaining power in the household, $\mathbf{x}_j = \{educ_j\}$ and there is a partial overlap between \mathbf{x}_j and \mathbf{z}_{jt} . In this specification \mathbf{d}_{1t} and \mathbf{d}_{2t} are left empty.

Stochastic fertility The arrival of children is stochastic and exogenously set to reproduce patterns in the PSID over the life-cycle. Children can affect individual choices in the family through: (i) their needs (they require more of the public good in the form of an equivalence scale $s(N_t)$), (ii) their direct impact on their parents’ time-use preferences $\kappa_2^{(n)}$, (iii) their direct impact on the budget constraint (children require child care if they are young and both parents work away from home), and (iv) their effect on the allocation of bargaining power between parents through $\eta^{(n)}$. To avoid increasing the state space beyond what is computationally feasible, I assume that only the age of

the youngest child (if any) matters for family choices, not the number of children in the household. The idea is that the family will always have to cater for the needs and costs of the youngest child regardless the age or number of older ones.

I assume there are 4 possible family composition (fertility) states in year t , summarized by the state variable N_t . State $N_t = 1$ corresponds to a family with no children under 18 years, $N_t = 2$ indicates a family whose youngest child is between $(0, 5)$ years old, $N_t = 3$ indicates a family whose youngest child is between $[5, 10)$ years old, and $N_t = 4$ is when the youngest child is between $[10, 18]$ years. At age 18 any child leaves the household with certainty. The marginal distribution of children, estimated in the PSID, depends on age and parental education.

The transition between fertility states depends on age, parental education, as well as the fertility state one period before. For a childless family at age t the probability that they have a child at age $t + 1$ is given by

$$\text{Prob}_{t+1}(N_{t+1} = 2 \mid N_t = 1, \text{educ}_1, \text{educ}_2).$$

I restrict the transition matrix to allow smooth transitions only: a family with $N_t = 1$ (no children) may next year have $N_{t+1} = 1$ again or progress to $N_{t+1} = 2$ (a child at the youngest age bracket), but not $N_{t+1} = 3$. Downwards transitions are not allowed with the exception of the arrival of a new child when an older one already exists (in this case I reinstate N to 2) or the departure of an older child from the household. These restrictions accord well with the patterns seen in the data.

Child care costs The function $CC_t(h_{2t}, N_t) = cc_t(h_{2t}, N_t) \times \text{Prob}_t(\text{costs} > 0 \mid N_t)$ describes child care costs a family must meet as a function of its composition and the hours the mother is away from home due to market work (recall that the father always works full-time). The function depends on time to reflect changing prices of child care over time as well as on the probability the family actually faces positive child care costs conditional on the age of

their child. I assume that pre-school children need child care for so long as the mother is away from home working. If she is present in the household for some time, then child care costs are 0 for that time. Young school-age children require some child care only following the schooldays as education is publicly provided whereas older school-age children do not require child care. To account for the fact that some families may have informal child care arrangements in place (such as a grandparent looking after a child) I multiply the costs function $cc_t(h_{2t}, N_t)$ by the probability that the family faces positive such costs. I allow the probability to depend on calendar time and the fertility state N_t and I estimate it directly in the PSID data.

Given that the mother can work either ‘full time’ (FT) or ‘part time’ (PT), the costs function $cc_t(h_{2t}, N_t)$ can be summarized by

$$cc_t(h_{2t}, N_t) = \begin{cases} FT \times cchrates_t & \text{if } N_t = 2 \text{ and } h_{2t} = FT \\ PT \times cchrates_t & \text{if } \begin{cases} N_t = 2 \text{ and } h_{2t} = PT \\ N_t = 3 \text{ and } h_{2t} = FT \end{cases} \\ 0 & \text{in all other cases.} \end{cases}$$

The hourly price of child care is $cchrates_t$ and varies with calendar time. Section 1.5.1 provides details on the estimation of $cchrates_t$ and the probability of positive child care costs.

1.3.4 Retirement

Retirement starts at time $T + 1$ and ends at time T^R for both spouses. During this period the individuals make no time allocation decisions: they are out of the labor force retired and do not engage in home production (thus their ‘productive’ time is entirely spent on leisure). They face no uncertainty regarding wages and productivity (as they earn no wages) or fertility (their children, if any, have grown up and left the household). They receive a pension income which, along with their savings (if any), they use to purchase market goods or save further. In the absence of wages or children their outside options remain

constant throughout retirement; intra-family bargaining power is fixed at its value in the last period of working life.^{27,28}

Retirement in this model serves as a stylized state towards the end of the partners' lifetime. It is not used or needed to infer behavior during the working period of life (which is the focus of this chapter). In the absence of retirement, however, individuals would probably need to accumulate fewer assets during the working period of life and, possibly, work less. In this case, the model would generate full- or part-time employment profiles less easily without pushing the disutility of work towards zero (in the context of the parametrization in (1.8)).

Adopting a compact formulation equivalent to expression (1.1'), the household can be seen during retirement as solving

$$\max_{\{Q_t, A_{t+1}\}_{t=T+1}^{t=T^R}} \sum_{t=T+1}^{T^R} \beta^t \left[\mu_{1T} U_1(Q_t, l_{1t}; \mathbf{z}_{1t}) + \mu_{2T} U_2(Q_t, l_{2t}; \mathbf{z}_{2t}) \right] \quad (1.1^R)$$

subject to

$$A_t + \sum_{j=1}^2 I_{jt} = Q_t + \frac{A_{t+1}}{1+r} \quad A_{t+1} \geq \underline{A}_{t+1} \quad (1.12)$$

$$l_{jt} = \mathcal{T} \quad j = \{1, 2\} \quad (1.13)$$

$$A_{T^R} = 0. \quad (1.14)$$

Most of the notation has been introduced previously. Preferences U_j are given by (1.8) and the Pareto weight μ_{jT} by (1.11); the vector of observable taste shifters is now $\mathbf{z}_{jt} = (N_t = 1, educ_j)'$ as there is no fertility. The budget constraint is slightly different from (1.3) in that earnings are replaced by pension income I_{jt} , the public good Q directly enters the constraint (as there is no

²⁷Allowing the outside options to depend on distribution factors implies that reallocations of bargaining power are in principle possible during retirement too. In practice I am making no use of distribution factors in this chapter, other than the within-household gender wage gap, and thus I fix the retirees' intra-family bargaining power to its last value prior to retirement.

²⁸In the absence of time-use choices during retirement (or, more generally, strictly private goods), the retirees' Pareto weights play in reality no role and their problem collapses to a unitary 'cake-eating' problem.

home production), and there are no child care costs. For each individual, pension income is set to a deterministic 75% of their average full-time earnings at the start of their working life. This figure only serves as an approximation to the actual pension income one would expect to receive; it is not stochastic, it is history independent (it ignores whether one has worked full-time throughout their lifetime or not at all), and it is not adjusted for wage growth that has occurred over time.

Recursive formulation Let $\mathcal{S}_t^R = \{\mathbf{z}_{1t}, \mathbf{z}_{2t}, \mathbf{x}_1, \mathbf{x}_2, w_{1T}, w_{2T}, \mathbf{d}_{1T}, \mathbf{d}_{2T}\}$ be the set of exogenous state variables at time t of the retirement stage; current period assets A_t is the endogenous state variable.²⁹ Q_t and A_{t+1} are the choice variables. Program (1.1^R) can be written recursively as

$$V_t^R(A_t, \mathcal{S}_t^R) = \max_{Q_t, A_{t+1}} \left\{ \sum_{j=1}^2 \mu_j(\mathbf{x}_1, \mathbf{x}_2, w_{1T}, w_{2T}, \mathbf{d}_{1T}, \mathbf{d}_{2T}, \mathbf{z}_{1T}, \mathbf{z}_{2T}) U_j(Q_t, l_{jt}; \mathbf{z}_{jt}) + \beta V_{t+1}^R(A_{t+1}, \mathcal{S}_{t+1}^R) \right\}$$

subject to the budget constraint (1.12), the time budget (1.13), and terminal condition (1.14). This is essentially a modified ‘cake-eating’ problem: in each period the partners maximize a fixed household welfare function deciding, without uncertainty, about current-period expenditure and savings.³⁰

Transition from working to retirement period At time T , the last period of working life, the household’s problem can be written recursively as

$$V_T(A_T, \mathcal{S}_T) = \max_{C_T, A_{T+1}} \left\{ \mathcal{U}(C_T, \mathcal{S}_T) + \beta V_{T+1}^R(A_{T+1}, \mathcal{S}_{T+1}^R) \right\}$$

subject to the constraints of the working period.

²⁹The variables \mathbf{x}_1 , \mathbf{x}_2 , w_{1T} , w_{2T} , \mathbf{d}_{1T} , and \mathbf{d}_{2T} enter the retirement state space through their effect on the Pareto weight in the last period of working life. These are included here for theoretical completeness as, in reality, the retirees’ problem is invariant to the Pareto weight; see footnote 28.

³⁰Fernández and Wong (2014) also model a ‘cake-eating’ and non-stochastic retirement period.

1.4 Model Solution and Simulation

In this section I describe the steps I take to solve and simulate the model developed in section 2.2. The finite horizon life-cycle model requires computation of the solution as a function of the entire state space, including age/time, as described at the start of section 1.3.3. A time period is taken to be 1 year.³¹

I solve the model starting at the end of the retirement period of life, assuming that households exhaust their assets and die without debts, and I move recursively backwards until the beginning of the working period. The solution in the retirement period is straightforward: this part of the model is a ‘cake-eating’ problem without uncertainty that involves the continuous choice of allocating contemporaneous assets and pension income between public consumption and future assets. The solution in the working period is more involved: it entails a mixture of discrete (time allocation) and continuous (consumption/assets) choices under uncertainty which I describe in more detail below.

I discretize the domain of all continuous state variables to reduce the dimensionality of the problem: these are assets A (applies to both the working and retirement periods of life) and idiosyncratic productivity v_1 and v_2 (applies to the working period of life only). I use a grid of 12 points in A , the domain of which depends on age/time, and a grid of 8 points in each of v_1 and v_2 . In generating the grids in v_1 and v_2 , I assume the spouses expect the gender wage gap between them to remain flat (mean stationary) over their life-cycle. This fundamental assumption enables the identification of the household structure and I discuss it further in section 1.5.2. I use information on the variance of each spouse’s log wage net of the variance of the economically-irrelevant transitory shock, the mean of the male’s log wage, and the gender wage gap at the start of the life-cycle. I trim the support of wages 3.25 standard deviations

³¹Currently the chapter focuses on one cohort only (those born between 1943 and 1955). The age of the male spouse is 30 at the start of working life in the model; in the data I match that to people aged 25-37 (mean age 30) in 1980. The age of the male spouse is 59 at the end of working life; in the data I match that to people who are 55-65 years old (mean age 59) in 2009.

above and below their applicable means; the grid points are then the mid-points of equiprobable adjacent intervals covering the applicable support.

The support of the discrete state variables is fully accounted for in the solution. Spouses' education $educ_1$ and $educ_2$ and age/time t are the discrete states in both working and retirement periods of life. Women's unobserved costs of work θ_2 , and the presence of children and the age of the youngest among them N , are additional discrete state variables in the working period.

At any given point of the state space, the solution in the working stage of life proceeds in two steps. In the first step I calculate the optimal consumption/future assets allocation conditional on every possible allocation of spouses' discrete time in the market and the household; these time allocations appear in table 1.2 of section 1.3.3. The second step involves the calculation of the value of the household objective across all possible allocations of time (given the corresponding optimal consumption/future assets allocation from the first step) and the selection of that time allocation that is associated with the highest value. Note that for any given realization of the state space the Pareto weight is known as it is mechanical transformation of the gender wage gap and family composition through (1.11).³² The solution in the retirement stage of life only involves the unconditional calculation of the optimal consumption/future assets allocation as there are no discrete time choices in this stage.

The calculation of the optimal consumption/future assets allocation in the first step requires knowledge of the stream of expected household utilities (weighted sums of spousal utilities) from the following period onwards. This expected future value is a function of today's information, the realization of the state space in the following period, and future assets (a choice variable today). Expectations are taken with respect to three stochastic components in the future period; these are future family composition (presence and age of youngest child) and the spouses' future labor market productivity. The

³²As a consequence, there is no separate grid for the Pareto weight (unlike, for example, Mazzocco et al., 2014).

transition matrices for the random components, *i.e.* the probabilities of moving from one point in today's grid to another grid point tomorrow, are estimated directly from the data given the parametric assumptions of section 1.3.3.

Once the expected future values are calculated, the conditionally optimal consumption/future assets allocation in the first step is obtained by maximizing the weighted sum of the spouses' utilities today and the discounted expected future value. The maximization proceeds in a 'table look-up' fashion where I evaluate the objective in proximate points on the applicable domain of all the relevant choice variables (consumption, future assets). I select the point that produces the maximum value and, using the immediately adjacent points, I generate a new finer grid that is concentrated therein. I reevaluate the objective function and I proceed likewise until I reach the optimal with an acceptable tolerance. This approach guarantees a global maximum if the conditional (on a time allocation) objective function is concave. Although the discrete time use choices in the future can, in principle, induce kinks in the expected future value, I overcome this thanks to sufficient uncertainty about the future state (uncertainty about family composition and labor market productivity).³³ Finally, I use linear interpolations to evaluate the expected future value function outside the asset grid points for which it is explicitly generated.

I simulate 10 replications of the life-cycle choices of 1279 households (a total of 12790 simulations). The simulations are based on initial conditions for spouses' education and family composition observed in the data. I draw initial (log) wages for men and women assuming they are normally distributed around their beginning-of-life means. I replicate the empirical covariance between the two netting out the covariance of measurement error (transitory shock). I produce random draws for the entire profile of permanent shocks and I use (1.10) to generate life-cycle profiles of wages in such a way so as to replicate the empirical profiles of wages and the gender wage gap. I trim the draws

³³I check concavity of the conditional (on a time allocation) objective function by verifying that the second derivative of the expected future value function is globally non-positive with respect to assets.

of such shocks 2.1 times above and below their annual means so as to ensure that the support of simulated wages falls within the wage grids used in the solution of the model.³⁴ I also draw profiles of ‘fertility’ shocks given the initial conditions and the fertility transition matrix estimated in the data. I use the model’s policy functions to infer the optimal choices associated with the random profiles of wages and fertility. This involves the interpolation of the policy functions outside the grid points that are explicitly constructed for. I interpolate linearly over the asset dimension only after selecting the slice of the policy functions that is closer to the simulated wage and fertility state at a given age/time. I start the simulations assuming households hold 0 initial assets.

The above solution and simulation routines are written in Julia.³⁵ With currently one education level only active for either spouse, they run in approximately 40 seconds in total on a 12-core Intel Xeon E5-2630 at a 2.3GHz clock speed.

1.5 Identification and Estimation

In this section I describe the steps I take to estimate the structural model of section 2.2. I follow a two-step procedure. In the first step I estimate the external blocks of the model, namely the wage process for each spouse, the fertility process, and the child care costs. Moreover, using reduced-form information on divorcees, I approximate the gender-specific value of divorce and I use it to normalize intra-family bargaining power in the first few years of the family life-cycle. The role of this approximation and normalization is further discussed below. In the second step, and conditional on results from the first step, I estimate the parameters of the structural model using the method of simulated moments.

³⁴Recall that the wage grids used in the model solution assume a fixed average gender wage gap over the lifetime.

³⁵Julia is a new high-performance programming language; documentation is available at <http://julialang.org>.

Section 1.5.1 discusses identification and estimation of the external blocks of the model; section 1.5.2 discusses identification and estimation of the model's parameters. This part requires information on the initial intra-family bargaining power, a matter that I discuss in section 1.5.3.

1.5.1 External Processes

Wages To construct the wage grids for the solution of the model, I require the mean (log) wage over the life-cycle as well as the wage variance net of the variation in measurement error (transitory shock) for each gender. To integrate out future uncertainty, I require the transition rule for wages, *i.e.* the probability of moving from one point on the wage grid to another, which, in turn, requires knowledge of the covariance matrix of permanent shocks over the life-cycle. To obtain simulated wage profiles, I also need the covariance matrix of spouses' transitory shocks in the first period (used for initial conditions).

The mean and the variance of wages are calculated directly in the data. Results are omitted for brevity but a graphical illustration of the mean appears in figure A.1 in the appendix. Results regarding the transition matrix for wages are also omitted (but available upon request).

Given the parametrization of the wage process in (1.10) the second moments of shocks can be readily identified from various second moments of spouses' contemporaneous, lagged, and lead wages. Meghir and Pistaferri (2004) and previous studies show that $\mathbb{E}[\Delta \ln w_{jt} \sum_{\tau=-1}^1 \Delta \ln w_{jt+\tau}]$ identifies the variance of individual j 's permanent shock at t and $\mathbb{E}[\Delta \ln w_{jt} \Delta \ln w_{jt+1}]$ identifies (minus) the variance of j 's transitory shock. In the first case the sum of consecutive wage growths removes the transitory elements; the remaining covariation between the sum and contemporaneous wage growth is due to the variance of the permanent shock. In the second case the covariation between consecutive wage growths picks up the variance of the mean-reverting transitory shock. Similar moments *between* spouses identify the covariance of spousal shocks.

To obtain estimates of the second moments of shocks I run a minimum

Figure 1.8: Simulated against actual wages (means)

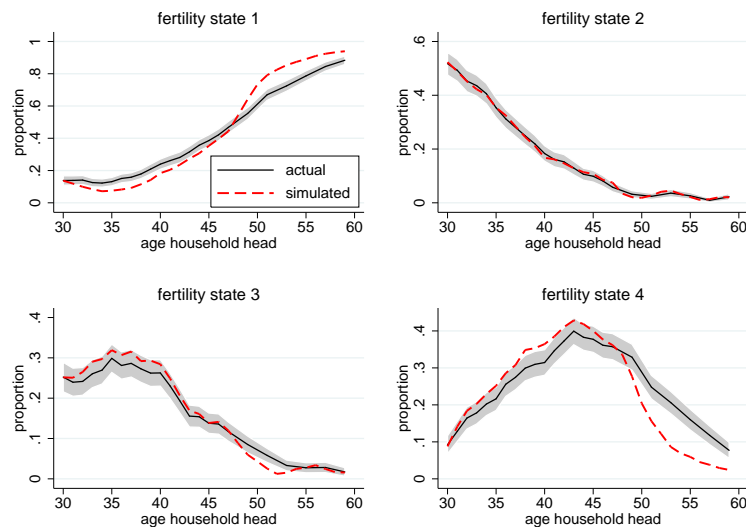
Notes: This figure plots the mean of simulated against actual wages over the life-cycle. Simulated wages are based on wage process (1.10). A 95% confidence interval around the empirical means appears in gray shade.

Source: Major PSID sample and own simulations.

distance estimation matching the empirical covariance matrix of (log) wages to its theoretical counterpart; I allow the second moments to vary over the life-cycle. I use equal weights across all moments (identity weighting matrix). Appendix A.3 reports the estimation details and a full table of estimates. To reduce the effect of wage measurement error in the second stage, I input into the structural model a 5-point two-sided moving average of the covariance matrix of shocks instead of the original point-estimates; figure A.2 in the appendix provides a graphical illustration of the moving average for the variances of men's and women's permanent shocks.

After estimating the wage process, I draw 12790 random profiles for men's and women's wages. Figure 1.8 plots the mean of the simulated wages over time against the empirical ones. This simulation naturally performs well although there are some small discrepancies due to the use of the smoothed second moments of shocks rather than the actual ones.

Fertility To integrate out future uncertainty while solving the model I require the transition rule for fertility, *i.e.* the probability that a family moves

Figure 1.9: Proportion of families in various fertility states

Notes: This figure plots the proportion of families in each fertility state in the actual and the simulated data over the life-cycle. A 95% confidence interval around the empirical means appears in gray shade.

Source: Major PSID sample and own simulations.

from one fertility state to another. These probabilities are obtained directly from the PSID data. I count the number of families reporting a given family composition at time t conditional on their composition at $t - 1$. This is done separately by age t ; the calculation only involves families that are observed in consecutive years and, therefore, uses a subset of the major PSID sample only.

To simulate the model I also require the categorical distribution of family composition at the beginning of the life-cycle. This is taken directly from the data. With this in hand and using the aforementioned transition rule I draw 12790 random life-cycle fertility profiles. I use those as input to the structural model. Figure 1.9 plots the proportion of families in each fertility state in the actual and the simulated data over the life-cycle. Again, this simulation performs well.

Child care costs The hourly rate of child care in function $CC_t(h_{2t}, N_t)$ is $cchrates_t$ and varies with *calendar* time. It is hard to find direct evidence on this. The PSID reports child care expenditure by households but any meaningful analysis of this measure would be incomplete for several reasons.

First, child care expenditure does not necessarily convey information about the price of child care; as an example, child care expenditure may increase for a given household expenditure due to increased demand (say, parents work longer hours, mothers switch from part to full time work etc.) but the hourly price may well have stayed constant. Second, only a fraction of households report positive such costs due to, possibly, one parent being available at home or some other informal child care arrangements. It is not clear how these households compare to the general population and standard selectivity issues arise.

To get around these problems, I exploit the fact that child care is a labor intensive industry and I assume that its sole cost is the hourly wage child care workers are paid. A study that provides a systematic analysis of the wages of such workers is Blau (1992). Based on Current Population Survey data between 1976 and 1986, the study finds that child care workers are paid approximately 50% of the mean wage of all other female workers in that period.³⁶ This number is somehow confirmed by Whitebook et al. (1993) who argue that “child care teaching staff in 1992, as in 1988, continue to earn less than half as much as comparably educated women”.

Given that the PSID data I use cover the years 1980-2009, I adopt the above percentage and I fix *cchr* in year 1981, the mid-year between 1976 and 1986, at \$6.59 (expressed in 2010 dollars).³⁷ A question remains regarding *cchr* prior to and after 1981. There is lack of consistent ‘hard’ evidence on the compensation of child care workers in the longer period. Blau (1992) finds a significant negative trend for wages in one child care sector (with trends in other sectors being insignificant); Whitebook et al. (1993, page 7) report

³⁶Blau (1992) selects a nationally representative sample of around 4,000 child care givers (all of whom are women) and divides them in 3 broad child care sectors (private household care, non-household care, teachers). Table 3 therein reports the average hourly wage in each of the three sectors alongside the average hourly wage of a random sample of other female workers. Based on these numbers I calculate the (weighted) average wage across all child care sectors and divide it by the average wage of other female workers to obtain a ratio of $0.496 \approx 0.5$.

³⁷*cchr*₁₉₈₁ = \$6.59 is 50% of the average female wage in the PSID over 1976-1986.

that a growing segment of the child care workforce has seen a decline in their real wages between 1988 and 1992. O'Neill and O'Connell (2001) report that real wages of child care have been flat or slightly decreasing over the 1977-1997 period. In light of this 'soft' evidence I calibrate *cchr* at a constant \$6.59 (expressed in 2010 dollars) throughout the 1980-2009 period (this period coincides with the life-cycle of the cohort the chapter focuses on). Whenever this rate is below the real federal minimum wage, I update *cchr* to reflect this.³⁸ Essentially this pattern implies that the hourly wage of child care workers declines relative to that of the general population (of both men and women) reflecting -what seems to be- a consensus that child care has steadily become less expensive in the last 3 decades.

Finally, I calculate the probability of a family facing positive child care costs by counting the number of families of a given fertility status that report non-zero such costs (the PSID collects information on child care expenditure after 1988). This is done separately by *calendar* time. In years when child care expenditures are missing I use the probabilities estimated in the closest period when data are available. Table A.2 in the appendix reports the estimated probabilities as well as the calibrated hourly price of child care over calendar time.

1.5.2 Model Parameters

The model is estimated by the method of simulated moments conditional on first-stage results for wages, fertility, child care costs, and a normalization of intra-family bargaining power (the Pareto weight) in the first 10 years of the family life-cycle. The details of this normalization are presented in section 1.5.3 after I discuss its role in the present section.

There is a number of parameters in the model that are kept fixed based

³⁸I update the hourly child care rate upwards to reflect a binding federal minimum wage in the following years: 1980-1987, 1992-1993, 1996-1999, and 2008-2009 (average upwards adjustment of \$0.69 and maximum upwards adjustment of \$2.27; all amounts are in \$2010). Historical data for the federal minimum wage rate are available by the US Department of Labor at www.dol.gov/whd/minwage/chart.htm.

Table 1.3: Fixed parameters

Parameter	Value	Reference
r	0.01	Attanasio et al. (2008)
β	0.98	Attanasio et al. (2008)
γ	1.5	Attanasio et al. (2008)
ϕ	0.8	Lise and Yamada (2014)
π	0.5	Lise and Yamada (2014)
φ	0.5	Lise and Yamada (2014)

on estimates available in the literature; these are summarized in table 1.3.

The interest rate r is set at 1% annually which is very close to the interest rate in Attanasio et al. (2008) and Blundell et al. (2013). The discount factor β is set at 0.98 (same as in the aforementioned papers) implying that families are slightly impatient as the corresponding discount rate is higher than the interest rate. The coefficient of relative risk aversion γ is set at 1.5 as for example in Attanasio et al. (2008). In principle γ could be estimated combining consumption data (from the Consumer Expenditure Survey for example) or assets data. Finally I fix 3 parameters pertaining to the technology of home production: the output elasticity of public expenditures ϕ is set at 0.8, the share of men's housework time π at 0.5, and the technology parameter φ at 0.5. These values accord with Lise and Yamada (2014) who estimate a home production function in a collective setting. The production parameters cannot be identified in the current setup because the output of household production is not observed and there are no observable factors that operate exclusively as production shifters.³⁹

There is a total of 26 remaining parameters⁴⁰: 4 pertaining to male pref-

³⁹Cherchye et al. (2012) identify the parameters of home production through variation in exclusive production shifters (number and age of children) that leave individual preferences unchanged. By contrast, in the present chapter the age and number of children affect individual tastes and the Pareto weight. Lise and Yamada (2014) identify the production parameters parametrically using the marginal rates of substitution between leisure, household work and private consumption (all of which they observe in their data). Such marginal rates of substitution cannot be used in the current setup as time is a discrete resource.

⁴⁰There are 26 parameters to be estimated when $educ_j$, the education state variable of each spouse, is suppressed.

erences ($\kappa_1^{(n)}$), 18 pertaining to female preferences ($\kappa_2^{(n)}$) including unobserved work types (θ_2), and 4 coefficients on the gender wage gap from the Pareto weight function ($\eta^{(n)}$). The estimation proceeds as follows. Given a set of parameter values I solve the life-cycle problem and simulate 12790 households. I compute a number of life-cycle moments pertaining to family time allocations in the simulated dataset and I do exactly the same using the actual data. Finally, I calculate a distance metric between the simulated and actual moments and I repeat the process until the metric is minimized. Formally, the estimated parameters $\hat{\Theta}$ solve the minimization problem

$$\hat{\Theta} = \arg \min_{\Theta} (\widetilde{\mathbf{M}}_n - \mathbf{M}_s(\Theta))' \mathbf{V}_n (\widetilde{\mathbf{M}}_n - \mathbf{M}_s(\Theta))$$

where $\widetilde{\mathbf{M}}_n$ is a $k \times 1$ vector of moments over n observations from the real data, $\mathbf{M}_s(\Theta)$ is a vector of moments over s simulations from the artificial data, and \mathbf{V}_n is the inverse of the diagonal of the covariance matrix of the data moments.⁴¹ For the optimization I use `NLopt` (Johnson) and a number of algorithms therein;⁴² I compute asymptotic standard errors as in Gourieroux et al. (1993) and Adda and Cooper (2003). In total I fit $k = 72$ moments; these are proportions of married men and women engaging in various time allocations by their family composition and period in their life-cycle. Specifically I fit the proportion of

- men in:
 - ‘low middle’ household work
- women in
 - ‘full time’ market work *and* ‘maximum’ household work

⁴¹I use the inverse *diagonal* of the covariance matrix of the data moments in the light of evidence about small-sample biases that arise when using the optimal weighting matrix; see Altonji and Segal (1996).

⁴²I use a local and a global derivative-free algorithm interchangeably. The local is the `Subplex` algorithm implemented by Rowan (1990); the global is a fast controlled random search algorithm described in Kaelo and Ali (2006).

- ‘full time’ market work *and* ‘high middle’ household work
- ‘part time’ market work *and* ‘maximum’ household work
- ‘part time’ market work *and* ‘high middle’ household work
- no market work *and* ‘maximum’ household work.

These moments are calculated separately by family composition (fertility) and over 3 different stages of the working stage of the life-cycle (beginning - first 10 years, middle - subsequent 10 years, end - last 10 years).


Identification works as follows. After normalizing intra-family bargaining power in the *first 10 years* of the family life-cycle, women’s employment rates in the same years identify their preferences over full-time and part-time market work $(k_2^{(n)}, \theta_2)$. Similarly, men’s and women’s rates of household work in the same years identify their respective preferences over home production time $(k_1^{(n)}, k_2^{(n)}, \theta_2)$. These preferences may differ by family composition and variation across fertility states identifies such differences.

Keeping preferences and the initial normalized Pareto weight fixed, I use the model to generate life-cycle profiles of family time allocations. Then, I allow intra-family bargaining power to shift parametrically with the gender wage gap in life-cycle years 11-30, so as to minimize the wedge, if any, between the model-generated and the empirical profiles of time allocations. This identifies $\eta^{(n)}$.⁴³ The identifying assumptions are two: (1) preferences conditional on family composition do not change with time; (2) changes in the gender wage gap over time are unexpected (i.e. entirely treated as shocks). This implies that individuals do not foresee the narrowing of the gender wage gap in the future; the extent to which they do foresee it, their choices should reflect this right from the beginning of their life-cycle and subsequent changes in the gender gap should not induce further behavioral responses.


These points can be made clearer using the graphical illustration in figure 1.10. Suppose the life-cycle consists of three periods only, **period 1**, **period 2**,

⁴³Identification of the structural parameters also obtains if the Pareto weight is fixed, alternatively, in any other year(s) in the family life-cycle, not necessarily the first 10 ones.


Figure 1.10: Graphical illustration of identification




(a) Normalize initial Pareto weight



(b) Identify preferences in period 1



(c) Generate entire profile of time-use



(d) Update Pareto weight

Notes: These four graphs illustrate how I achieve identification of the household structure (individual preferences and changes in intra-family bargaining power). The graphs serve as illustration only and do not correspond to actual profiles from the model or the data.

and **period 3**. Suppose also that the solid black line in figure 1.10 depicts the empirical mean of female full-time market work over the life-cycle. I normalize the Pareto weight in **period 1**, shaded gray in graph 1.10a. Conditional on it, I use the empirical profile in **period 1** to identify women’s preferences for full-time market work. Holding preferences and the initial normalization fixed, I use the model to generate an entire profile of female full-time market work over the life-cycle (red dashed and solid line in graphs 1.10b and 1.10c). In **period 2** and **period 3** I shift the Pareto weight in response to the gender wage gap in order to minimize the distance between model-generated and empirical profiles of female full-time market work (graph 1.10d).

How do I normalize the initial intra-family bargaining power? I discuss this extensively in section 1.5.3. As a brief illustration here, I proxy decision power of a spouse by the ratio of his/her lifetime earnings over the sum of the couple’s lifetime earnings should the spouses get divorce. To construct this ratio I project each spouse’s lifetime earnings in the hypothetical scenario of divorce given their observable characteristics and using information on divorcees from the PSID.⁴⁴ This normalization is appealing because it uses information on divorcees while divorce is the relevant outside option. Alternatively I could have used an arbitrary normalization as in Fernández and Wong (2014).

1.5.3 Initialization of the Pareto Weight

The projection of lifetime earnings if spouses get divorced requires information on the earnings of divorcees. I obtain this information from the minor PSID sample of singles whose discussion I postponed previously in data section 1.2.1. This sample consists of single men and women and mimics many of the selection criteria applied to the major PSID sample of married. I restrict my attention to persons aged 25 to 65 from the core ‘Survey Research Center’ sample between years 1980 and 2009 irrespective of their cohort. I select

⁴⁴Other papers too have used information on divorcees/singles in order to obtain identification in the context of the collective model; for example Browning et al. (2013) use singles to identify a version of equivalence scales and intra-household bargaining power.

Table 1.4: Minor sample descriptive statistics

	Men	Women
% some college	0.50	0.48
Annual earnings	49759	31762
Annual work hours	2133	1826
Num. of kids	0.26	0.94
Observations (person-year)	4561	7614

Notes: ‘some college’ is defined as any education above the 12th grade. Annual earnings are expressed in 2010 dollars.

Source: Minor PSID sample.

individuals who report being divorced, who work in the labor market (as I require information on their earnings), and whose earnings do not fall below 1% or above 99% of the (gender- and time-specific) distribution.⁴⁵ The resulting dataset consists of 4561 divorced male-year and 7614 divorced female-year observations. Some key descriptive statistics are presented in table 1.4 and appendix A.1 discusses this sample at a greater length. A few differences are apparent between married (table 1.1) and divorced individuals: the latter are on average less likely to have been to college, divorced men work and earn less than their married counterparts and divorced women earn roughly the same but work more than their married counterparts. A stark contrast is the number of kids each group has with those continuously married having on average more kids.

My first goal is to form an estimate of the expected flow of lifetime earnings (expected human wealth) of *married* men and women in the major PSID sample *should they get divorced*. Specifically, for each married individual and each time period I want to calculate

$$\text{Human Wealth}_{jt} \text{ if Divorced} = E_t(Y_{jt}^d) + \frac{E_t(Y_{jt+1}^d)}{1 + \rho} + \frac{E_t(Y_{jt+2}^d)}{(1 + \rho)^2} + \dots$$

where Y_{jt}^d is j 's earnings as divorcee at time t . The main difficulty is to estimate

⁴⁵This sample includes those separated alongside those formally divorced.

expected future earnings and for that I use the minor PSID sample of divorcees.

Following Blundell et al. (2015, 2016) on the estimation of expected lifetime earnings, I pool earnings of divorcees, separately by gender, for all years and ages. I regress their earnings Y on two types of individual attributes: attributes that are fixed over time (race and education) and attributes that change with time in a deterministic way (a quadratic polynomial in age and its interactions with race and education). Specifically I regress

$$Y_{kt} = \boldsymbol{\chi}'_k \mathbf{a}_k + \boldsymbol{\psi}'_{kt} \mathbf{b}_k + \epsilon_{kt}, \quad k = \{1, 2\} \quad (1.15)$$

where $\boldsymbol{\chi}_k$ points to the first set of attributes and $\boldsymbol{\psi}_{kt}$ to the second; subscript k reflects the gender of the divorcee. \mathbf{a}_k and \mathbf{b}_k are linear regression parameters. To obtain a time- t estimate of future earnings at $t+s$ should *married* individual j get divorce, I use estimates from (1.15) along with information known at t about the concerned person j . Specifically I construct $\widehat{Y}_{jt+s} = \boldsymbol{\chi}'_j \widehat{\mathbf{a}}_j + \boldsymbol{\psi}'_{jt+s} \widehat{\mathbf{b}}_j$, $j = \{1, 2\}$, and I use it in place of $E_t(Y_{jt+s}^d)$. To generate the sequence of expected future earnings I assume that individuals work until age 65 and that $(1 + \rho)^{-1} = \beta$.

My second goal is to generate a proxy for intra-family bargaining power in the first 10 years of the family life-cycle. The spouses' projected human wealth serves as an approximation to the value of their outside option in the event of divorce. I employ a simple mapping from the approximate value of divorce to bargaining power. Specifically, I define intra-family bargaining power of the *male* spouse at time t (μ_{1t}) as

$$\text{Human Wealth}_{1t} \text{ if Divorced} / (\text{Human Wealth}_{1t} \text{ if Divorced} + \text{Human Wealth}_{2t} \text{ if Divorced})$$

where subscript $_1$ points to the male spouse and subscript $_2$ to the female. This expression is bounded in the unit interval, it is increasing in the value of men's human wealth and decreasing in the value of women's human wealth.⁴⁶

⁴⁶The way intra-family bargaining power is constructed favors mechanically the youngest

Table 1.5: Initialized intra-family bargaining power of men in 1980-1989

	mean	median	std.dev.	min	max	N
Men's barg. power	0.6358	0.6208	0.0673	0.4022	0.9106	5970

Notes: This table reports descriptive statistics for the derived intra-family bargaining power of married men in the major PSID sample in years 1980-1989. 'N' refers to the number of married household-year observations.

The intuition behind these projections is the following. At any given time, the spouses observe perfectly how divorced people of various ages and characteristics fare in life and form expectations about how they would fare, should they get divorced themselves. In other words, the spouses have perfect knowledge of equation (1.15) at any time t and they use it to form expectations about themselves, assuming that the distribution of errors ϵ_{kt} is the same among married and divorced persons of the same gender. Obviously such projections depart from the theoretically-correct value of divorce (*i.e.* the *utility* of the divorcee) but I expect this departure to be innocuous for the initial normalization.


The results from regressions (1.15) appear in table A.3 in the appendix. I normalize intra-family bargaining power in the first 10 years of the family life-cycle so these results use information (earnings, individual attributes) on divorcees from calendar years 1980-1989 (these years coincide with the first 10 years of the married household's life-cycle given the restriction of my estimation sample to one cohort only). Table 1.5 reports basic descriptive statistics for the derived bargaining power of married males in years 1980-1989; figure 1.11 is a graphical representation of its cross-sectional distribution.

I estimate an average $\bar{\mu}_1 = 63.58\%$ in years 1980-1989; the median is slightly lower at 62.08%.⁴⁷ The maximum value is approximately 91% and

spouse in the household (as for such spouse the sequence of earnings forming human wealth would be longer). This can be a problem when the age difference between spouses is large. To remove this undesirable feature I replace human wealth as divorcee by an equivalent annual annuity; specifically I divide 'Human Wealth _{jt} if Divorced' by $\rho^{-1} \times (1 - (1 + \rho)^{-T_j})$ where ρ is the discount rate and T_j is the remaining years of individual j until age 65.

⁴⁷Voena (2015) estimates the male Pareto weight at 0.7 using earlier years of the PSID.

Figure 1.11: Histogram of the initialized intra-family bargaining power of men in 1980-1989



Notes: This figure plots the distribution of the derived intra-family bargaining power of married men in the major PSID sample in 1980-1989. A reference line is plotted at 0.5 (equal bargaining powers between spouses).

there are households in which women are relatively stronger (*i.e.* for which men's bargaining power is less than 50%). The derived bargaining power varies with education; this is consistent with the structural model that postulates that $\mathbf{x}_j = \{educ_j\}$ affects the initial allocation of bargaining power in the household. Appendix table A.4 reports how the derived intra-family bargaining power correlates with each spouse's education, race, and age.

In light of this evidence, I normalize men's intra-family bargaining power in the first 10 years of the life-cycle at 0.6208 (the median value from table 1.5 across all education states). Naturally, the estimates of the structural parameters are conditional on this initial bargaining power and it remains an open question how sensitive they are to a different normalization. One would expect that an alternative normalization would naturally affect the magnitudes of the parameters but likely leave unchanged the gradient of each spouse's parameters with respect to leisure or the sign of coefficients in the Pareto weight function.

1.6 Results

This section presents the estimates of the parameters of the structural model and displays the model's overall fit. Table 1.6 reports the estimates for the parameters of preferences, namely the components of the $g_j(\cdot)$ functions in (1.8). Panel A reports women's preferences for market work and panel B reports men's and women's preferences for household work.

Reading through table 1.6 note that *within* a type of time use, for example within women's 'full-time work' in panel A(I), the parameters corresponding to different family compositions (rows 1-4) are mutually exclusive inside $g_j(\cdot)$. Notice also that the parameters *across* types of female market work in panel A do not increment but they too are mutually exclusive. Finally, note that positive and larger values of these parameters imply that work, in the market or in the household (panel B), induces greater disutility as utility in (1.8) is negative due to γ being set to 1.5. In this case leisure can be seen as a substitute good to consumption.

The parameters of female full-time market work (panel A(I), rows 1-4) turn out positive and with a good spread across fertility regimes. Women with very young kids (up to age 5) suffer the greatest disutility from work whereas childless women the least. This evidence is in line with Blundell et al. (2013) who use a similar preference specification for UK women. The parameters of female part-time market work (panel A(II), rows 1-4) are everywhere lower than those of full-time work implying that the former causes less disutility than the latter. Again, women with young kids (up to age 10) suffer the greatest disutility from part-time work compared to their counterparts with older or no children. Interestingly, whether childless women work full- or part-time in the market makes little difference in terms of the disutility these two types of market work induce.

Rows (5) and (6) report the estimates for women's unobserved random costs of work θ_2 as they materialize in the case of market work. Recall that θ_2 is drawn from a two-point discrete distribution separately for each type of

Table 1.6: Estimates of preference parameters

		A. Women's market work	
		(I) full-time work	(II) part-time work
		value	st.error
(1)	No children	0.174	(0.0023)
(2)	Youngest child: up to 5	0.262	(0.0025)
(3)	Youngest child: 5-10	0.227	(0.0029)
(4)	Youngest child: 10+	0.222	(0.0024)
<i>Unobserved types:</i>			
(5)	Type I: utility cost of work	0.131	(0.0023)
(6)	Type II: utility cost of work	-0.321	(0.0099)
		B. Household work	
		(I) men 'low middle'	(II) women 'maximum'
		value	st.error
(1)	No children	0.0600	(0.0004)
(2)	Youngest child: up to 5	0.0613	(0.0000)
(3)	Youngest child: 5-10	0.0608	(0.0001)
(4)	Youngest child: 10+	0.0620	(0.0007)
<i>Unobserved types:</i>			
(5)	Type I: utility cost of work		0.012
(6)	Type II: utility cost of work		-0.004

Notes: This table reports estimates of preference parameters. Asymptotic standard errors are reported in parentheses.

market work (full-time, part-time) and note that these costs *increment* to the parameters in rows 1-4. Row (5) refers to the ‘low type’ of market work (*i.e.* the type who dislikes work the most) and row (6) refers to the ‘high type’ (*i.e.* the type who favors work). The probability attached to each type is such that the average of θ_2 per type of market work is zero over the population. Taken together with rows 1-4 in panel A(I), these estimates suggest that for approximately 25% of women in the sample function $g_2(\cdot)$ is negative rendering consumption and leisure complements. These are likely to be highly educated women for whom the costs of *not* working are significantly higher than the utility benefits and who are, therefore, highly attached to the labor market.

The parameters of preferences for home production time (panel B, rows 1-4) also turn out positive. For women these estimates are everywhere lower than the parameters of part-time market work implying that ‘maximum’ household work is relatively more attractive to them despite the higher amount of hours it entails (see table 1.2). Comparing men’s and women’s household work preferences it seems at first that men’s disutility from household work is less than women’s due to the lower estimates for men’s household work parameters. However, the amount of time each gender devotes to household work is very different by construction. Women in the model are restricted to devote more than 3 times as much time as men and this is likely to be driving the big wedge between their preferences.

Rows (5) and (6) in panel B report the estimates for women’s unobserved random costs of work θ_2 as they materialize in the case of ‘maximum’ household work. Recall that θ_2 *increments* to the parameters in rows 1-4, panel B(II). There are no unobserved costs for men’s household work because these cannot be identified when men have one binary time choice only. The results suggest that there is not much heterogeneity in women’s household work preferences; indeed shutting down such heterogeneity does not affect the model fit or the other estimates too much.

Table 1.7 reports the coefficients $\eta^{(n)}$ on the gender wage gap from the

Table 1.7: Estimates of the parameters in the Pareto weight

		parameter $\eta^{(n)}$	
		value	st.error
(1)	No children	0.0490	(0.0033)
(2)	Youngest child: up to 5	-0.0171	(0.0069)
(3)	Youngest child: 5-10	0.1143	(0.0187)
(4)	Youngest child: 10+	-0.0538	(0.0046)

Notes: This table reports estimates of the coefficients on the gender wage gap in the Pareto weight function (1.11). Asymptotic standard errors are reported in parentheses.

Pareto weight function (1.11). Note again that, as these are parameters corresponding to different family compositions (rows 1-4), they are mutually exclusive in (1.11). Positive and larger values of these parameters imply that a narrower (lower) gender wage gap in a given period reduces the argument of the logistic function (1.11), lowers men's bargaining power in the household, and increases women's by an equal amount.

It is hard to judge the magnitude of these parameters directly; instead one can look at the implications of the bargaining effects for family time allocations and infer indirectly the importance of those numbers (this is section 1.7). The sign of the parameters is such that the modal with respect to fertility woman sees her bargaining power improve with the narrowing gender wage gap over time.⁴⁸ Two caveats are due here. First, the restriction of the estimation sample to stable households only (*i.e.* to those that do not actually divorce) is likely to push these parameters towards 0 (no effect of gender wage gap on bargaining power). Stable households are those for whom re-allocation should take place less frequently as opposed to the general population.⁴⁹ Second, tar-

⁴⁸Using the estimates in table 1.7 I find that the modal with respect to fertility woman sees her Pareto weight improve by up to 1.3% over her lifetime. For this calculation I use the parametrization of the Pareto weight (1.11) and I assume wages in the household change intertemporally according to the gender wage gap in figure 1.6. The magnitude of this change, unlike its direction, is not of practical interest as it is subject to the initial normalized Pareto weight and the cardinalization of preferences.

⁴⁹The restriction of the estimation sample to stable households is not uncommon in the literature of dynamic household decision making. Lise and Yamada (2014) also select a sample of families that do not experience divorce. This restriction is unavoidable in their

Table 1.8: Proportions of people in different time allocations

	Data	Model
A. Men		
‘low middle’ household work	0.705	0.703
‘low’ household work	0.295	0.297
B. Women		
<i>FT market work and</i>		
‘maximum’ household work	0.295	0.318
‘high middle’ household work	0.352	0.370
<i>PT market work and</i>		
‘maximum’ household work	0.121	0.098
‘high middle’ household work	0.039	0.037
<i>No market work and</i>		
‘maximum’ household work	0.165	0.148
‘high middle’ household work	0.029	0.029

Notes: This table reports the proportion of spouses across different time allocations in the actual and simulated data. The definitions of ‘maximum’, ‘high middle’, ‘low middle’, and ‘low’ household work refer to different amounts of time put into home production; see table 1.2 for details.


Source: Major PSID sample and model simulations.

getting ‘aggregate’ moments of time allocation *separately* for men and women may not convey adequate information to uncover shifts in intra-family power. Possibly, joint moments of time allocation in the household (for example, the proportion of households in which men supply ‘low middle’ hours to home production and women work full-time in the market) may be more appropriate for uncovering shifts in bargaining power (however, I check similar joint moments as a means of over-identification; see below).

The overall model fit is good. Table 1.8 reports selected moments of time allocation in the data alongside their counterparts from the model simulations. These moments are not targeted directly in the estimation but only indirectly as they are weighted means of targeted moments. The full set of targeted moments in the data, alongside their model counterparts, appears in tables A.5-A.6 in the appendix.

paper as they estimate a dynamic collective model using the first order conditions. In this chapter this restriction is not unavoidable thanks to the use of a dynamic programming solution and can be relaxed in future research.

Figure 1.12: Life-cycle model fit



Notes: This figure plots the life-cycle model fit across time allocations of men and women averaged over different family compositions. A 95% confidence interval around the data means appears in gray shade. The definitions of 'maximum', 'high middle', and 'low middle' refer to different amounts of time put into household work; see table 1.2 for details.

Source: Major PSID sample and model simulations.

Figure 1.12 illustrates the life-cycle model fit across different time allocations of men and women (averaged over various family compositions). The most noticeable discrepancy between data and model occurs for women who work full-time in the market and also supply ‘maximum’ hours to the household sector (figure 1.12b). The model overestimates the proportion of women who work full-time in the market during the last few years of their working life. This discrepancy is mirrored for women who do not work in the market but supply ‘maximum’ hours to the household (figure 1.12f). There is currently no mechanism in the model inducing early retirement, such as a compulsory receipt of social security benefits that crowds out labor earnings (French, 2005), and allowing for such a mechanism is likely to rectify this.

Finally, I check a big number of non-targeted joint and dynamic moments of time allocation as a means of over-identification. These are transition probabilities, namely probabilities that an individual engages in a given time allocation conditional on what they or their partner did one or two periods in the past. These appear in figure A.3 in appendix A.4.

1.7 Implications of Model for Behavior

This section discusses the implications of a narrowing gender wage gap for married people’s time allocations by illustrating the various effects (income, substitution, bargaining) the gender wage gap induces on their behavior.

Changes in the gender wage gap (and therefore changes in spousal wages) induce a number of effects on family behavior:

1. A higher female wage is likely to render female labor supply relatively more attractive (especially so for women with young children for whom child care costs may have previously been prohibitive). This is the standard sum of own income and substitution effects operating on labor supply with the latter outperforming the former. On the contrary, a higher male wage may render female labor supply less attractive due to a standard income or added worker effect (Lundberg, 1985).

2. Conditional on labor supply, increasing family wages is likely to increase public expenditure which, depending on the home production technology, may crowd out or boost spouses' household work. Whether this effect is symmetric across spouses or not depends on the nature of complementarity between τ_1 and τ_2 in the production function.⁵⁰
3. Shifts in relative wages in the household can alter the task specialization the spouses engage in. For example, a spouse whose wage increases in relative terms may engage fully in the labor market while her spouse may increase his involvement in home production.
4. Changes in relative wages can alter spouses' value of divorce, shift intra-family bargaining power, and induce bargaining effects across all choices made by them.

In the data, the average within-family gender wage gap, which I feed into the model through budget constraint (1.3) and intra-family bargaining power (1.11), narrows down on average by approximately 10% over the family lifetime (graph 1.6b). This narrowing induces effects on family time allocations that can be categorized in two broad groups: income and substitution effects (corresponding to points 1-3 above) and bargaining effects (corresponding to point 4).

My aim is to separate and quantify the two groups of effects. I proceed as follows. Using the preference estimates from section 2.5 and the observed wage and fertility dynamics over the family life-cycle, I simulate 12790 random households *prohibiting* intra-family bargaining power from shifting with the gender wage gap. In this case the Pareto weight remains fixed at its initial normalized level throughout life. I compare the resulting life-cycle family time allocations to the original ones (the original model fit). Any difference between the two identifies the bargaining effects of a narrowing gender wage

⁵⁰An inspection of the production function (1.9) yields $\frac{d\tau_{jt}}{dK_t} < 0$ and $\frac{d\tau_{2t}}{d\tau_{1t}} < 0$ for $\phi \in (0, 1)$ and $\pi_j > 0$. The inputs to home production are all substitutes.

gap; alternatively it answers *How important are shifts in bargaining power in response to the gender wage gap?*


Subsequently, I simulate a new set of 12790 random households assuming men's and women's wages grow *similarly* over the entire family life-cycle and, therefore, relative wages remain unchanged throughout.⁵¹ I compare the resulting life-cycle family time allocations to the original ones (the original model fit); the difference between the two identifies the sum of income/substitution *and* bargaining effects of the narrowing gender wage gap. Netting out the bargaining effects (identified above) isolates the income and substitution effects.

The results from the first application suggest that the narrowing gender gap induces small bargaining effects on women's labor supply and sizeable effects on men's and women's household work. Figure 1.13 depicts life-cycle profiles of family time allocations when intra-family bargaining power does not respond to the gender wage gap (blue dashed lines through the X's). It superimposes them over the original model-generated profiles (red dashed lines) and the empirical profiles (solid lines). Table 2.8 quantifies the differences: it reports how proportions (averages) of people in various time allocations change when bargaining effects are prohibited, and it does so over various age bands in the life-cycle. The original model's baseline rates (expressed in %) appear in square brackets on the side.

In the first application, women's bargaining power is not allowed to improve alongside the narrowing gender wage gap. This induces *up to 3.06%* more women into working 'maximum' hours in the household and *up to 4.98%* fewer men into supplying 'low middle'. When compared to the baseline rates in the original model, the first figure corresponds to an increase in women's rate by 6.48% and a decrease in men's by 6.95%.⁵² Up to 1.03% more women

⁵¹Specifically I assume that women's wages grow according to men's, relatively slower, wage growth. The within-family gender wage gap remains on average constant at its beginning of life level.

⁵²The drop in men's rate of 'low middle' household hours is mitigated as time goes by possibly due to women supplying more household work (spousal hours are substitutes in home production) or the family affording higher public expenditures (time and public expenditures are substitute inputs in home production).

Figure 1.13: Bargaining effects of a narrowing gender wage gap

Notes: This figure plots life-cycle profiles of family time allocations when intra-family bargaining power does not respond to the gender wage gap (bargaining effects shut; blue dashed line through the X's). The solid line depicts the original data and the red dashed line depicts the original model fit when bargaining effects are allowed. The definitions of 'maximum' or 'low middle' refer to different amounts of time put into home production; see table 1.2 for details.

Source: Major PSID sample and model simulations.

are also induced into full-time market work.

The idea behind these results is the following: keeping intra-family bargaining fixed prevents women from improving their bargaining power as the gender wage gap narrows down in their favor. In that case women would supply more time to both the market and the household (compared to what they actually do), and thus enjoy less leisure. The opposite holds for men.

The results from the second application suggest that keeping the average gender wage gap constant at its beginning-of-life level has, through its income and substitution effects, important implications for women's labor supply and household work but less so for men's household work. Figure 1.14 presents life-

Table 1.9: Bargaining effects: changes in proportions of people in various time allocations

	(1) women full time work		(2) women part time work		(3) women ‘max’ home work		(4) men ‘LM’ home work	
ages 40-49	+0.46	[70.3%]	-0.58	[13.6%]	+2.72	[55.7%]	-4.98	[71.7%]
ages 50-59	+1.03	[79.2%]	-1.40	[10.5%]	+3.06	[47.2%]	-0.47	[70.9%]

Notes: The table reports how proportions of people in different time allocations change when bargaining effects are shut. The original model’s baseline proportions (in %) appear in square brackets on the side. There are no changes at ages 30-39 (first 10 years of life-cycle) because bargaining power in those years does not change anyway (see the normalization of the initial Pareto weight in section 1.5.3). The definitions of ‘maximum’ or ‘low middle’ refer to different amounts of time put into home production; see table 1.2.

Source: Model simulations.

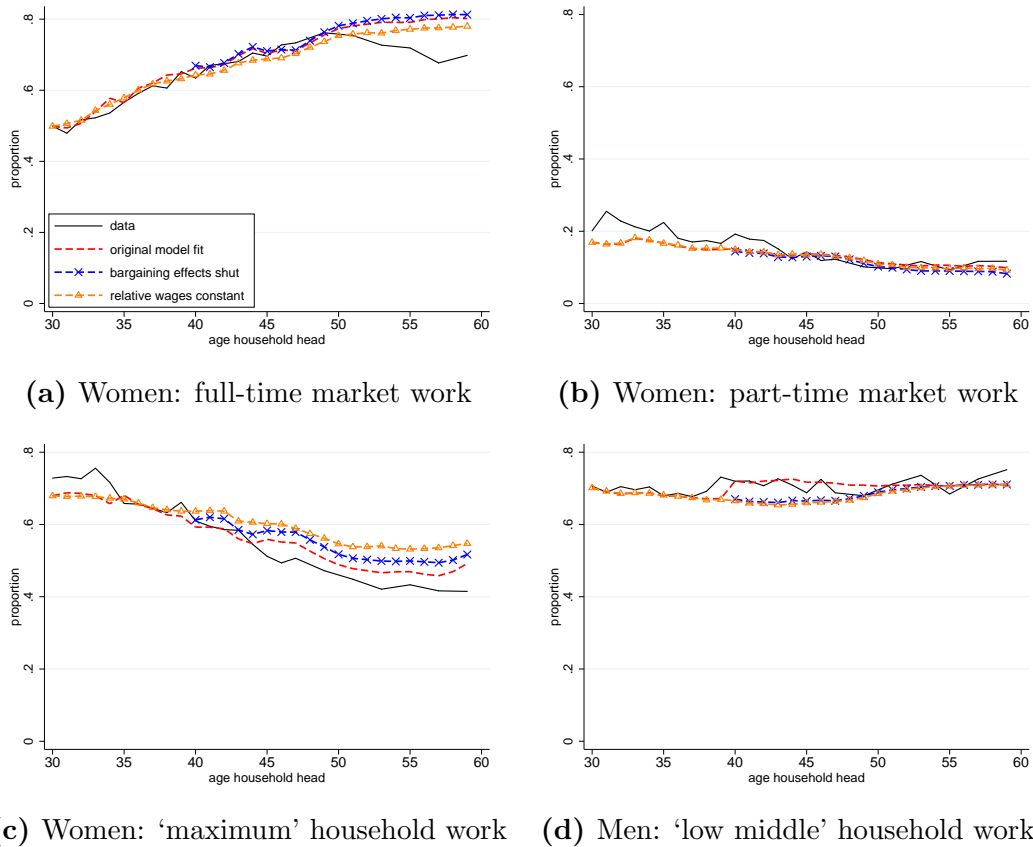
cycle profiles of family time allocations when women’s wages grow according to men’s observed wage growth over time and, as a result, the gender wage gap remains unchanged (orange lines through the triangles). Table 1.10 reports how the proportions (averages) of people in various time allocations change when the gender wage gap remains fixed at its beginning-of-life level. The table reports *exclusively* the sign and magnitude of the income/substitution effects should the gender wage gap remain fixed; model rates (in %) factoring in the bargaining effects appear in square brackets on the side.⁵³

The results suggest that up to 3.39% fewer women would work full-time in the market when the gender gap does not narrow down. Only some of them would work part-time and the majority would not participate in the market at all. ‘Maximum’ household work would attract up to 3.55% more women whereas men’s housework is less responsive: up to 0.52% fewer men would supply ‘low-middle’ housework hours.

The narrowing of the gender wage gap appears important for women’s household work as it accounts for approximately 1/2 of its observed drop over the life-cycle. Half of its effect is due to the higher monetary reward

⁵³The figures presented in the square brackets are the model’s original baseline rates adding the bargaining effects of table 2.8. When the gender wage gap narrows down, it induces income/substitution and bargaining effects. Removing the narrowing of the gender gap, as this application suggests, removes both types of effects. Table 1.10 reports exclusively the former as the latter are already reported in table 2.8.

Figure 1.14: Income/substitution effects of a narrowing gender wage gap



Notes: This figure depicts life-cycle profiles of family time allocations when women’s wages grow according to men’s observed wage growth over time and the gender wage gap within the family remains constant at its initial normalization throughout life (orange lines through the triangles). The solid line depicts the original data and the red dashed line depicts the original model fit. The blue dashed line through the X’s depicts life-cycle profiles when the gender wage gap narrows down but intra-family bargaining power does not respond to it. The definitions of ‘maximum’ or ‘low middle’ refer to different amounts of time put into home production; see table 1.2 for details. *Source:* Major PSID sample and model simulations.

Table 1.10: Income and substitution effects: changes in proportions of people in various time allocations

	(1) women full time work		(2) women part time work		(3) women ‘max’ home work		(4) men ‘LM’ home work	
ages 30-39	-0.19	[57.0%]	+0.16	[16.3%]	+0.19	[66.2%]	-0.03	[68.3%]
ages 40-49	-2.30	[70.8%]	+0.48	[13.0%]	+2.13	[58.4%]	-0.52	[66.7%]
ages 50-59	-3.39	[80.2%]	+0.78	[9.2%]	+3.55	[50.3%]	-0.17	[70.4%]

Notes: The table reports how the proportions (averages) of people in various time allocations change when the average gender wage gap within the family remains fixed at its initial normalized level. The table reports exclusively the changes due to the income and substitution effects. The figures presented inside the square brackets are the model’s rates factoring in (adding) the bargaining effects of table 2.8. When the gender wage gap narrows down, it induces income/substitution and bargaining effects. Removing the narrowing of the gender gap, as this application suggests, removes both effects. This table reports exclusively the former after netting out the bargaining effects reported in table 2.8. The definitions of ‘maximum’ or ‘low middle’ refer to different amounts of time put into home production; see table 1.2.

Source: Model simulations.

of women’s market work, thus to women switching to some form of market work. The other half is due to women becoming relatively stronger in the household decision process, and therefore able to enjoy more leisure (illustrated by the bargaining effects above). The narrowing of the gender gap also appears important for women’s (full-time) market work contributing to its rise over the life-cycle. Two opposite forces operate here: the higher monetary reward pushes women’s market work up (dominating force) whereas women’s improved bargaining power pushes work down replacing it with leisure. Finally, as the income women bring in the household rises, the spouses are in a better financial position to replace household chores, such as child care or cleaning, with similar services purchased from the market. In principle this would benefit men too by cutting down their household work (monetary effect). In reality, however, men keep their household work unchanged as is seen in the PSID because their weakening bargaining position counterbalances the monetary effect. As women become more powerful thanks to a narrower gender wage gap, they shift household work from themselves to their husbands.

Note, finally, that prohibiting the gender wage gap from narrowing does not alter the overall shape of the life-cycle profiles of family time allocations

as figure 1.14 illustrates. This suggests that the fertility dynamics and child care costs spouses face over their lifetime are still jointly quite important even after accounting for the gender wage gap.

1.8 Implications of Gender Wage Equality

In this section I investigate the likely implications that establishing gender wage equality ('equal pay') would have for married couples' allocation of time. This is a realistic counterfactual that policy and business leaders pledge to implement.

In the United States the gender wage gap has been criticized on grounds of discrimination against women. During a weekly radio address on April 12, 2014, President Obama called the lack of equal pay between men and women in the same profession and with the same education "an embarrassment".⁵⁴ In the same month, President Obama took executive action requiring federal contractors to publish data on their employees' pay by race and gender whereas earlier, in February 2010, he announced the establishment of a National Equal Pay Enforcement Task. He has also signed into bill the related Paycheck Fairness Act.⁵⁵

Using the model of section 2.2 and the parameter estimates of section 2.5, I investigate the implications of equal pay through three counterfactual experiments. The experiments have a similar 'flavor' in that across all three of them I shift the mean of the distribution of female wages towards the mean of men's wage distribution, leaving the latter unchanged; the timing that this shift occurs within married people's lifetime differs across experiments.⁵⁶

⁵⁴The full speech of President Obama is available at www.whitehouse.gov/blog/2014/04/12/weekly-address-ensuring-equal-pay-equal-work.

⁵⁵Detailed information is available at www.whitehouse.gov/issues/equal-pay. Other countries too have followed, or led, the campaign for gender wage equality. For example, Prime Minister Cameron in the UK declared on July 14, 2015 his intention to "end the gender pay gap in a generation" (www.gov.uk/government/news/prime-minister-my-one-nation-government-will-close-the-gender-pay-gap).

⁵⁶So far I have estimated the model with one only education/schooling level active for each spouse. In a more detailed estimation across multiple schooling levels, the counterfactual experiments would involve equalizing the *mapping* from education to wages between genders.

In the first counterfactual, I make women earn on average the same wage as men over their entire life-cycle. Specifically, I shift women's mean wage up so that $\bar{W}_{2t} = \bar{W}_{1t}$ at all times (where \bar{W}_{1t} is the mean of men's wages at t). In the second counterfactual, men and women start off their working lives with gender-specific wages at their observed average levels at the start of life, that is $\bar{W}_{1t=0}$ and $\bar{W}_{2t=0}$ respectively. Subsequently, female wages grow rapidly and catch up with men's during the *last* 1/3 of their life-cycle (*i.e.* in life-cycle year 21 out of 30). Once they catch up, men's and women's wages grow in parallel. The third experiment is a repetition of the second one but now women catch up with men just after the *first* 1/3 of the life-cycle (*i.e.* in life-cycle year 11 out of 30); thereafter the spouses earn the same on average wage, that is men's, until the end of their working lives. Across all counterfactuals the spouses are faced with the *observed* fertility dynamics and child care costs, and with their estimated gender-specific variation in wages. Table 1.11 summarizes the three counterfactual experiments.

Gender wage equality has important implications for men's and women's time allocations. The results across all experiments are illustrated numerically in table 1.12 and visually in figure 1.15. Figure 1.15 plots life-cycle profiles of time allocations; across all graphs therein the black solid lines depict the real data whereas the red dashed lines depict the original model fit. Counterfactual #1 is depicted by the blue lines through the crosses, counterfactual #2 by the orange lines through the triangles, and counterfactual #3 by the purple lines through the hollow circles.

In a nutshell, equal pay induces women's entry in the labor market even during the child bearing years. It increases their rates of full-time market work by up to 32% in certain years and decreases part-time work, albeit by less. It lowers women's likelihood of supplying 'maximum' hours to home production by 22.1% and increases men's merely by 3.2% in certain years.

The most sizeable effects on married people's time allocations are seen in counterfactual #1 where men's and women's wages are on average equal

Table 1.11: Equal pay counterfactuals

	When do women catch up with men?	Bargaining effects allowed?
#1	equal average wages throughout life-cycle	no; average wage gap constant
#2	in year 21 of 30	yes
#3	in year 11 of 30	yes

throughout the entire life-cycle. Equal pay makes women 18.36 percentage points -or by 32% relative to the baseline- more likely to be in full-time market work at ages 30-39. Compared to a model baseline rate of 57% at those ages,⁵⁷ equal pay implies that up to 75.36% of women would work full-time. The effect is more profound if one looks at specific early ages: at 30, for example, equal pay renders women 21.3 percentage points more likely to work full-time. Equal pay makes women aged 30-39 up to 2.28 percentage points *less* likely to work part-time; therefore the big increase in full-time work comes from women entering the labor market when they previously did not participate. Interestingly, equal pay induces women to enter the market and work full-time even though they would generally be in their child bearing years. Apparently they must value their increased earnings more than the higher costs of child care that their prolonged absence from home would imply.

The proportion of women aged 30-39 supplying ‘maximum’ hours to home production drops 14.55 percentage points (from a baseline of 66.2% to 51.7%), whereas the proportion of men supplying ‘low middle’ increases 2.18 points (from a baseline of 68.2% to 70.4%). The first figure corresponds to a 22% drop in women’s ‘maximum’ household work rate, whereas the second to a mere 3.2% increase in men’s ‘low middle’ rate. Less time is now devoted jointly to home production possibly because the couple substitutes household time with higher public expenditures. The decline in women’s rates of household work

⁵⁷The ‘model baseline rates’ refer to the proportions observed in the originally simulated data after fitting the model to the PSID.

at ages 30-39 is equivalent to up to 7 hours of home work less per week (see figure 1.15e).⁵⁸

The effects of equal pay remain important in later periods of life, albeit less profound; the gender wage gap anyway narrows down in the real data as time goes by and the effects of equal pay become inevitably less noticeable in later periods. Nevertheless, equal pay still implies an average increase of 12.29 (8.64) percentage points in women's likelihood of full-time market work at ages 40-49 (50-59). The proportion of women supplying 'maximum' household work drops on average 7.27 (3.63) percentage points at ages 40-49 (50-59) whereas the proportion of men supplying 'low middle', perhaps surprisingly, drops approximately 2 points at 40-49 and remains unchanged afterwards.

Counterfactual #1 generates a higher and flatter female life-cycle profile of full-time market work compared to the data. If men and women earned on average the same wages, the model predicts that more than 75% of women would work full-time right from the beginning of their working life; that would rise to more than 82% in the middle years as women become less restricted by young children and to approximately 88% in the later years. With equal pay women's profile of full-time market work tends to mimic men's toward the end of life. The life-cycle profile of part-time work is lower and also flatter compared to the data, as also is the profile of women supplying 'maximum' hours to the household. Men are the least responsive to the elimination of the gender wage gap; the life-cycle profile of men supplying 'low middle' hours tends to be lower in the middle years but the changes are a fraction of those of women.

Counterfactuals #2 and #3 both induce similar eventually effects as counterfactual #1 does: higher female full-time market work, lower female part-time market work and household work, higher male household work. There are little differences between all three experiments in the last 10 years of the

⁵⁸An additional calibration is carried out in order to translate *rates* of household work into home work *hours* because the level of household hours is not targeted in the structural estimation. The details of this calibration are omitted for brevity but are available upon request.

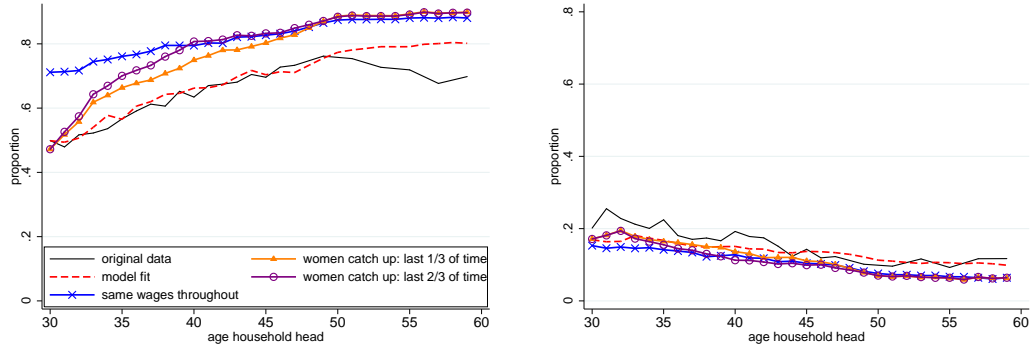
Table 1.12: Counterfactual family time allocations

	(1) women full time work	(2) women part time work	(3) women 'max' home work	(4) men 'LM' home work
<i>Experiment 1: equal average wages throughout life-cycle</i>				
ages 30-39	+18.36 [57.0%]	-2.28 [16.3%]	-14.55 [66.2%]	+2.18 [68.2%]
ages 40-49	+12.29 [70.3%]	-3.00 [13.6%]	-7.27 [55.7%]	-1.99 [71.7%]
ages 50-59	+8.64 [79.2%]	-3.71 [10.5%]	-3.63 [47.2%]	+0.03 [70.9%]
<i>Experiment 2: women catch up with men in year 21</i>				
ages 30-39	+5.69 [57.0%]	+0.48 [16.3%]	-3.08 [66.2%]	+1.16 [68.2%]
ages 40-49	+10.06 [70.3%]	-2.46 [13.6%]	-6.72 [55.7%]	+0.46 [71.7%]
ages 50-59	+10.04 [79.2%]	-4.06 [10.5%]	-7.09 [47.2%]	-0.13 [70.9%]
<i>Experiment 3: women catch up with men in year 11</i>				
ages 30-39	+8.79 [57.0%]	-0.52 [16.3%]	-6.03 [66.2%]	+1.47 [68.2%]
ages 40-49	+12.98 [70.3%]	-3.69 [13.6%]	-9.04 [55.7%]	+0.34 [71.7%]
ages 50-59	+9.91 [79.2%]	-4.03 [10.5%]	-5.51 [47.2%]	-1.71 [70.9%]

Notes: The table reports how baseline proportions of people in various time allocations change across three counterfactual experiments. Baseline proportions from the fitted model appear in square brackets. The definitions of 'max' or 'low middle' refer to different amounts of time put into home production; see table 1.2.


Source: Model simulations.

Figure 1.15: Counterfactual life-cycle profiles of time allocations




(a) Women: full-time market work

(b) Women: part-time market work



(c) Women: 'maximum' household work

(d) Men: 'low middle' household work



(e) Women: continuous household hours

(f) Men: continuous household hours

Notes: This figure illustrates the life-cycle profiles of family time allocations across three counterfactual experiments. The real data appear in the black solid line and the original model fit appears in the red dashed line; experiment #1 is depicted by the blue lines through the crosses, experiment #2 by the orange lines through the triangles, and experiment #3 by the purple lines through the hollow circles.

Source: Model simulations.

life-cycle (when for all three average wages are equal between spouses), but there are some noticeable differences in the earlier years.

Counterfactual #2 generates a steeper profile of female full-time market work over the first 20 years of the life-cycle. As women's wages gradually catch up with men's, the proportion of women in full-time work increases from 62.7% at ages 30-39 to 80.4% at 40-49 and reaches eventually 89.2% at 50-59. These numbers imply an overall 5.7 to 10.1 points increase in full-time market work rates compared to the baseline model (table 1.12, experiment #2, column 1). Experiment #3 generates the steepest full-time work profile predicting a jump from 65.8% at ages 30-39 to 83.3% at 40-49 as women's wages catch up quickly in the first 10 years of the life-cycle. These numbers imply an overall 8.8 to 13 percentage points increase in women's rates of full-time work compared to the original model (tab.1.12, exp.#3, col.1). During the first 20 years the rates of full-time work in experiment #3 are everywhere higher than in #2.

Counterfactual #2 generates a downward profile in part-time market work similar to what we see in the data: the proportion of part-time work decreases from 16.8% at ages 30-39 to 11.1% at 40-49 and 6.4% in the later years. These numbers correspond to a drop of up to 4.1 percentage points in women's rates of part-time work compared to the original model (tab.1.12, exp.#2, col.2). Counterfactual #3 generates the steepest downwards profile in part-time work: the proportion decreases from 15.8% at ages 30-39 to 9.9% at 40-49 and 6.5% in the later years (tab.1.12, exp.#3, col.2). During the first 20 years women's rate of part-time work in experiment #3 is everywhere lower than in #2.

Regarding women's household work, experiment #3 generates the steepest downward profile in the early years from a proportion of 60.2% at ages 30-39 to 46.7% at 40-49 and 41.7% afterwards (tab.1.12, exp.#3, col.3). Counterfactual #2 generates a slightly less steep profile over the first 20 years (actually quite similar to what we see in the data) and the proportion of women supplying 'maximum' hours is almost everywhere higher than in experiment #3. Finally, the proportion of men supplying 'low middle' hours is usually higher

in both experiments than it is in the original model simulations (but with little differences between them).

Overall, equal pay has important implications for married people's time allocations with the most striking changes concentrated around women's entry in the labor market and full-time market work during child rearing. The early effects are greatly mitigated when equal pay is established later on in their lives. A higher proportion of women working full-time and a higher hourly wage result in additional income for their households, a boost to their savings, and a more equal allocation of household work between men and women (even though women would still work more in the household than men). The time mothers spend with children is not modeled in this chapter and it remains an open question how equal pay would impact on this important dimension (Del Boca et al., 2014).

1.9 Discussion and Conclusions

Over the last three decades the wage gap between men and women in the United States narrowed down dramatically. This chapter asks how this closing of the gender wage gap affects family time allocations across market work, household work, and leisure. I hypothesize that the narrowing has a direct monetary effect (by increasing women's monetary reward for market work) as well as an indirect one through shifting bargaining power in the household. Currently the chapter investigates this hypothesis on one only cohort whose working life spans the period 1980-2009 when the closing of the gender gap had been the most profound. The data (PSID) reveal that both mechanisms (monetary reward and changes in bargaining power) are important.

I develop a life-cycle collective model of public consumption, savings and time allocation for spouses who differ in preferences but share a common budget constraint. In the model, spouses' hourly wages affect choices through entering (i) the intertemporal budget constraint and (ii) intra-family bargaining power. The latter is so because wages shift the utility spouses can enjoy

outside their household in the event of divorce. To estimate the model I use cross-sectional variation in wages and family composition as well as the narrowing of the gender wage gap since 1980 which I treat entirely as shock. I utilize data on married as well as divorced individuals.

The empirical life-cycle profiles of family time allocations are reproduced closely. To achieve this, the model assigns women a higher intra-family bargaining power as the gender wage gap narrows down in their favor. The improvement in women's bargaining power affects primarily spousal time into home production shifting household work from women to their husbands. A closing of the gap by 10% since 1980 decreased women's household work by approximately 14% and increased their rate of full time market work by 4%. Half of the decrease in women's household work is due to the higher monetary reward of market work, thus to women switching to some form of market work. The other half is due to women becoming relatively stronger in the household decision process, and therefore able to enjoy more leisure.

I use the model to assess the likely implications that gender wage equality would have for family time allocations. If women earn on average their husband's wage, women's rate of full-time work increases dramatically throughout the life-cycle. The increase is more striking in the early child bearing years when the model predicts that 75.36% of women would work full time compared to 57% in the data. This is primarily due to women entering the labor market when they previously did not participate. Equal pay makes the allocation of time into home production more equal between spouses but it also decreases the overall amount of time invested therein.

This study is subject to a number of limitations. The current focus on a single cohort is probably the most serious one as it prohibits separating time and age/life-cycle effects. Economists tend to think of the gender wage gap as evolving *over time*; nevertheless, this chapter investigates how the gender gap affects family time allocations *over a particular life-cycle*. Extending this chapter to multiple cohorts will enable to study how the gender gap ultimately

affects time allocations *over time* and, therefore, investigate its role for the sharp increase in female labor supply and decrease in household work over the last 3 decades. Wages are taken as exogenous and the model abstracts from human capital; this raises ‘reverse causality’ concerns especially in the present context of family time allocations: a narrower gender wage gap may be driving family time allocations but also women’s gradual entry into the labor market may well be driving the narrowing of the gender wage gap. Allowing for human capital presents challenges and jeopardizes identification in the current model where intra-family bargaining power is a function of wages. Solving for the full dynamic problem of divorced persons is likely to overcome this challenge but will also result in detaching the Pareto weight from the gender wage gap. However, an additional advantage of solving for divorcees explicitly is that it allows the characterization of divorce and removes the restriction of *stable* households on the estimation sample.

A number of extensions are desirable and, possibly, feasible. Modeling fertility as an endogenous choice (in the spirit of Francesconi (2002) in the unitary model) is likely to be important for the use of the model in assessing counterfactual policies. Parental time with children is certainly a big component of parents’ time use (Knowles, 2013) but is not modeled herein due to lack of consistent data over time. A better treatment of the price of household appliances, as in Greenwood et al. (2005), or services, beyond the price of child care, is likely to be important for the patterns of household work. Finally, use of consumption data can help identify a number of parameters currently imposed onto the model from the literature and possibly characterize consumption allocations between spouses.

Chapter 2

Consumption Dynamics and Allocation in the Family

2.1 Introduction

The transmission of idiosyncratic wage/income changes into household consumption, or the link between wage/income and consumption inequality, has been the focus of a large body of ongoing research (eg. Blundell and Preston, 1998; Krueger and Perri, 2006; Blundell et al., 2008). A consistent empirical finding is that household consumption is considerably smoothed with respect to wage/income shocks, even with respect to shocks to after-tax income.¹ Recent economic research attempts to understand theoretically and quantify empirically the mechanisms behind household consumption smoothing (eg. Heathcote et al., 2014; Blundell et al., 2016). But, although these studies explore a rich set of possible consumption smoothing mechanisms, such as family labor supply or external insurance, there is little work on *intra*-household consumption inequality or, more generally, on the internal workings of the household.² The goal of the present chapter is precisely this: investigate the mechanisms of consumption smoothing within a collectivity of potentially different individuals that act together under common constraints (that is, within

¹Blundell and Etheridge (2010) and Heathcote et al. (2010) study the empirical patterns of inequality in consumption and various measures of income since the 1980s.

²Lise and Seitz (2011) is an exception; I discuss below how this chapter relates to it.

a household) and thus respect the fundamental principle of methodological individualism (Manser and Brown, 1980; McElroy and Horney, 1981; Chiappori, 1988). Treating the household as a collectivity of individuals opens up the discussion to new issues, such as how resources are shared within the household or how intra-household bargaining power influences consumption smoothing.

This chapter develops a collective life-cycle model for consumption, savings, and labor supply of a household of two decision-making spouses.³ The model takes into account a number of mechanisms that potentially come into play at the intersection of wages, earnings, and consumption. First, it allows for self-insurance through borrowing or saving in the credit markets. Second, it gives each spouse an endogenous labor supply choice at the intensive margin. Third, it gives each spouse an endogenous consumption choice over multiple consumption goods, namely over a private (or rival) good and a public (or non-rival) good. Fourth, it assigns intra-family bargaining (or decision) powers to the decision-making spouses conditional on which all other choices are made. Such bargaining powers reflect individual and household characteristics, as well as market conditions, and can change with them across as well as within households.

The life-cycle model has four important features. First, the model is collective in the spirit of the static implementations of Chiappori (1988, 1992) and, given the presence of a public good, of Blundell et al. (2005). The decision-making spouses have their own, egoistic, gender-specific preferences over their leisure and private consumption, and the household public good; and they devise a mechanism, summarized by their intra-family bargaining powers, to reach agreements between them. Second, each spouse's preferences are potentially nonseparable across the three goods they draw utility from. Such non-separability permits a rich pattern of complementarities among goods at the individual and household level. Third, each spouse/earner receives an

³Throughout this chapter, I use the terms 'household' and 'family' interchangeably to refer to a two-member collectivity of financially dependent individuals. Whenever I need to make reference to single-member households, I clearly state so. I also use interchangeably the terms 'spouses', 'partners', 'individuals'.

idiosyncratic hourly market wage, the stochastic component of which is decomposed into a spouse-specific permanent and transitory shock (Abowd and Card, 1989; MaCurdy, 1982; Meghir and Pistaferri, 2004). These shocks are correlated between spouses reflecting positive or negative assortative mating (Blundell et al., 2016). Market wages are assumed exogenous to individual choices (albeit functions of age, gender, and other observables) and serve as the primitive source of uncertainty conditional on which the spouses make joint decisions. Fourth, the model features lack of spousal commitment to future allocations of resources in the spirit of Mazzocco (2007) and Lise and Yamada (2014). Formally, in each period the spouses stay together as members of a common household, they do so because they satisfy, separately, their participation constraints in the household. Such participation constraints take the form of lower bounds, often micro-founded as the value of the outside option/divorce (Voena, 2015), that the utility each spouse enjoys from within their joint household must respect.⁴ These lower bounds may move with observable characteristics or market conditions triggering shifts in intra-family bargaining power. In this model, market wages enter the outside options and, therefore, affect intra-family bargaining power. Intra-family bargaining power may increase reflecting an improvement in one's market wage, although the direction of this effect is a priori unrestricted. Lack of commitment nests the full commitment benchmark, where intra-family bargaining power remains fixed over a given family's life-cycle, and allows tests of one setup against the other.

I derive closed-form analytical expressions for spousal earnings (labor supply) and consumption as functions of wage shocks, relying on Taylor approximations to the problem's first order conditions and the intertemporal budget constraint. Similar approximations within the unitary context appear in Blundell and Preston (1998) and a sequence of papers thereafter (Attanasio et al., 2002; Blundell et al., 2008, 2015, 2016). These expressions are convenient be-

⁴Non-cooperation in the household may be seen as an alternative outside option available to spouses. Lechene and Preston (2011) and Cherchye et al. (2015) are recent studies of static noncooperative household models.

cause they provide a neat picture for the contribution of various components (preferences, intra-family bargaining power, consumption sharing, etc.) to the transmission of wage shocks into earnings and consumption without resorting to any specific parametrization of individual preferences. To the best of my knowledge the present chapter is the first one to implement such approximations in the context of the collective model.

Identification of the household structure is challenging for a number of reasons that are either inherent to the household model or relate to the availability of appropriate data. The first reason pertains to the non-observability of two outcome variables, namely of *individual* private consumption of each spouse. Unlike Cherchye et al. (2012) for the Netherlands and Lise and Yamada (2014) for Japan, there is no household survey in the US with comprehensive information on assignable or exclusive expenditure within the household. This is why, throughout this study, I assume one observes *total* private consumption in the household, but not its sharing between spouses. A second reason pertains to the rich complementarities allowed for by the non-separabilities in spousal preferences. The demand for any given good is a function of the quantity demanded for every other good separately, and, in the analog of a demand analysis, the demand system out of this household problem exhibits no exclusion restrictions. A third reason, inherent to the collective model, pertains to the indistinguishability of preferences from intra-family bargaining power. The literature has often resorted to distribution factors, that is variables that affect bargaining power but not preferences, to overcome this problem (Chiappori et al., 2002; Voena, 2015). This problem is only aggravated here due to wages entering the budget set *and* intra-family bargaining powers. A final reason, not uncommon across applied consumer economics, is the lack of observable cross-sectional variation in the price of consumption (in here, in the prices of two consumption goods).⁵

With the aim to identify the household structure, that is preferences and

⁵A general overview of identification in the static collective model is given in Chiappori and Ekeland (2009).

the parameters of consumption sharing and intra-family bargaining power, I use information on single individuals of either gender before they form a joint household or after they separate from it. This approach has been used in Barmby and Smith (2001), Vermeulen (2006), Vermeulen et al. (2006), or, more recently, Browning et al. (2013). Assuming that each spouse's preferences are the same as their single counterparts', the full set of gender-specific preferences, the parameters of consumption sharing, and the effects of wages on household outcomes through their impact on intra-family bargaining power are identified. The rationale is the following: observability of singles suffices for recovering individual preferences over leisure and all types of consumption. When a given wage shock hits, and conditional on those preferences, the unobserved response of each spouse's private and public consumption *should they be single* becomes known. The observed response of consumption at the household level is then a weighted sum of the two individual responses; the weight is nothing but the consumption sharing between spouses in the case of private consumption or (a transformation of) their intra-family bargaining power in the case of public consumption.

The empirical implementation requires panel data on individual earnings and hours of work, household-level consumption, and household-level assets of single- and multi-member households. The Panel Study of Income Dynamics (PSID) is suitable as it collects such data biennially since 1999. I utilize 7 waves in total (1999-2011).

The results suggest sizeable Frisch own-elasticities of labor supply for both men and women, with women's labor supply being considerably more elastic than men's. I also find sizeable Frisch consumption substitution elasticities, with, perhaps surprisingly, the public good (comprising items such as housing services for renters and owners, household utilities, and children's expenditure) exhibiting a much higher elasticity than the private good. This may in part reflect higher variability in public expenditure due to house prices. The evidence from the cross-elasticities suggests that market hours are Frisch sub-

stitutes with all types of consumption and statistically significantly so far as public consumption is concerned. Note that this Frisch substitutability concerns the intensive margin of labor supply only. Overall, men seem to differ from women in their labor supply preferences but not in their consumption preferences. Finally, the bargaining effects of wages, that is the effects on household outcomes through the impact of wages on bargaining powers, are economically significant but not statistically so. At first sight, this suggests one cannot reject full commitment between spouses, but the result may mask, among other factors, unobserved heterogeneity across households.

This chapter is related to two distinct strands of literature. A large literature in consumer analysis and labor economics is devoted to understanding household decisions respecting methodological individualism. The most prominent approach, imposing efficiency of household decisions, has been the collective model introduced by Chiappori (1988, 1992) and Apps and Rees (1988). These first static implementations triggered a series of follow-up studies extending the basic model into home production (Chiappori, 1997), income taxation and discrete labor supply (Donni, 2003; Blundell et al., 2007), public goods (Blundell et al., 2005), multiple consumption items (Chiappori, 2011), and numerous other features. Browning and Chiappori (1998), Chiappori and Ekeland (2006, 2009), and Bourguignon et al. (2009) provide a theoretical overview of the collective model's features and implications for household behaviour.

Early empirical implementations of the static model include Bourguignon et al. (1993), Browning et al. (1994), or Chiappori et al. (2002). These studies were mostly concerned with testing the collective model or showcasing the relevance of distribution factors for household behaviour. More recently, Lise and Seitz (2011) use the collective model to infer the evolution of intra-household inequality in Britain. This chapter differs from Lise and Seitz (2011) in that I develop a life-cycle collective model (theirs is static) with the aim to investigate the *mechanisms* of consumption smoothing and inequality transmission

in the household. In a somehow related study, Browning et al. (2013) estimate the allocation of consumption and bargaining power in the household, as well as concepts of economies of scale due to living with a partner (as opposed to staying single) or equivalence scales for comparing the same individual across the two states of life (*i.e.* in a couple and as single).⁶ Cherchye et al. (2012) estimate a collective model with home production while relying on variation in observed production shifters to achieve identification. Lewbel and Pendakur (2008) and Dunbar et al. (2013) estimate the household sharing rule, namely each spouse's share of household expenditure, using restrictions on Engel curves. Finally, Cherchye et al. (2009, 2011, 2015) apply a somehow distinct empirical methodology to the collective model, namely they set-identify the sharing rules from alternative versions of the model relying on revealed preferences principles.

The extension to the dynamics of the collective model opens up the discussion to new issues such as issues pertaining to inter-temporal commitment between spouses. Mazzocco (2007), the first implementation of the dynamic collective model, rejects full commitment on data from the US post 1980. Ever since, the dynamic collective model has been used to study a number of issues, such as household behaviour post the liberalization of divorce and property division laws in the US in the 1970s (Voena, 2015) and education investments of women (Bronson, 2014; Chiappori et al., 2015). Lise and Yamada (2014), who also provide evidence against full commitment, and Mazzocco et al. (2014) extend the dynamic collective model to home production and general preferences.⁷

Another literature, at the intersection of macro- and microeconomics, studies the relation between wage/income and household consumption growth. Recent papers include Blundell and Preston (1998), Attanasio et al. (2002), Krueger and Perri (2006), Blundell et al. (2008), Kaplan and Violante (2010),

⁶Chiappori and Meghir (2014) provide an overview of these concepts.

⁷Browning et al. (2014) and Chiappori and Mazzocco (2014) provide an overview of the static and intertemporal collective models.

Heathcote et al. (2014), and, more recently, Blundell et al. (2016). Most of these papers have a life-cycle focus and investigate the transmission of wage/income changes, expected and unexpected, into household consumption in a variety of alternative environments (stretching from exogenous income and no insurance against permanent shocks to endogenous labor supply and partial insurance). What they all have in common, however, is their modeling of the household as a single economic agent, and therefore, their abstracting from issues pertaining to the internal workings of the household or intertemporal commitment.⁸

The present chapter advances this literature into the realm of the collective model. It contributes by studying the transmission of idiosyncratic wage/income changes into household consumption, or the link between wage/income and consumption inequality, when household members have their own preferences and an efficient decision mechanism, which reflects in part their outside options, is put in place among them.

The chapter proceeds as follows: section 2.2 presents the life-cycle collective model and discusses identification. Section 2.3 presents a parallel model for singles and discusses how the two models can be of use together. Section 2.4 is concerned with the empirical implementation of the model(s) whereas section 2.5 presents the results. Section 2.6 concludes.

2.2 A Life-Cycle Collective Model without Commitment

This section presents the life-cycle collective model without commitment for the family of two decision-making partners. Specifically, section 2.2.1 develops the model, section 2.2.2 overviews the solution to the household problem, section 2.2.3 illustrates the family response to wage shocks, and section 2.2.4 is dedicated to the identification of the household structure.

⁸Jappelli and Pistaferri (2010) and Meghir and Pistaferri (2011) provide an overview of this literature.

The model of this section describes the household *after* the spouses have decided to form one and *before* they possibly decide to dissolve it. Specifically, the model is conditional on the spouses meeting, say, in the marriage market, and deciding to live together in one household. Assortative mating is taken care by conditioning observables and unrestricted correlations between spousal wage shocks. Divorce or, more generally, household dissolution exists as threat, but I do not model spousal choices in case those events actually materialize. Instead I model spousal choices for so long as the spouses find their joint arrangement (their joint household) superior to their outside options, *i.e.* for so long as they do not exercise the threat of divorce or household dissolution.

2.2.1 The Model

Formally, let H_{jit} be the continuous hours of market work of spouse j in family i at time t ; to fix ideas suppose a male-female couple where $j = 1$ denotes the male spouse ('the husband') and $j = 2$ the female spouse ('the wife'). I normalize the total time endowment available to each spouse to 1 so that $H_{jit} \in (0, 1]$ (time can only be allocated between market work and leisure). I abstract from modeling labor supply at the extensive margin too, as in Blundell et al. (2007), because this would render the preferred model solution method, namely the Taylor approximations to the first-order conditions and the intertemporal budget constraint, inapplicable.

Let C_{it} be *total private* consumption of the household at t and K_{it} be its *public* consumption; both C_{it} and K_{it} are composite Hicksian goods. C_{it} is the sum of C_{1it} and C_{2it} where C_{jit} is spouse j 's *individual private* consumption.

A good is private if consumption of one unit by one spouse deprives the other spouse from that particular unit. Consider a meal at the local coffee shop or a subway ticket; both goods are private in that their consumption (or use) by one spouse renders the goods unavailable to the other spouse.

On the contrary, a good is public if consumption by one spouse does not reduce the amount available to the other spouse. Consider a car journey to the countryside, or the gas required to heat the house in the winter; if a spouse

consumes (purchases) a unit of these goods, the exact same unit is always available for the other spouse to enjoy. In a seminal paper, Blundell et al. (2005) consider children's expenditure in the household as a public good for their parents.⁹ The existence of public goods generates economies of scale in the family as the spouses share the expenditure on them, an expenditure which they would otherwise have to bear by themselves alone.

Throughout this study I assume the econometrician observes C_{it} separately from K_{it} given the itemized household expenditure data available nowadays. However, the econometrician does not observe the partition of C_{it} into C_{1it} and C_{2it} as consumption information is usually collected at the *household* rather than the *individual* level.¹⁰

Spouse j 's within-period and gender-specific preferences are given by

$$U_j(K_{it}, C_{jit}, 1 - H_{jit}; \mathbf{z}_{jit})$$

where $1 - H_{jit}$ denotes leisure and \mathbf{z}_{jit} an $m_j \times 1$ vector of observable taste shifters such as race, age, or education.¹¹ I do not parameterize U_j but I assume it has continuous second-order derivatives in K , C_j , H_j , and $U_{j,K} > 0$, $U_{j,C_j} > 0$, $U_{j,H_j} < 0$ (first derivatives), $U_{j,KK} < 0$, $U_{j,C_j C_j} < 0$, $U_{j,H_j H_j} > 0$ (second derivatives).

The household solves over the life-cycle

$$\max_{\{K_{it}, C_{jit}, H_{jit}, A_{it+1}\}_{\forall t, j}} \mathbb{E}_0 \sum_{t=0}^T \beta_t U_1(K_{it}, C_{1it}, 1 - H_{1it}; \mathbf{z}_{1it}) \quad (2.1)$$

⁹Even though the couple may have children in this model, children themselves are not decision makers.

¹⁰There exists a limited only number of household surveys with individual consumption information; Cherchye et al. (2012) gather detailed assignable expenditure through the Longitudinal Internet Studies for the Social sciences whereas Lise and Yamada (2014) use the Japanese Panel Survey of Consumption with detailed information on individual expenditures in the household. To the best of my knowledge there is no household survey in the US that provides such information comprehensively.

¹¹I treat leisure as private good. Allowing for publicness in spousal leisure, as in Fong and Zhang (2001), together with publicness of K , will likely jeopardize the model's empirical tractability.

subject to constraints

$$\mathbb{E}_0 \sum_{t=0}^T \beta_t U_2(K_{it}, C_{2it}, 1 - H_{2it}; \mathbf{z}_{2it}) \geq \mathcal{U}_2(\mathbf{X}_{i0}) \quad (2.2)$$

$$U_1(K_{it}, C_{1it}, 1 - H_{1it}; \mathbf{z}_{1it}) \geq \bar{U}_{1t}(W_{1it}, \mathbf{d}_{it}) \quad \forall t \quad (2.3)$$

$$U_2(K_{it}, C_{2it}, 1 - H_{2it}; \mathbf{z}_{2it}) \geq \bar{U}_{2t}(W_{2it}, \mathbf{d}_{it}) \quad \forall t \quad (2.4)$$

$$A_{it} + W_{1it}H_{1it} + W_{2it}H_{2it} = K_{it} + P_t C_{it} + \frac{A_{it+1}}{1+r}, \quad A_{iT+1} = 0, \quad \forall t. \quad (2.5)$$

Constraint (2.2) is a promise keeping constraint, set out at time $t = 0$, that spouse 2's lifetime utility will not fall below a threshold $\mathcal{U}_2(\mathbf{X}_{i0})$. \mathbf{X}_{i0} is a set of covariates known at $t = 0$ (or expected at time $t = 0$) that determine the promise keeping threshold; possible examples are spousal education, their age difference, their occupation or expected lifetime earnings.

Constraints (2.3) and (2.4), repeated over time, are the partners' participation constraints in the family; they ensure that each spouse is always at least as well off in their joint household as they could possibly be if they 'fall back' on their outside option at t , \bar{U}_{jt} . I do not have to specify the context of the outside option explicitly although one can think of \bar{U}_{jt} as the value of divorce (for example see Voena, 2015). I assume that spousal hourly market wages W_{1it} and W_{2it} , as well as distribution factors \mathbf{d}_{it} , affect the value of the outside option.¹²

Constraint (2.5) is the sequential budget constraint; A_{it} is common household assets at t , P_t is the relative price of the private good when the price of the public good is normalized to 1 (I assume P_t exhibits no cross-sectional variation), and r is the non-stochastic and known real interest rate.¹³ For

¹²By distribution factors I refer to any variable that affects the spouses' outside options but not their preferences or the budget set (Bourguignon et al., 2009). Chiappori et al. (2002) and Voena (2015) use the sex ratio in the local marriage market or laws governing divorce and property sharing as distribution factors.

¹³If P_t^k is the price of the public good at t , and P_t^c is the price of the private good, then the relative price of private consumption is defined as $P_t = P_t^c/P_t^k$. In other words, P_t^k is the deflator of all other monetary figures in the model, including assets and wages.

tractability, discounting β_t is common between spouses. The household problem can be decentralized using personal (Lindahl) prices for the public good (Chiappori and Meghir, 2014).

Let $\vartheta_{2it=0}$ be the Lagrange multiplier on (2.2), $\tilde{\vartheta}_{1it}$ be the Lagrange multiplier on (2.3), and $\tilde{\vartheta}_{2it}$ on (2.4). Then the above household problem is equivalent to

$$\max_{\{K_{it}, C_{jit}, H_{jit}, A_{it+1}\}_{\forall t, j}} \mathbb{E}_0 \sum_{t=0}^T \beta_t \left[\mu_{1it} U_1(K_{it}, C_{1it}, 1-H_{1it}; \mathbf{z}_{1it}) + \mu_{2it} U_2(K_{it}, C_{2it}, 1-H_{2it}; \mathbf{z}_{2it}) \right] \quad (2.1')$$

subject to the budget constraint (2.5) only. Here, $\mu_{1it} = 1 + \tilde{\vartheta}_{1it} \beta_t^{-1}$ denotes spouse 1's intra-family bargaining power at t and $\mu_{2it} = \vartheta_{2it=0} + \tilde{\vartheta}_{2it} \beta_t^{-1}$ spouse 2's intra-family bargaining power. I normalize the sum of the powers to 1 so that $\mu_{2it} = 1 - \mu_{1it}$. Therefore, μ_{1it} and μ_{2it} can also be seen as the Pareto weights a social planner attaches to each partner's preferences in the social planning problem (2.1') subject to (2.5).¹⁴

The two formulations of the household problem, (2.1) and (2.1'), are equivalent so long as reservation utilities \bar{U}_{1t} and \bar{U}_{2t} do not depend on endogenous choices (Chiappori and Mazzocco, 2014). Formulation (2.1') is more convenient for empirical purposes as it subsumes a number of constraints into intra-family bargaining powers μ_{1it} and μ_{2it} .

2.2.1.1 Intra-Family Bargaining Powers

Intra-family bargaining powers μ_{jit} are functions of the factors entering the constraints to which their components, $\vartheta_{2it=0}$ and $\tilde{\vartheta}_{jit}$, serve as Lagrange multipliers. As a result I model μ_{1it} as

$$\mu_{1it} = \mu(W_{1it}, W_{2it}, \mathbf{d}_{it}; \mu_{1it-1}, \mathbf{X}_{i0}) \quad (2.6)$$

where the dependence on μ_{1it-1} and \mathbf{X}_{i0} will become clearer below. Starting from a generic $\mu_{jit} = \frac{1}{2}$, spouse j becomes relatively more (less) powerful when

¹⁴This normalization is possible if β_t is appropriately redefined in each period to avoid distorting the intertemporal incentives when $\mu_{1it} + \mu_{2it} \neq \mu_{1is} + \mu_{2is}$ for $t \neq s$.

$\mu_{jit'} > \frac{1}{2}$ ($\mu_{jit'} < \frac{1}{2}$) for a future period t' .

With full commitment between spouses, the partners stick to their initial intra-family bargaining powers throughout their life-cycle. In this case

$$\mu_{1it} = \mu_{1it-1} = \dots = \mu_{1i0}(\mathbf{X}_{i0})$$

and revelation of new information about wages or distribution factors does not affect the intra-family bargaining powers. Notice that the beginning-of-life bargaining power μ_{1i0} still depends on \mathbf{X}_{i0} and is household-specific.

If full commitment is not possible, two possibilities open up. *Limited* commitment, on one hand, implies that the spouses update their bargaining powers whenever any of the participation constraints (2.3) or (2.4) bind as a result of changing wages or distribution factors. In this case

$$\mu_{1it} = \begin{cases} \mu(W_{1it}, W_{2it}, \mathbf{d}_{it}; \mu_{1it-1}, \mathbf{X}_{i0}) & \text{if (2.3) or (2.4) bind} \\ \mu_{1it-1} & \text{otherwise.} \end{cases}$$

The partners reallocate power between them to reward the spouse who receives favorable news, say a sudden promotion and a higher permanent wage. As one's outside option gets more attractive due to the favorable news, meaning an increase in \bar{U}_{jt} in (2.3)-(2.4), one must be rewarded with a higher intra-family bargaining power so as to keep satisfying one's participation constraint and remain in the family. This case is introduced in Kocherlakota (1996) and Ligon et al. (2002) and explored empirically in Mazzocco et al. (2014).

Lack of commitment, on the other hand, implies that the spouses update their bargaining powers at any time their wages or distribution factors shift. This case, originally explored in Mazzocco (2007), is equivalent to a series of static interactions between spouses if other dynamic features in the family, such as savings, are assumed away. Full commitment is nested within lack of commitment *and* limited commitment; however lack of commitment is not nested within limited commitment or vice versa. In this chapter, strictly

speaking, I model *lack* of commitment because I allow intra-family bargaining powers to respond to wages (and distribution factors) at all times; however I cannot distinguish empirically lack of- from limited commitment.

As Browning et al. (2014, section 6.2.2) point out, the outcome of problem (2.1') is ex-post (second-best) efficient under lack or limited commitment; once news regarding wages and distribution factors are revealed and bargaining powers determined, no other allocation of resources, other than the chosen one, can induce Pareto improvements without violating the prevailing budget constraint.

2.2.1.2 Prices

There are two types of prices in the model: one is hourly wages W_{1it} and W_{2it} the spouses earn in the labor market, the other is the relative price of private consumption P_t and the interest rate r . As I cannot exploit cross-sectional variation in the latter prices, a feature that would help considerably in the identification of the household structure and the estimation of the model, the rest of this subsection is devoted exclusively to spousal wages.¹⁵

The primitive source of uncertainty the partners face is the hourly wages they earn. Other things remaining fixed, the partners are hit by unexpected changes in W_{jit} and respond by shifting their choices appropriately.

I adopt a permanent-transitory process for each earner's wage in the household. The log of spouse j 's real hourly wage at t is given by

$$\begin{aligned}\ln W_{jit} &= \mathbf{x}_{jit}^{W'} \boldsymbol{\zeta}_{jt}^W + \ln w_{jit}^P + u_{jit} \\ \ln w_{jit}^P &= \ln w_{jit-1}^P + v_{jit}.\end{aligned}$$

The vector \mathbf{x}_{jit}^W contains observable characteristics known at t such as race, age, or education; $\boldsymbol{\zeta}_{jt}^W$ is the vector of time-varying parameters. The permanent

¹⁵There is variation in P_t over time but this is absorbed by conditioning private consumption on time dummies (see section 2.4.3). Cross-sectional variation in P_t could be obtained by exploiting regional (state or county) variation in prices or by constructing family-specific prices given the basket of goods a household consumes (as in Kiefer, 1984).

component $\ln w_{jit}^P$ follows a unit-root and v_{jit} is the (permanent) shock to this process. Transitory deviations from one's wage profile are captured by the transitory shock u_{jit} . Rewriting I get

$$\Delta w_{jit} = v_{jit} + \Delta u_{jit} \quad (2.7)$$

where $\Delta w_{jit} = \Delta \ln W_{jit} - \Delta(\mathbf{x}_{jit}^{W'} \boldsymbol{\zeta}_{jt}^W)$ captures wage growth net of growth in observable characteristics and Δ is the first difference operator.

Deviations from the deterministic path for wages ($\mathbf{x}_{jit}^{W'} \boldsymbol{\zeta}_{jt}^W$) occur because permanent and transitory shocks hit the partners. A permanent shock shifts the value of one's skills in the market permanently (for example, a promotion shifting wages permanently up); a transitory shock is mean reverting (for example, a short illness affecting productivity). When shocks hit, I assume the spouses can perfectly observe and distinguish between them; moreover they hold no advance information on them.

All shocks have zero cross-sectional means and second moments given by

$$\mathbb{E}(v_{jit}v_{kit+s}) = \begin{cases} \sigma_{v_j,t}^2 & \text{if } j = k, s = 0 \\ \sigma_{v_j v_k,t} & \text{if } j \neq k, s = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbb{E}(u_{jit}u_{kit+s}) = \begin{cases} \sigma_{u_j,t}^2 & \text{if } j = k, s = 0 \\ \sigma_{u_j u_k,t} & \text{if } j \neq k, s = 0 \\ 0 & \text{otherwise} \end{cases}$$

for $j, k \in \{1, 2\}$. Permanent shocks are independent of transitory, $v_{jit} \perp u_{kit+s} \forall j, k, i, t, s$. The second moments are time-varying to reflect, for instance, the possibility that some periods of time are more turbulent than others. Shocks of the same type are correlated between spouses reflecting assortative mating (positive correlation) or risk sharing agreements (negative correlation). Finally, shocks are serially uncorrelated.

Although the permanent-transitory process for wages may seem restric-

tive, it fits the PSID data well (for example Blundell et al., 2016). Recent generalizations of this process or alternative formulations include Blundell et al. (2015), Guvenen (2007), and Browning et al. (2010).

2.2.2 The Solution to the Family Problem

Without parameterizing or restricting spousal preferences in (2.1'), I obtain analytical closed-form expressions linking the growth rates of earnings (labor supply) and consumption to wage shocks that hit the household. To achieve this, I apply first- and second-order Taylor approximations to the problem's first order conditions and the intertemporal budget constraint following Blundell and Preston (1998) and a sequence of papers thereafter (Attanasio et al., 2002; Blundell et al., 2008, 2015, 2016).

Specifically, I derive the necessary first order conditions of the household problem (2.1'), subject to (2.5), assuming an interior solution. These are given by

$$\begin{aligned}
[K_{it}] : \quad & \mu_{1it}U_{1,K} + \mu_{2it}U_{2,K} = \lambda_{it} \\
[C_{1it}] : \quad & \mu_{1it}U_{1,C_1} = \lambda_{it}P_t \\
[C_{2it}] : \quad & \mu_{2it}U_{2,C_2} = \lambda_{it}P_t \\
[H_{1it}] : \quad & \mu_{1it}(-U_{1,H_1}) = \lambda_{it}W_{1it} \\
[H_{2it}] : \quad & \mu_{2it}(-U_{2,H_2}) = \lambda_{it}W_{2it} \\
[A_{it+1}] : \quad & \lambda_{it} = \beta(1+r)\mathbb{E}_t\lambda_{it+1}
\end{aligned} \tag{2.8}$$

where λ_{it} is the Lagrange multiplier on the sequential budget constraint at t (the marginal utility of wealth) and μ_{jit} , $j = \{1, 2\}$, is given by (2.6). U_{j,x_j} , the first order partial derivative of U_j with respect to $x_j = \{K, C_j, H_j\}$, $j = \{1, 2\}$, is i - and t -specific but I omit these subscripts to ease the notation.

There are five *intra*-temporal and one *inter*-temporal optimality conditions in each period t . With unrestricted spousal preferences the first order conditions for C_{jit} and H_{jit} are functions of K_{it} , C_{jit} , and H_{jit} , whereas the first-order condition for K_{it} is a function of K_{it} , C_{1it} , C_{2it} , H_{1it} , and H_{2it} . The

latter is so because K_{it} bridges partners' preferences and is therefore interrelated with all other choice variables.

Applying Taylor approximations to the first order conditions (more on this below) enables me to relate the growth rates of labor supply and consumption to wage shocks *and* an innovation to the marginal utility of wealth. As the latter is not appealing empirically, I apply a Taylor approximation to the intertemporal budget constraint (more on this below), which enables me to replace the aforementioned innovation to λ_{it} by a function of wage shocks. I combine these approximations into a single set of equations for the growth rates of household outcomes as functions of permanent and transitory wage shocks only.

To render such equations empirically useful, 1.) I add up individual private consumption, which is not observed in the data, to total private consumption in the household, and 2.) I replace market hours with earnings using the identity $Y_{jit} = W_{jit}H_{jit}$, $j = \{1, 2\}$. The final set of equations is

$$\begin{pmatrix} \Delta k_{it} \\ \Delta c_{it} \\ \Delta y_{1it} \\ \Delta y_{2it} \end{pmatrix} \approx \mathbf{T}_{it} \begin{pmatrix} v_{1it} \\ v_{2it} \\ \Delta u_{1it} \\ \Delta u_{2it} \end{pmatrix} \quad (2.9)$$

where lower case variables denote logged outcomes net of observable characteristics (*i.e.* $\Delta k_{it} = \Delta \ln K_{it}$ net of the change in covariates).¹⁶

\mathbf{T}_{it} is a 4×4 matrix of 'transmission parameters' (loading factors) of shocks into outcomes. Such transmission parameters are nonlinear functions of a large set of gender-specific Frisch (λ -constant) elasticities of each spouse plus additional parameters pertaining to financial and human wealth in the family, as well as bargaining power and the allocation of private consumption between spouses. I report the transmission parameters in \mathbf{T}_{it} analytically in appendix B.4 after I present the details of the approximations to the first

¹⁶I assume that changes in observable covariates are anticipated by partners (and ex ante contracted upon) and therefore already accounted for in their optimal choices.

Table 2.1: Frisch elasticities of spouse j **j's labor supply elasticities...** $\eta_{j,h,w}$: with respect to own wage W_j η_{j,h,p^c} : with respect to the price of private consumption P^c η_{j,h,p^k} : with respect to the price of public consumption P^k **j's private consumption elasticities...** $\eta_{j,c,w}$: with respect to own wage W_j η_{j,c,p^c} : with respect to the price of private consumption P^c η_{j,c,p^k} : with respect to the price of public consumption P^k **j's public consumption elasticities...** $\eta_{j,k,w}$: with respect to own wage W_j η_{j,k,p^c} : with respect to the price of private consumption P^c η_{j,k,p^k} : with respect to the price of public consumption P^k

Notes: This table presents the full set of Frisch (λ -constant) elasticities of spouse j . These elasticities constitute an ordinal representation of j 's preferences U_j . There are 9 elasticities per gender in total, 3 own-price and 6 cross-price elasticities. The rule governing the notation of Frisch elasticities in this chapter: $\eta_{j,x,\chi}$ is individual j 's elasticity of own outcome variable $x = \{h_j, c_j, k\}$ with respect to price $\chi = \{w_j, p^c, p^k\}$.

order conditions in appendix B.2 and the intertemporal budget constraint in appendix B.3.

The gender-specific Frisch elasticities ordinally describe spousal preferences over public & private consumption, and leisure. Table 2.1 introduces the full set of such elasticities and appendix B.1 defines them analytically.

The advantage of the aforementioned approximations over, for example, a full numerical solution to the household problem, is summarized in (2.9). Such analytical equations provide a neat picture for the contribution of wage shocks, preferences, bargaining power, consumption sharing etc. to consumption and earnings growth and, therefore, a straightforward way to identify the household structure through, for example, the covariance matrix of (2.9).

The rest of this section is devoted to a more detailed description of the steps I take to reach equations (2.9):

1. I approximate the intra-temporal first order conditions in (2.8) around

$t - 1$'s prices, choices, and intra-family bargaining power. These approximations appear in detail in appendix B.2. The approximated equations constitute a 5×5 system in (the $t - 1 \rightarrow t$ growth in) K , C_1 , C_2 , H_1 , and H_2 . Solving for the choice variables results in approximate closed-form expressions for each one of them as functions of (the growth in) wages, intra-family bargaining power, and the marginal utility of wealth. This is equation (B.8) in the appendix.

2. As the growth in intra-family bargaining power in (B.8) is unobserved, I exploit (2.6), which relates bargaining power to wages and distribution factors, and log-linearize μ_{1it} to get

$$\Delta \ln \mu_{1it} \approx \eta_{\mu, w_1} \Delta \ln W_{1it} + \eta_{\mu, w_2} \Delta \ln W_{2it} + \sum_n \eta_{\mu, d^n} \Delta \ln d_{it}^n. \quad (2.10)$$

Here the ‘surplus extraction’ elasticity η_{μ, w_j} , $j = \{1, 2\}$, captures the sensitivity of function $\mu(\cdot)$ to each spouse’s wage and η_{μ, d^n} captures the sensitivity to distribution factor d^n . If $\mu(\cdot)$ in (2.6) remains unrestricted, then η_{μ, w_j} and η_{μ, d^n} vary with μ_{1it-1} and last period’s wages and distribution factors.

I assume that only *permanent* wage shocks shift intra-family bargaining power in (2.10). As transitory shocks are rapidly mean-reverting, one can expect that they do not reflect serious enough events that can alter the allocation of bargaining power between spouses.¹⁷

3. As the growth in the marginal utility of wealth in (B.8) is unobserved, I apply a second order Taylor approximation to the Euler equation in (2.8) and decompose $\Delta \ln \lambda_{it}$ into two additive terms (see appendix B.2 and equation (B.9) for details). The first term, ω_{it} , reflects spousal motives

¹⁷Lise and Yamada (2014) find that unpredicted deviations in the spouses’ relative wages during marriage impact on intra-family bargaining powers but *not* for all households. They report that the mode of revisions (to the Pareto weight) is one, indicating no revisions at all. This evidence motivates the exclusion of transitory shocks, namely the least important shocks over the family life-cycle, from the revisions to μ_{1it} .

for precautionary savings over their life-cycle. The second component ϵ_{it} is a 0-mean innovation term which captures idiosyncratic revisions to λ_{it} in response to wage shocks. Blundell et al. (2013, appendix A.1) show that ω_{it} converges to a constant as the variance of ϵ_{it} tends to 0. To achieve tractability I assume that $\omega_{it} \equiv \omega_t$ does not vary cross-sectionally (Blundell et al., 2016, appendix).

4. Combining results from points 2-3 above, namely equations (2.10) and (B.9), I update equation (B.8) and write (the $t - 1 \rightarrow t$ growth in) K , C_1 , C_2 , H_1 , and H_2 as functions of wage shocks and the components of λ_{it} only. This is equation (B.10) in the appendix; to obtain it I have assumed away any variation in distribution factors.¹⁸
5. I use the intertemporal budget constraint (the sequential version of which appears in (2.5) above) to link ϵ_{it} to the wage shocks that hit the spouses. My aim is to replace the innovation to the marginal utility of wealth ϵ_{it} by an expression involving wage shocks only so as to render equation (B.10) empirically useful.

Specifically, I apply a first order Taylor approximation to each side of the intertemporal budget constraint around the path that would be followed if wage shocks were zero. The approximation draws on Campbell (1993); Blundell et al. (2013, 2016) are recent applications of such approximation to budget constraints that are similar to (2.5). If either spouse's current earnings are negligible compared to their expected lifetime earnings, the approximation enables me to map *permanent* wage shocks v_{jit} , $j = \{1, 2\}$, into ϵ_{it} . The details appear in appendix B.3.

The mapping from permanent wage shocks to ϵ_{it} , given by (B.11) in the appendix, involves spousal preferences (the full set of gender-specific Frisch elasticities in table 2.1), parameters pertaining to the allocation of private consumption (more on this below) and bargaining power in

¹⁸Lise and Yamada (2014) find no effect of distribution factors other than the spouses' relative wages on the allocation of resources in the household.

the household, as well as three ‘quasi-reduced-form’ parameters pertaining to expectations about the allocation of future household resources. These are: $\xi_{it} \approx \frac{\text{Lifetime Spending on } K_{it}}{\text{Lifetime Spending on } K_{it} + C_{it}}$ is the share of public in total expected lifetime expenditure, $s_{it} \approx \frac{\text{Lifetime Earnings}_{1it}}{\text{Lifetime Earnings}_{it}}$ is the share of the male spouse’s expected lifetime earnings (his human wealth) in total lifetime earnings in the family, and $\pi_{it} \approx \frac{\text{Assets}_{it}}{\text{Assets}_{it} + \text{Lifetime Earnings}_{it}}$ is the ‘partial insurance’ parameter (term due to Blundell et al., 2008) that reflects the share of assets (financial wealth) in the family’s total financial and human wealth combined.

6. A final step is needed to obtain equations (2.9); this involves replacing individual private consumption C_1 and C_2 , which is not observed in the data, by *total* private consumption using the approximation

$$\Delta \ln C_{it} \approx \varphi_{it-1} \Delta \ln C_{1it} + (1 - \varphi_{it-1}) \Delta \ln C_{2it}. \quad (2.11)$$

Here $\varphi_{it} = \frac{C_{1it}}{C_{it}}$ is the male spouse’s share of private consumption; φ_{it} will serve as a parameter to be identified in the data.

2.2.3 The Family Response to Wage Shocks

Wage shocks induce three types of effects on household outcomes, namely on public and private consumption, and spousal earnings (or labor supply). Permanent shocks induce static, bargaining, and wealth effects whereas transitory shocks induce static effects only.

The static effects summarize the impact a price change, in this case a *wage* change, has on outcomes in the exact same period when the price change occurs. Such effects are, by definition, identical across a dynamic model, like the present one, and a static one-period model. They reflect the standard substitution effect the price change induces by tilting the *contemporaneous* budget constraint and altering the marginal rate of substitution among the different goods the spouses enjoy. Permanent and transitory shocks (for a given spouse) induce identical static effects as there is no distinction between

such shocks in an one-period environment.

On top of their static effects, permanent shocks induce bargaining effects through shifting intra-family bargaining power in (2.10). Changes in intra-family bargaining power further alter the marginal rate of substitution among goods promoting those goods that the relatively powerful spouse prefers more.

Finally, permanent shocks induce wealth/income effects through shifting the *intertemporal* budget constraint as a result of a permanent change in household earnings. As household earnings are an endogenous household outcome, and therefore depend on the spouses' response to shocks through the aforementioned static and bargaining channels, the wealth effects themselves depend on (are a function of) the static and bargaining effects.

With reference to the transmission matrix \mathbf{T}_{it} in (2.9), its first two columns are the transmission parameters of permanent wage shocks comprising the sums of the static, bargaining, and wealth effects by each spouse's permanent wage shock. The last two columns are the transmission parameters of transitory wage shocks, therefore comprising the static effects only.

Bargaining effects do not operate in the unitary model. Static and wealth effects do (Blundell et al., 2016). However, in contrast to the unitary model, static effect in the collective model are explicit functions of both spouses' preferences (Frisch elasticities), the prevailing intra-family bargaining power, as well as the prevailing allocation of private consumption (the latter matters for the transmission of shocks into private consumption only). Wealth effects, being a function of static and bargaining effects, are themselves structurally richer in the collective as opposed to the unitary model.

To help build intuition, I illustrate how a transitory wage shock transmits into selected household outcomes. Without loss of generality, consider the ceteris-paribus¹⁹ transmission of $\Delta u_{1it} > 0$, a positive transitory shock to the male wage at t , into public consumption K_{it} and his leisure L_{1it} . I study the response of leisure to the transitory shock, instead of earnings or labor supply,

¹⁹Ceteris-paribus implies $v_{1it} = v_{2it} = \Delta u_{2it} = 0$.

because leisure facilitates the present illustration; recall that there is 1-to-1 mapping between leisure and earnings/labor supply in the absence of home production.

From appendix B.4 the response of those outcomes to Δu_{1it} is given analytically by

$$\begin{aligned}\Delta k_{it} &\approx \eta_{1,k,w} \times \nu_{it-1} \times \frac{\eta_{2,k,p^k}}{\bar{\eta}_{k,p^k}} \times \Delta u_{1it} \\ \Delta l_{1it} &\approx -\kappa \times \left(\eta_{1,h,w} - \eta_{1,k,w} \times (1 - \nu_{it-1}) \times \frac{\eta_{1,h,p^k}}{\bar{\eta}_{k,p^k}} \right) \times \Delta u_{1it}\end{aligned}\tag{2.12}$$

where $\eta_{1,k,w}$ is the male elasticity of the public good with respect to his wage, $\eta_{1,h,w}$ is his own-wage labor supply elasticity, and η_{1,h,p^k} is his labor supply elasticity with respect to the price of the public good (see table 2.1). ν_{it-1} is a mixture of preferences and the prevailing (*i.e.* last period's) intra-family bargaining power; $\bar{\eta}_{k,p^k} = (1 - \nu_{it-1})\eta_{1,k,p^k} + \nu_{it-1}\eta_{2,k,p^k}$ is a weighted average of the spouses' elasticities η_{j,k,p^k} of the public good with respect to its price.^{20,21}


In figure 2.1 I draw the male spouse's hypothetical preferences for leisure and the public good assuming that these goods are imperfect substitutes (this assumption is used for the purposes of this illustration only). The bold downward sloping curve ('Indifference curve 1') depicts his preferences across the two goods at the beginning of t , that is before the realization of the transitory shock. The straight line passing through A represents the original contemporaneous budget constraint at the same time. The scale of the axes is irrelevant. The spouse consumes initially at A (L_A and K_A respectively).

The positive transitory wage shock does not shift the contemporaneous budget constraint outwards. Section 2.2.2 and appendix B.3 show that, under weak assumptions, transitory shocks do not induce wealth/income effects and, therefore, leave the total available budget at t unchanged. However, such

²⁰ $\nu_{it-1} = \mu_{1it-1}U_{1,K}(\mu_{1it-1}U_{1,K} + \mu_{2it-1}U_{2,K})^{-1}$ where $U_{j,K}$ is j 's marginal utility of the public good at t . Expect that $\nu \in [0, 1]$.

²¹In the absence of home production Δl_{1it} is obtained from (2.9) as $\Delta l_{1it} = -\kappa \Delta h_{1it} = -\kappa(\Delta y_{1it} - \Delta w_{1it})$ where $\kappa > 0$ reflects Jensen's inequality when translating logged market hours to logged leisure. As a result, κ in expression (2.12) reflects that leisure elasticities are a *scaled* translation of labor supply elasticities.

Figure 2.1: The transmission of $\Delta u_{1it} > 0$ into public consumption and leisure



shock, once it realizes, tilts the budget constraint around the initial indifference curve rendering leisure relatively more expensive (because $\Delta u_{1it} > 0$ and the wage is the price of leisure) and the public good relatively cheaper. The new budget constraint is the straight line passing through B.

If the male spouse can adjust his leisure and the public good freely, he will move consumption to bundle B on the plane (L_B and K_B respectively) in response to the tilting of the budget constraint. In this case the change in his leisure is proportional to his own-wage labor supply elasticity, namely $\Delta l_{1it} = -\kappa \times \eta_{1,h,w} \times \Delta u_{1it}$, and he cuts down leisure from L_A to L_B (assuming $\eta_{1,h,w} > 0$). The change in the public good is given by his elasticity of the public good with respect to his wage, namely $\Delta k_{it} = \eta_{1,k,w} \times \Delta u_{1it}$, and he increases public consumption from K_A to K_B (assuming $\eta_{1,k,w} > 0$ in this case). These responses stem from (2.12) after setting $\nu_{it-1} = 1$ and observing that, as a result, $\frac{\eta_{2,k,p^k}}{\eta_{k,p^k}} = 1$; $\nu_{it-1} = 1$ corresponds to the extreme case where $\mu_{1it-1} = 1$ and the male spouse acts as a ‘dictator’, *i.e.* he completely imposes his preferences onto household outcomes. Consumption bundle B is one extreme possible outcome

in response to the positive transitory shock.

In the opposite extreme case, the male spouse holds no power in the household and his partner becomes the ‘dictator’. This case can be obtained when $\mu_{2it-1} = 1$ implying $\nu_{it-1} = 0$. From (2.12), $\Delta l_{1it} = -\kappa \times (\eta_{1,h,w} - \eta_{1,k,w} \times \frac{\eta_{1,h,p^k}}{\eta_{1,k,p^k}}) \times \Delta u_{1it}$ whereas $\Delta k_{it} = 0$. The (lack of) response of public consumption is straightforward: a shock to *his* wage does not alter the female spouse’s tradeoffs (marginal rates of substitution) among the goods she enjoys and, as a result, she does not change her (the household) demand for the public good. The response of his leisure is less straightforward. As μ_{2it-1} tends to 1 and the public good tends to remain at K_A , it is suboptimal for the male spouse to shift his leisure like he would if he was the ‘dictator’: in such case he would consume at C (L_B and K_A respectively) which is below the applicable budget frontier and a profound deterioration in his welfare. Instead he adjusts his leisure so that he ends up on the budget frontier given $k_{it} = K_A$: that is consumption bundle D. This adjustment differs from the first best response $-\kappa \times \eta_{1,h,w} \times \Delta u_{1it}$ by $\kappa \times \eta_{1,k,w} \times \frac{\eta_{1,h,p^k}}{\eta_{1,k,p^k}} \times \Delta u_{1it}$, that is by the product of his optimal response of K had he been able to adjust it freely ($\eta_{1,k,w}$) and his sensitivity of leisure relative to the public good ($\kappa \times \frac{\eta_{1,h,p^k}}{\eta_{1,k,p^k}}$).

In more realistic cases, neither partner would act as ‘dictator’ and spouse 1 should land somewhere between the extreme bundles B and D. The final response of the public good, given by (2.12), is a fraction ν_{it-1} of his unrestricted response $\eta_{1,k,w} \times \Delta u_{1it}$ weighed by a term $\frac{\eta_{2,k,p^k}}{\bar{\eta}_{k,p^k}}$ that reflects the female spouse’s relative sensitivity for the public good. The last term appears because, as the public good enters both spouses’ preferences, changes in its demand have to be agreed by both. The final response of his leisure, given by (2.12), is equal to his unrestricted response $-\kappa \times \eta_{1,h,w} \times \Delta u_{1it}$ plus an additional feedback term in response to the household-wide shift in the public good; such adjustment is captured by $\kappa \times \eta_{1,k,w} \times (1 - \nu_{it-1}) \times \frac{\eta_{1,h,p^k}}{\bar{\eta}_{k,p^k}} \times \Delta u_{1it}$ and exists because 1.) leisure and public consumption are nonseparable in spousal preferences, 2.) the response of public consumption, whatever that may be, reflects not just

his but *both* spouses' preferences.

Extending the discussion to all outcome variables is straightforward. The difference between the impact of Δu_{1it} on k_{it} , l_{1it} (or y_{1it}), c_{1it} on one hand and l_{2it} (or y_{2it}), c_{2it} on the other is that there are direct and indirect static effects induced on the former set, whereas there are only indirect ones induced on the latter. The direct effects capture the direct response to a transitory wage shock which, in all cases except leisure, exists because of non-separabilities in spousal preferences. The indirect effects capture adjustments in the demand of a good in response to household-wide shifts in public consumption.

The response to a permanent shock follows a similar discussion. In addition to their static effects, permanent shocks induce bargaining effects, which, in the analog of figure 2.1, restrict or expand the admissible range of outcomes between points B and D, and wealth/income effects that move the budget constraint inwards or outwards.

2.2.4 Identification of the Family Structure

In this section I discuss the identification of the household structure from panel data on consumption (household-level public and private consumption), spousal earnings, and spousal wages only. In each period t there are 6 wage variances and covariances ($\sigma_{v_1,t}^2$, $\sigma_{v_2,t}^2$, $\sigma_{v_1v_2,t}$, $\sigma_{u_1,t}^2$, $\sigma_{u_2,t}^2$, $\sigma_{u_1u_2,t}$), 18 Frisch elasticities (2 spouses \times 9 elasticities each; see table 2.1), a parameter reflecting the allocation of private consumption between spouses (φ_{it}), the Pareto weight on male preferences (μ_{1it}), 2 'surplus extraction' elasticities (η_{μ,w_j}), and a parameter reflecting a mixture of preferences and intra-family bargaining power (ν_{it}). These amount to 29 parameters in total. Three additional parameters, ξ_{it} , s_{it} , and π_{it} , enter the wealth effects but, as they are obtainable directly from the data (section 2.4.2), I treat them as known in this discussion.

Several parameters vary cross-sectionally with observables; for example the allocation of private consumption φ_{it} varies with the *level* of consumption C_{it} , the Pareto weight μ_{1it} varies with initial conditions or last period's Pareto weight μ_{1it-1} , the 'surplus extraction' elasticities η_{μ,w_j} vary with last period's

wage levels W_{jit-1} and the Pareto weight μ_{1it-1} . I discuss identification of such parameters assuming they do not vary with observables; once identification is established in such case, the extension to when they vary with observables is straightforward. Identification does not require any parameter to be stationary over time.

The parameters of the wage process are identified by wage moments only (*i.e.* independently of preferences). Specifically

$$\begin{aligned}
\sigma_{v_j,t}^2 &= \mathbb{E}[\Delta w_{jit}(\Delta w_{jit-1} + \Delta w_{jit} + \Delta w_{jit+1})] \\
\sigma_{v_1v_2,t} &= \mathbb{E}[\Delta w_{1it}(\Delta w_{2it-1} + \Delta w_{2it} + \Delta w_{2it+1})] \\
\sigma_{u_j,t}^2 &= \mathbb{E}[\Delta w_{jit}\Delta w_{jit+1}] \\
\sigma_{u_1u_2,t} &= \mathbb{E}[\Delta w_{1it}\Delta w_{2it+1}]
\end{aligned} \tag{2.13}$$

where Δw_{jit} is given by (2.7), $j = \{1, 2\}$, and $\mathbb{E}[\cdot]$ denotes the first moment. Identification follows the logic in Meghir and Pistaferri (2004) and earlier studies. $\sum_{t-1}^{t+1} \Delta w_{jit}$ strips Δw_{jit} of its transitory component and, as a result, the covariance $\mathbb{E}[\Delta w_{jit} \sum_{t-1}^{t+1} \Delta w_{jit}]$ identifies the variance of the permanent shock. Similarly, the covariance of Δw_{jit} and $\sum_{t-1}^{t+1} \Delta w_{kit}$, $j \neq k$, identifies the covariance between the partners' permanent shocks. The covariance of Δw_{jit} and Δw_{jit+1} identifies the variance of the transitory shock because consecutive transitory shocks are autocorrelated due to mean reversion. Similarly, the covariance of Δw_{jit} and Δw_{kit+1} , $j \neq k$, identifies the covariance between transitory shocks. Overidentifying restrictions are available.

The average transmission parameters of shocks into consumption and earnings, namely the components of matrix \mathbf{T}_{it} in (2.9), are identified by second moments of the joint distribution of wages, earnings, and consumption. Consider, as an example, the transmission of shocks into the public good and define the following moments

$$\begin{aligned}
m_t^{k1} &= \mathbb{E}[\Delta k_{it}(\Delta w_{1it-1} + \Delta w_{1it} + \Delta w_{1it+1})] \\
m_t^{k2} &= \mathbb{E}[\Delta k_{it}(\Delta w_{2it-1} + \Delta w_{2it} + \Delta w_{2it+1})]
\end{aligned}$$

$$m_t^{k3} = \mathbb{E}[\Delta k_{it} \Delta w_{1it+1}]$$

$$m_t^{k4} = \mathbb{E}[\Delta k_{it} \Delta w_{2it+1}]$$

where Δk_{it} is given by (2.9). The average transmission parameters into the public good are identified as

$$\begin{aligned} \mathbb{E}[\tau_{it}^{11}] &= \frac{m_t^{k1} \sigma_{v_2,t}^2 - m_t^{k2} \sigma_{v_1 v_2,t}}{\sigma_{v_1,t}^2 \sigma_{v_2,t}^2 - \sigma_{v_1 v_2,t}^2} & \mathbb{E}[\tau_{it}^{13}] &= -\frac{m_t^{k3} \sigma_{u_2,t}^2 - m_t^{k4} \sigma_{u_1 u_2,t}}{\sigma_{u_1,t}^2 \sigma_{u_2,t}^2 - \sigma_{u_1 u_2,t}^2} \\ \mathbb{E}[\tau_{it}^{12}] &= \frac{m_t^{k1} \sigma_{v_1 v_2,t} - m_t^{k2} \sigma_{v_1,t}^2}{\sigma_{v_1 v_2,t}^2 - \sigma_{v_1,t}^2 \sigma_{v_2,t}^2} & \mathbb{E}[\tau_{it}^{14}] &= -\frac{m_t^{k3} \sigma_{u_1 u_2,t} - m_t^{k4} \sigma_{u_1,t}^2}{\sigma_{u_1 u_2,t}^2 - \sigma_{u_1,t}^2 \sigma_{u_2,t}^2} \end{aligned}$$

where $(\tau_{it}^{11}, \tau_{it}^{12}, \tau_{it}^{13}, \tau_{it}^{14})$ is the first row of the transmission matrix \mathbf{T}_{it} in (2.9).

Identification rests on the following idea: if permanent wage shocks transmit into public consumption, the contemporaneous covariance between spouse j 's wage growth (stripped of its transitory components) and public consumption growth picks up the variance of spouse j 's permanent shock scaled by its transmission parameter into public consumption. Similarly, the covariance between contemporaneous growth in public consumption and next period's wage growth picks up (minus) the variance of the mean reverting transitory shock weighed by its loading factor onto public consumption. In both cases adjustments are made to account for the correlation of wages in the family.

Identification of the remaining transmission parameters (into total private consumption and individual earnings) follows the same logic and requires analogous moment conditions. A total of 16 transmission parameters are identified.

The question that emerges is whether such transmission parameters convey sufficient information for the identification of the underlying household structure. The answer is negative; there are only 16 transmission parameters comprising a total of 29 structural parameters and most one can identify is a few uninformative ratios of Frisch elasticities. Reasons behind the lack of identification include 1.) the lack of variation in the price of the consumption goods, 2.) the non-observability of individual private consumption, 3.) the ex-

istence of a public good bridging spousal preferences and the interdependence of goods allowed for by non-separabilities in spousal preferences, and 4.) the indistinguishability between preferences and intra-family bargaining power.

I explore two possibilities in order to overcome the lack of identification. In section 2.2.4.1 I impose restrictions on spousal preferences, rendering the public good additively separable from private consumption and leisure, and enforce full commitment. In section 2.2.4.2 I maintain nonseparable spousal preferences and lack of commitment but I bring additional information into (2.9), namely information on single individuals of each gender. In both cases I obtain partial identification of the household structure.

2.2.4.1 Separable Public Good and Full Commitment

Additive separability of the public good implies that, conditional on household income, public consumption does not change with wages or the price of private consumption, leading to $\eta_{j,k,w} = \eta_{j,k,p^c} = \eta_{j,h,p^k} = \eta_{j,c,p^k} = 0$ for $j = \{1, 2\}$. Changes in wages still trigger shifts in the public good through the budget constraint only. Full commitment implies that bargaining power in a given household remains fixed over time ($\mu_{1it} = \mu_{1it=0}, \forall t$) leading to $\eta_{\mu,w_j} = 0$ and 0 bargaining effects.

The transmission parameters of wage shocks into household outcomes are obtained from the expressions in appendix B.4 plugging in the aforementioned restrictions. Note that private consumption and leisure remain nonseparable.

It is straightforward to see that the response of spousal earnings to own transitory shocks identifies the own-wage labor supply elasticities $\eta_{j,h,w}$, $j = \{1, 2\}$. Such response reflects the direct static effect on earnings (the indirect static effect, which would feed back should public consumption also shift, is 0) which is proportional to one's labor supply elasticity.

The response of private consumption to transitory shocks identifies the private consumption elasticities $\eta_{j,c,w}$ *up to scale* where the scale is given by a spouse's share of private consumption in the household. Specifically, the response of private consumption to the male transitory shock identifies

$\mathbb{E}[\varphi_{it}]\eta_{1,c,w}$ whereas the response to the female shock identifies $\mathbb{E}[1 - \varphi_{it}]\eta_{1,c,w}$. Again, such responses reflect the direct static effect on private consumption, which is proportional to one's own-wage private consumption elasticity *and* their share of private consumption. It is not possible to separate the two.

To identify the hours elasticities with respect to the price of private consumption (η_{j,h,p^c}) I exploit a natural set of theoretical restrictions, that is symmetry of each spouse's matrix of substitution effects. After a constant marginal-utility-of-wealth price change (essentially a transitory shock), the matrix of substitution effects, that is (B.1) in the appendix, is equal to the (scaled) inverse Hessian of the individual utility function. Symmetry of this matrix follows from symmetry of the Hessian. In the present context, this symmetry implies one linear restriction per spouse involving their non-zero cross-price elasticities, namely $\eta_{1,h,p^c} = -\eta_{1,c,w} \times \varphi_{it} \times \frac{P_t^c C_{it}}{W_{1it} H_{1it}}$ and $\eta_{2,h,p^c} = -\eta_{2,c,w} \times (1 - \varphi_{it}) \times \frac{P_t^c C_{it}}{W_{2it} H_{2it}}$.²²

The response of consumption to permanent shocks identifies two quasi-reduced-form consumption substitution elasticities at the household level. Specifically, the response of public consumption identifies

$$\tilde{\eta}_{k,p^k} = \mathbb{E} \left[\eta_{1,k,p^k} \eta_{2,k,p^k} / \left((1 - \nu_{it-1}) \eta_{1,k,p^k} + \nu_{it-1} \eta_{2,k,p^k} \right) \right]$$

whereas the response of private consumption identifies

$$\tilde{\eta}_{c,p^c} = \mathbb{E} \left[\varphi_{it-1} \eta_{1,c,p^c} + (1 - \varphi_{it-1}) \eta_{2,c,p^c} \right].$$

The former parameter coincides with the *household-level* public consumption substitution elasticity, *i.e.* the substitution elasticity of public consumption when the household acts as a single economic agent (the case of the unitary model). The latter parameter coincides with the *household-level* private

²²See appendix B.1 for details. Frisch symmetry is not the only way to identify η_{j,h,p^c} , $j = \{1, 2\}$. In principle, those elasticities can also be identified from the transmission of *permanent* shocks into household outcomes. In practice the transmission of permanent shocks alone could not pin down η_{j,h,p^c} precisely.

consumption substitution elasticity, *i.e.* the substitution elasticity of private consumption in a unitary household. The response of earnings to permanent shocks provides overidentifying restrictions for $\tilde{\eta}_{k,p^k}$ and $\tilde{\eta}_{c,p^c}$.

A note of caution is due here: even though additive separability of public consumption delivers identification of a number of structural parameters, the demand for ‘public’ items, such as water or electricity, is likely to co-move with hours of work (see, for example, Browning and Meghir, 1991). Casual arguments suggest that water or electricity consumption may increase if individuals work less and spend more time at home. In the present model, additive separability of K also implies that one partner’s earnings do not vary with their spouse’s wage, thus stripping earnings of the ‘added worker’ effect (Lundberg, 1985). This is because such effect operates here through the interdependence of preferences due to a nonseparable public good and it vanishes otherwise. Finally, in the absence of bargaining effects, full commitment renders the model indistinguishable from a unitary one unless $\tilde{\eta}_{k,p^k}$ varies with past wages through its dependence on ν_{it-1} and last period’s Pareto weight.²³

2.2.4.2 Nonseparable Public Good and Lack of Commitment

Suppose single and married individuals of gender $j = \{1, 2\}$ have the same preferences across the two states of life (as singles, in couples). Suppose, also, that the entire set of Frisch elasticities per gender $\eta_j = \{\eta_{j,h,w}, \eta_{j,h,p^c}, \eta_{j,h,p^k}; \eta_{j,c,w}, \eta_{j,c,p^c}, \eta_{j,c,p^k}; \eta_{j,k,w}, \eta_{j,k,p^c}, \eta_{j,k,p^k}\}$ is available through, say, observing singles. Then the second moments of the joint distribution of wages, earnings, and consumption in the family, whose theoretical counterpart is the covariance matrix of (2.9), suffice for the (over)-identification of the remaining household structure.

Once spousal preferences (the Frisch elasticities) are known, the response of public consumption to either spouse’s transitory shock identifies $\mathbb{E}[\nu_{it-1}]$.

²³One cannot distinguish this model from the unitary if φ_{it} alone and, therefore $\tilde{\eta}_{c,p^c}$ or the response of private consumption to transitory shocks, varies with wages or other observables. The non-separability of leisure and private consumption implies that the allocation of private consumption will vary with wages even when the unitary model is the ‘correct’ representation of the household.

ν_{it-1} is a ‘preference aggregator’; it determines the extent to which shifts in the household demand for the public good reflect one or another spouse’s preferences. Similarly, the response of private consumption to either spouse’s transitory shock identifies the average consumption allocation between them $\mathbb{E}[\varphi_{it-1}]$. The response of earnings provides over-identifying restrictions.

In both cases the rationale behind identification is the following: knowledge of spousal preferences implies knowledge of how each spouse would, hypothetically, change their consumption *as single*. The actual consumption response at the family level is a weighted sum/product of the two individual responses; the weight is given by ν in the case of public consumption and by φ in the case of private consumption. An equivalent interpretation is the following: the extent to which the response of family consumption to wage shocks resembles the way one or another partner would have responded had they been single is informative about the decision process for the public good as well as the allocation of the private good.

The response of household outcomes to permanent shocks identifies 8 bargaining effects (4 bargaining effects -on public consumption, private consumption, and earnings- by *each* spouse’s permanent shock). Permanent shocks induce bargaining and wealth effects; as the latter are functions of the former *and* the static effects that are identified from transitory shocks and singles, the response to permanent shocks just identifies the bargaining effects.²⁴ It

²⁴Without loss of generality consider the bargaining effects $\beta_{k,w_1}, \beta_{c,w_1}, \beta_{y_1,w_1}, \beta_{y_2,w_1}$ and the wealth effects $\gamma_{k,w_1}, \gamma_{c,w_1}, \gamma_{y_1,w_1}, \gamma_{y_2,w_1}$ induced by men’s permanent shock. I write

$$\begin{pmatrix} \beta_{k,w_1} + \gamma_{k,w_1} \\ \beta_{c,w_1} + \gamma_{c,w_1} \\ \beta_{y_1,w_1} + \gamma_{y_1,w_1} \\ \beta_{y_2,w_1} + \gamma_{y_2,w_1} \end{pmatrix} = \begin{pmatrix} \varpi_{k,w_1} \\ \varpi_{c,w_1} \\ \varpi_{y_1,w_1} \\ \varpi_{y_2,w_1} \end{pmatrix}$$

where $(\varpi_{k,w_1}, \varpi_{c,w_1}, \varpi_{y_1,w_1}, \varpi_{y_2,w_1})' \neq \mathbf{0}$ is the difference between the transmission parameters of permanent and transitory shocks to men’s wage. Conditional on the static effects, the wealth effects γ are linear functions of the bargaining effects (see appendix B.4). Manipulating this system appropriately, I write

$$\mathbf{M}_{it} (\beta_{k,w_1}, \beta_{c,w_1}, \beta_{y_1,w_1}, \beta_{y_2,w_1})' = \mathbf{m}_{it}.$$

\mathbf{M}_{it} is a data-dependent matrix that is unlikely to be singular. \mathbf{m}_{it} is unlikely to be the zero vector. There is a unique solution for $(\beta_{k,w_1}, \beta_{c,w_1}, \beta_{y_1,w_1}, \beta_{y_2,w_1})'$ given by $\mathbf{M}_{it}^{-1} \mathbf{m}_{it}$.

is not possible to delve deeper and identify components of the bargaining effects, namely η_{μ,w_j} or μ_{1it-1} . The ‘surplus extraction’ elasticities η_{μ,w_j} cannot be separately identified from the (prevailing) Pareto weight μ_{1it-1} . Indeed, inspecting the bargaining effects in appendix B.4, these elasticities appear always together with and multiplicatively to the Pareto weight.

But how can one obtain η_j ? The following section (section 2.3) develops a life-cycle model for consumption and leisure of singles. Based on such model I show how the entire set of Frisch elasticities per gender η_j can be identified and used as input to the model for couples.

2.3 A Life-Cycle Model for Singles

Single individuals of gender $j = \{1, 2\}$ live alone and consume the same set of goods as their non-single counterparts, namely ‘public’ consumption K , ‘private’ consumption C , and leisure $1 - H$. Public consumption is no longer ‘public’ in the sense that it is enjoyed together with a spouse because such spouse does not exist; however, it still comprises the same goods that spouses would generally consume together (such as last section’s car journey to the countryside or the gas to heat the house). Private consumption C also comprises the same goods as in section 2.2 (e.g. a meal at the local coffee shop or a subway ticket). Although a distinction between K and C is now less meaningful, I do not collapse them to a single commodity because I want to maintain consistency with the family problem of section 2.2.

I illustrate the main points of the model focusing on single j . Using a time index s , the single individual’s problem is given by

$$\max_{\{K_{jis}, C_{jis}, H_{jis}, A_{jis+1}\}_{\forall s}} \mathbb{E}_0 \sum_{s=0}^S \beta_s U_j(K_{jis}, C_{jis}, 1 - H_{jis}; \mathbf{z}_{jis}) \quad (2.14)$$

subject to the sequential budget constraint

$$A_{jis} + W_{jis}H_{jis} = K_{jis} + P_t C_{jis} + \frac{A_{jis+1}}{1+r}, \quad A_{jiS+1} = \tilde{A}, \quad \forall t. \quad (2.15)$$

Here \mathcal{S} indicates the horizon of the single's lifetime and deserves some attention. In the case of a young single individual who may get married in the future, \mathcal{S} reflects the uncertain future time when he/she gets married; \tilde{A} reflects assets, positive or negative, the single transfers to his/her future family. In the case of an older single who was married in the past but is now divorced or widowed, \mathcal{S} reflects the end of lifetime (assuming no remarriage) and $\tilde{A} = 0$ is the usual no-Ponzi terminal condition. (2.14)-(2.15) could also reflect the problem of a person who remains single for their entire lifetime, in which case \mathcal{S} and \tilde{A} are like for the older single above.

The public good is now subscripted by j to indicate its assignability to single j ; assets are also assignable and thus also subscripted by j . Private consumption C_{jis} is observed and fully assignable to individual j . The wage process is given by (2.7). The rest of the notation, as well as the properties of U_j , remain exactly the same as in section 2.2.1.

I solve the problem following the same procedure like before (section 2.2.2): I obtain the problem's first order conditions, I apply first and second order Taylor approximations to them, I log-linearize the singles' intertemporal budget constraint. I obtain a set of closed-form equations for (growth in) singles' consumption and earnings as functions of their permanent and transitory wage shocks. These equations are given by

$$\begin{pmatrix} \Delta k_{jis} \\ \Delta c_{jis} \\ \Delta y_{jis} \end{pmatrix} \approx \begin{pmatrix} \eta_{j,k,w} + (\eta_{j,k,w} + \eta_{j,k,p^c} + \eta_{j,k,p^k})\ell_{jit}^s & \eta_{j,k,w} \\ \eta_{j,c,w} + (\eta_{j,c,w} + \eta_{j,c,p^c} + \eta_{j,c,p^k})\ell_{jit}^s & \eta_{j,c,w} \\ \eta_{j,h,w} + (\eta_{j,h,w} + \eta_{j,h,p^c} + \eta_{j,h,p^k})\ell_{jit}^s + 1 & \eta_{j,h,w} + 1 \end{pmatrix} \begin{pmatrix} v_{jis} \\ \Delta u_{jis} \end{pmatrix} \quad (2.16)$$

and they are the singles' analog of equations (2.9) in the family problem. The transmission of the permanent shock differs from that of the transitory shock solely by the wealth/income effect it induces as there are no bargaining effects in a single-member household. The wealth/income effect is captured by ℓ_{jit}^s , which stems from the log-linearization of the intertemporal budget constraint and is defined analytically in appendix B.3, and the terms multiplying

it.²⁵

The response of public consumption to the transitory wage shock identifies the public consumption wage elasticity $\eta_{j,k,w}$. Specifically, $\mathbb{E}[\Delta k_{jis} \Delta w_{jis+1}]$ identifies (minus) this elasticity scaled by the variance of the transitory shock; the latter is identified by (2.13). Similarly, the response of private consumption (earnings) identifies the private consumption wage elasticity $\eta_{j,c,w}$ (the own-wage labor supply elasticity $\eta_{j,h,w}$). Identification of $\eta_{j,k,w}$ and $\eta_{j,c,w}$, and symmetry of the matrix of substitution effects (illustrated in section 2.2.4.1 and appendix B.1), brings about identification of the ‘reciprocal’ labor supply elasticities $\eta_{j,h,p^k} = -\eta_{j,k,w} \frac{P_s^k K_{jis}}{W_{jis} H_{jis}}$ and $\eta_{j,h,p^c} = -\eta_{j,c,w} \frac{P_s^c C_{jis}}{W_{jis} H_{jis}}$.

The response of consumption to permanent shocks identifies two sums of consumption substitution elasticities, namely $\eta_{j,k,p^c} + \eta_{j,k,p^k}$ and $\eta_{j,c,p^c} + \eta_{j,c,p^k}$. Importantly, identification of the sums is robust to misspecification in the single person’s intertemporal budget constraint due to over- or under-estimating the true horizon \mathcal{S} of the single.²⁶

It is not possible to separate the components of those sums, namely η_{j,k,p^c} from η_{j,k,p^k} or η_{j,c,p^c} from η_{j,c,p^k} , even if one restricts the ‘reciprocal’ elasticities η_{j,k,p^c} and η_{j,c,p^k} by exploiting the symmetry in the matrix of substitution effects. Lack of variation in the prices of the two consumption goods is the fundamental reason. But even if such variation was obtainable in the data, it would be even harder to obtain *independent* variation in price P^c from price P^k . This observation motivates the restriction $P^c = P^k = P$ which implies $\eta_{j,k,p^c} = \eta_{j,k,p^k}$ and $\eta_{j,c,p^c} = \eta_{j,c,p^k}$.²⁷ These additional restrictions complete the

²⁵Transitory shocks do not induce wealth/income effects provided that the lifetime horizon *as single* is sufficiently long and earnings within a given period are negligible compared to lifetime earnings as single.

²⁶The response of earnings to permanent shocks identifies $\mathbb{E}[\ell_{jit}^s]$ conditional on $\eta_{j,h,w}$, η_{j,h,p^c} , and η_{j,h,p^k} . Even though ℓ_{jit}^s depends on the singles’s horizon \mathcal{S} through (B.12) in the appendix, identification of $\mathbb{E}[\ell_{jit}^s]$ does not. Indeed one could replace ℓ_{jit}^s by a flexible variable b_{jit} (which, unlike ℓ_{jit}^s , has an unknown analytical representation), so that the response of earnings to permanent shocks identifies $\mathbb{E}[b_{jit}]$. Conditional on this, the response of consumption to permanent shocks identifies the aforementioned sums $\eta_{j,k,p^c} + \eta_{j,k,p^k}$ and $\eta_{j,c,p^c} + \eta_{j,c,p^k}$. Of course, if the true horizon of the single *is* known, then the analytical representation of ℓ_{jit}^s has the advantage of providing overidentifying restrictions.

²⁷It also implies $\eta_{j,h,p^c} = \eta_{j,h,p^k}$.

identification of the entire set of Frisch elasticities per gender.

Conditional on these elasticities and as illustrated in section 2.2.4.2, the response of consumption to transitory shocks in a two-member household identifies, among other things, the sharing of private consumption between spouses; the response of consumption to permanent shocks identifies the bargaining effects on household outcomes. In practice, one would estimate the gender-specific preferences (Frisch elasticities) and the parameters of the collective household *jointly* across married couples, single males, and single females; therefore the Frisch elasticities would need to satisfy a big number of over-identifying restrictions across singles *and* couples.

The critical assumption that makes this identification strategy work is that spousal, gender-specific, preferences do not change across the two states of life, namely across single and married individuals, or change only through their dependence on taste shifters and observables. A number of studies have employed this assumption to identify the structure of the collective household, such as Barmby and Smith (2001), Vermeulen (2006), Vermeulen et al. (2006), or, more recently, Browning et al. (2013).

There are a few different ways to weaken such assumption empirically. For example, one may restrict the estimation sample to singles who have preferences for marriage, or otherwise for forming a household, by only admitting those who were married in the past or will get married in the future. Alternatively, one may admit singles in a short window of time around the ‘marriage cutoff’, that is for a few only periods before one gets married or after one divorces. I impose the first restriction in the subsequent empirical application (section 2.4).

Note that a class of preferences are excluded from U_j upon assuming preferences are state-of-life invariant. Spouse j in the family cannot be altruistic, therefore caring for his/her partner, as that would imply a change of preferences across the two states of life. As single, one cannot care for another person because there is no one else in the household to care for. When partnered, one

will still have to *not* care for their partner's utility, other than caring through the common public good, because, otherwise, preferences would change violating the aforementioned assumption.

2.4 Empirical Implementation

The model is estimated on seven waves of the PSID between 1999 and 2011; the 1997 wave is used for initial conditions.²⁸ The PSID started in 1968 interviewing a -then- nationally representative core sample of roughly 3,000 households; repeated annually until 1997 the survey collected information on employment, income, health, education and other demographics of the adult household members and their linear descendants should they split off and establish their own households. A second smaller sample of low income households, consisting roughly of 2,000 units in 1968, was also interviewed consistently. I estimate the model on the core sample only.

The survey becomes biennial after 1997 but, starting in 1999, it collects richer information including information on expenditure, wealth, philanthropy, and numerous other topics. The sample size has grown consistently over the years reaching 5,495 core sample households in 2011 (this reflects tracking of an increasing number of the original families' first and subsequent generations split-offs).

The PSID is suitable for the model in this chapter due to a number of desirable features: (i) detailed data on household assets and spending are available after 1999, along with data on earnings, hours of work, and demographics for the main earners, (ii) consecutive information on the same households is available, and (iii) multi-member, as well as single-member households are interviewed.

²⁸More information on the PSID, as well as access to all the data, is available online at psidonline.isr.umich.edu.

2.4.1 Sample Selection and Definition of Variables

I consider opposite-sex married couples; I exclude those cohabiting (a tiny fraction in the data). I assign the model spouse $j = 1$ to the male spouse in the data and the model spouse $j = 2$ to the female. Their single counterparts in the data are single men and single women respectively.

I select a baseline sample of stable²⁹ couples such that both partners are between 25 and 65 years old and have no missing demographics such as race, education, employment status, or state of residence. Both spouses must participate in the labor market and earn an hourly wage at least equal to half the applicable state minimum wage. The family must consume non-zero amounts of both the public and the private goods and report usable information on their wealth (precise definitions of the variables follow). I remove observations with extreme values for wages, earnings, or family consumption; these variables must also not experience extreme jumps from one period to another (jumps that probably signal measurement error).

Singles are selected according to the same exactly criteria like couples; their only difference is that they live in single-member households. I restrict my attention to single men and women who were married in the past or will get married in the future. In other words, I drop singles who may have tastes against marriage, something which would go against the assumption that preferences are state-of-life invariant.

In total there are 10,232 married couple-year observations satisfying the above selection criteria, 853 single male-year, and 1,109 single female-year (the totals include calendar year 1996).³⁰ There are more single women than men because of more widowed or divorced women in all years of the data.

Table 2.2 presents information on demographics and labor market out-

²⁹By stable households I refer to households whose lead couple (husband-wife) remains intact. If a new spouse enters a household, because, say, of divorce and remarriage, I drop the household the period this occurs and reinstate it in the following period as a new household. I also drop intermittent households, *i.e.* households that appear in the survey only once or in waves too far apart.

³⁰The PSID data are retrospective, *i.e.* information in the 1997 survey wave refers to calendar year 1996. I report descriptive statistics respecting this feature.

Table 2.2: Demographics and labor market outcomes by year

	Couples		Singles		Couples		Singles	
	Men	Women	Men	Women	Men	Women	Men	Women
	1996				1998			
# of children	1.21		0.25	0.92	1.13		0.29	1.02
Age	41.46	39.49	33.82	33.60	42.62	40.60	36.04	36.15
Race: white %	90.25	90.82	90.00	76.71	90.19	90.97	86.61	81.90
Race: black %	4.69	3.94	3.33	17.81	4.60	3.74	6.25	12.93
Years of schooling	14.12	14.11	14.18	14.44	14.10	14.17	14.18	14.18
Been to college %	64.67	66.82	66.67	75.34	64.49	67.76	66.07	68.97
Hours of work	2245.64	1702.04	2092.92	2080.75	2314.40	1693.95	2277.21	2016.11
Earnings	62128.75	33595.15	49882.41	36042.82	66903.14	34819.23	54236.60	37753.33
Hourly wage	28.69	19.91	24.23	17.23	29.70	20.72	24.30	19.07
Observations	1067		60	73	1284		112	116
	2000				2002			
# of children	1.06		0.24	0.92	0.99		0.30	0.84
Age	43.36	41.35	37.45	37.72	43.52	41.50	38.50	39.99
Race: white %	89.72	91.33	86.67	80.00	89.82	91.04	84.89	78.91
Race: black %	4.53	3.68	6.67	12.31	4.13	3.45	7.19	12.24
Years of schooling	14.15	14.27	14.05	14.05	14.21	14.38	14.09	13.92
Been to college %	65.31	68.69	62.22	66.15	65.70	70.90	65.47	65.31
Hours of work	2276.58	1716.42	2325.59	1915.23	2266.93	1741.58	2291.68	1836.51
Earnings	70501.86	37083.09	56672.66	38274.12	67114.87	39305.18	58799.85	31587.42
Hourly wage	31.65	21.70	26.03	20.14	30.28	22.91	27.03	17.32
Observations	1303		135	130	1306		139	147

(previous table continued)

	Couples		Singles		Couples		Singles	
	Men	Women	Men	Women	Men	Women	Men	Women
	2004		2006		2006		2010	
# of children	0.96		0.37	0.95	0.95		0.29	0.98
Age	43.79	41.93	40.14	40.54	44.03	42.25	42.16	41.08
Race: white %	90.15	90.68	87.79	78.85	90.63	90.99	89.91	78.48
Race: black %	4.06	3.46	6.11	12.82	3.82	3.39	6.42	13.92
Years of schooling	14.31	14.50	14.04	13.89	14.33	14.60	13.99	14.02
Been to college %	68.50	73.76	64.89	65.38	68.49	75.99	60.55	67.72
Hours of work	2280.20	1748.17	2329.37	1876.85	2247.03	1714.82	2091.02	1985.60
Earnings	68983.63	39717.41	63084.30	32968.85	70084.23	39904.72	53837.94	35567.90
Hourly wage	30.98	22.98	27.71	17.42	31.48	23.41	26.54	18.70
Observations	1330		131	156	1387		109	158
	2008		2008		2010		2010	
# of children	0.97		0.22	0.85	0.98		0.20	0.85
Age	43.84	41.99	46.20	42.86	45.14	43.37	48.99	45.02
Race: white %	90.33	91.53	85.71	77.30	89.98	92.18	85.51	75.69
Race: black %	4.02	3.25	8.16	10.81	4.04	3.16	8.70	11.11
Years of schooling	14.36	14.66	13.90	14.17	14.40	14.72	13.77	14.19
Been to college %	69.58	77.21	58.16	68.65	69.86	77.33	56.52	70.83
Hours of work	2133.30	1708.51	2067.77	1825.87	2135.16	1723.74	1962.77	1845.67
Earnings	68921.86	41683.20	60421.64	36851.53	69804.70	42263.06	49002.61	36506.96
Hourly wage	32.81	24.17	28.00	19.98	33.66	24.26	23.90	19.55
Observations	1417		98	185	1138		69	144

Notes: This table presents averages of demographics and labor market outcomes by year (1996-2010), gender (male-female), and marital state (married, single). '# of children' indicates the average number of children. 'Been to college %' indicates the proportion of people with more than 12 years of formal schooling. Monetary amounts (earnings and hourly wages) are expressed in \$2010.

comes by year (1996-2010), gender (male-female), and marital status (married, single). Mean age of individuals in couples is in the early 40s (with women approximately two years younger than their husbands). Mean age of singles trends upwards over time reflecting that the majority of young singles in the later survey waves are not included in the sample because I cannot observe them getting married in the future. Approximately 90% of couples comprise white individuals; this figure is slightly lower among single men (-4pct) and considerably lower among single women (-12pct). Average years of education, as well as the likelihood of having been to college, increases over time for married men and women; the latter almost always outperform the former. There is mixed evidence on the educational attainment of singles: single men have progressively fewer years of schooling, possibly reflecting changes in the age composition of single men, whereas single women maintain stable schooling levels despite that.

Earnings are defined on an annual basis and include labor income (including tips and overtime) and the labor part of business income from unincorporated businesses. Hourly wages are defined as annual earnings over annual hours of market work. Women earn consistently less than men across both marital states but they also work fewer hours. Single men work approximately the same hours as men in couples but they earn less; single women work more than their married counterparts but earn slightly less.

The PSID collects information on numerous elementary expenditure items (see Blundell et al., 2016, for how the PSID compares in that respect to the National Income and Product Accounts). To meet the requirements of the model I categorize and aggregate the elementary items into a private and a public good considering which items may be rival among family members and which may not. There is no easy way to draw a line between private and public and I treat the following categorization as *baseline* – private consumption comprises food at home, food out, transport on public means, medical services excluding health insurance, and prescriptions; public consumption comprises housing

services, home insurance, health insurance, utilities including gas, electricity, water and sewer, children's schooling and child care, and vehicle utilization costs including motor fuel.³¹

Tables 2.3 and 2.4 present summary statistics for private and public consumption respectively. Across the two tables, the first row presents by year average real expenditure on goods by couples (expressed in \$2010). The second and third rows present average real expenditure by single males and females respectively as a *percentage* of the expenditure of couples in the same year. In the subsequent rows I break down expenditure *of couples* to its elementary components. I do not repeat for single men or women as their small sample sizes along with the infrequency of purchases of disaggregated items prohibits such illustration.

The average private expenditure of couples, expressed in \$2010, is approximately \$11,500 over the period from 1998 to 2010; interestingly this average does not fluctuate over time. The average private expenditure of single men is approximately 65% of the couples' average; single women's is slightly less but still above 60% of the couples' average.

The average public expenditure of couples, expressed in \$2010, is \$26,244 in 1998; this increases steadily to \$38,057 in 2006 but drops sharply in the years of the financial crisis. The average public expenditure of single men is approximately 60%-65% of the couples' average; this fluctuates considerably over time reflecting, possibly, the compositional changes in the age of singles in the sample. The average public expenditure of single women is again slightly lower than single men's and also fluctuates with time.

Expenditure on public goods is consistently the largest part of household expenditure across marital states; it amounts approximately to 72.1% of a married couple's total expenditure, 70.3% of a single male's, and 71.8% of a single female's. Interestingly, for all marital states the share of public expenditure in total expenditure increases, slowly but steadily, over time.

³¹This categorization excludes goods that were added in the PSID after wave 2005 such as clothing and recreational goods.

Table 2.3: Private consumption by year and marital status

	1998	2000	2002	2004	2006	2008	2010
<i>couples \$</i>	11155.9	11561.9	11593.9	11965.5	11683.3	11243	11961.8
<i>single males %</i>	63.9%	65.7%	69.0%	74.7%	65.9%	64.4%	60.8%
<i>single females %</i>	64.3%	64.4%	61.3%	59.8%	59.7%	64.2%	66.3%
Breakdown of Private Consumption of Couples:							
Food at Home	6959.5	6947.4	6935	6920	6779.3	6654.5	7033.3
prepared at home	6781.1	6767.4	6720.4	6731.5	6616.4	6505.7	6904.3
delivered at home	178.3	180	214.5	188.4	162.9	148.9	129
Food Out	2791.2	2942.7	2952.8	3111	2875.3	2615.7	2896
Public Transport	262.2	382.5	304.5	251.8	234.4	205.4	167.9
buses and trains	53.7	66.7	58.4	60.3	54.1	64.5	82.7
all other means	208.5	315.9	246.1	191.4	180.3	140.9	85.2
Medical Services	895.3	934.2	1008.4	1201.3	1319.3	1316.4	1391
nurses-hospitals	262.2	245.6	267.6	333.9	433.4	429.1	497.1
professionals	633.2	688.6	740.8	867.4	886	887.4	893.9
Prescriptions	247.7	355.1	393.2	481.5	474.9	450.9	473.7

Notes: The first row presents the average real private expenditure on goods by couples. The second and third rows present average real private expenditure by single males and females respectively as a *percentage* of the expenditure of couples in the same year. In the subsequent rows I break down private expenditure of couples to its elementary components. Food at home comprises food prepared at home and food delivered at home (for recipients and non-recipients of food stamps). Public transport comprises buses and trains and all other means (including taxicabs). Medical services comprise payments to nurses, hospitals, physicians, and other professionals. All monetary figures are expressed in \$2010.

Table 2.4: Public consumption by year and marital status

	1998	2000	2002	2004	2006	2008	2010
<i>couples \$</i>	26243.9	29417.9	30923.2	36532.1	38057.3	34046.4	36820.3
<i>single males %</i>	66.2%	61.9%	60.1%	59.1%	60.6%	65.3%	49.0%
<i>single females %</i>	62.0%	61.8%	55.1%	57.8%	61.1%	65.0%	59.5%
Breakdown of Public Consumption of Couples:							
Housing Services	11742.5	13215.9	14422.9	16939.9	17719.4	15392.3	15424.7
renters	953.6	993.6	965.3	973.5	1102.8	1200.7	1171
owners	10710.5	12155	13397.7	15885.7	16464.5	14098.1	14171.4
Home Insurance	583.9	630.9	685.1	737.5	758	746.4	861
Health Insurance	1341.6	1596.7	1845.2	2075.3	2088.7	2198.4	2392.4
Utilities	3053.4	3393.1	3190.5	5443.9	5705.8	5860.9	6172.1
heating-electricity	2312.6	2602.1	2378.5	2517.3	2622.9	2680	2742.6
water and sewer	408	442.9	440.7	426.6	436	460.5	515.7
Children's Education	2598.6	2917.6	2991	3126.3	3366.8	2745.2	3115.6
Child Care	937.5	941.5	988.1	941	800.2	936.8	1150.9
Auto Vehicles	5986.4	6722.2	6800.5	7268.2	7618.5	6166.5	7703.7
motor fuel	1898.1	2692.2	2415.5	3204.9	3727.1	2759.7	4037.4
insurance	2051.8	1960.7	2438.6	2277.5	1966.7	1674.7	1787.2
repairs	1965.8	2002.9	1869.6	1708.9	1862.2	1669.1	1825.2

Notes: The first row presents the average real public expenditure on goods by couples. The second and third rows present average real public expenditure by single males and females respectively as a *percentage* of the expenditure of couples in the same year. In the subsequent rows I break down public expenditure of *couples* to its elementary components. Housing services comprise services rendered to renters and services rendered to owners. I proxy the latter as 6% of the self-reported house value per year. For those who have been offered public or similar housing I utilize a self-reported estimate of a rent-equivalent. Utilities comprise gas, electricity, water and sewer, and miscellaneous items. Auto vehicles comprise vehicle insurance, motor fuel, repair costs, and parking fees. All monetary figures are expressed in \$2010.

Table 2.5: Household wealth by year and marital status

	1998	2000	2002	2004	2006	2008	2010
<i>couples</i> \$	330469.6	348666.8	328168.7	390873.9	441436.4	332819	339367.7
<i>single males</i> %	32.3%	44.5%	46.3%	56.9%	44.8%	93.1%	53.8%
<i>single females</i> %	22.2%	24.9%	59.2%	35.8%	36.3%	34.2%	31.4%
Breakdown of Wealth of Couples:							
Home Equity	91892	105776	113968	142933	149007	108541	102634
Other Assets	247732	253485	225693	261094	306048	239540	252372
other real estate	29480	39093	29075	42415	54764	42212	33960
vehicles	24592	24093	24257	23544	22850	20580	21703
farms-businesses	53745	59818	55969	51621	74252	53620	54297
stocks-shares	55653	44376	37055	46719	47198	31978	34606
IRA-annuities	46895	50014	43123	49036	61905	44288	67739
savings accounts	20509	20151	21778	25394	29171	26339	26341
other assets	16858	15939	14437	22364	15909	20522	13725
Other Debts	9155	10594	11492	13153	13619	15263	15638

Notes: The first row presents average net worth of couples by year. The second and third rows present average net worth of single males and females respectively as a *percentage* of the net worth of couples in the same year. In the subsequent rows I break down *couples'* net worth to its elementary components. Net worth comprises home equity (house value net of any mortgages) and value of other assets net of other debts. Other assets comprise other real estate, vehicles, farms and businesses, stocks, shares and other investments, individual retirement accounts and annuities, savings accounts, and miscellaneous assets. Other debts comprise credit card debt, student loans, medical and legal bills, and loans to relatives, but excludes vehicle loans. All monetary figures are expressed in \$2010.

Information on assets is needed for the construction of the ‘partial insurance’ parameter π_{it} . The PSID collects data on home equity (house value net of mortgages), value of other real estate, vehicles, farms and businesses, shares, stocks and other investments, savings accounts and bond holdings, individual retirement accounts and annuities, and miscellaneous assets. Data on household debt are also collected including credit card debt, student loans, medical and legal bills, and loans to relatives. I am interested in the household net worth, therefore I aggregate the various asset components into one figure (‘wealth’) that captures total household assets and home equity net of outstanding debts (excluding vehicle loans).

Table 2.5 mimics the style of the consumption tables and presents summary statistics for assets and their components (expressed in \$2010). The average net worth of couples is positive at \$330,470 in 1998; this increases steadily until 2006 but drops sharply afterwards. The averages for singles are less reliable as, due to the small sample sizes, they are often impacted by extreme values. Nevertheless, net worth of singles appears to trend upwards with time possibly reflecting the shifting age composition of their samples.

2.4.2 Pre-Estimated Parameters

Three parameters pertaining to the log-linearization of the intertemporal budget constraint (appendix B.3) are estimated outside the structural model. These are

$$\begin{aligned}\xi_{it} &\approx \frac{\mathbb{E}_t [\text{Lifetime Expenditure on } K_{it}]}{\mathbb{E}_t [\text{Lifetime Total Expenditure}_{it}]} \\ s_{it} &\approx \frac{\mathbb{E}_t [\text{Lifetime Earnings}_{1it}]}{\mathbb{E}_t [\text{Lifetime Earnings}_{it}]} \\ \pi_{it} &\approx \frac{\text{Assets}_{it}}{\text{Assets}_{it} + \mathbb{E}_t [\text{Lifetime Earnings}_{it}]}\end{aligned}$$

where ‘Lifetime’ here spans the period between $t = 0$ and $t = T$. Similar parameters are also defined for singles. Identification of the household structure does not, however, require a log-linearization of the singles’ budget constraint

(section 2.3) and, as a result, I do not estimate the singles' counterparts of those parameters.

To deal with the expectations, first notice that

$$\mathbb{E}_t [\text{Lifetime Expenditure on } K_{it}] = K_{it} + \sum_{\varsigma=1}^T \frac{\mathbb{E}_t K_{it+\varsigma}}{(1+r)^\varsigma}.$$

I estimate $\mathbb{E}_t K_{it+\varsigma}$ by pooling K_{it} across all periods of time and regressing it on a set of predictable characteristics including each spouse's race, education, year of birth, a quartic polynomial in age, and a rich set of interactions. This regression can be written as $K_{it} = \mathbf{Q}_{it}^{k'} \boldsymbol{\beta}^k + \varepsilon_{it}^k$ where the notation is obvious. To obtain the time t -expected household public consumption, say, at $t+2$ ($\varsigma = 2$), I use $\mathbb{E}_t K_{it+2} = \mathbf{Q}_{it+2}^{k'} \hat{\boldsymbol{\beta}}^k$. I set the interest r at 2%. I repeat the same for private consumption C_{it} . It follows that $\mathbb{E}_t [\text{Lifetime Total Expenditure}_{it}] = \mathbb{E}_t [\text{Lifetime Expenditure on } K_{it}] + \mathbb{E}_t [\text{Lifetime Expenditure on } C_{it}]$.

As for s_{it} , notice that

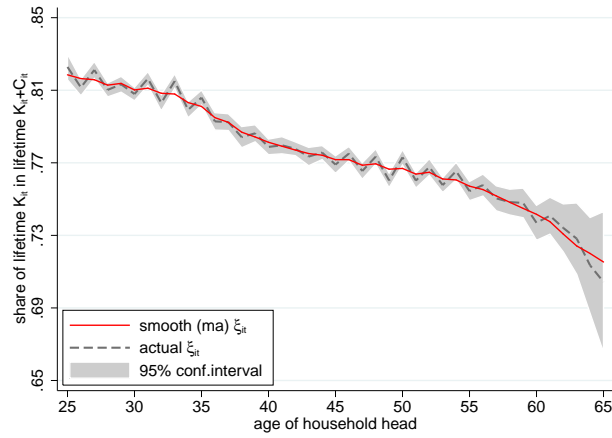
$$\mathbb{E}_t [\text{Lifetime Earnings}_{jit}] = Y_{jit} + \sum_{\varsigma=1}^T \frac{\mathbb{E}_t Y_{jit+\varsigma}}{(1+r)^\varsigma}.$$

I pool spouse j 's earnings over the years and I regress them on his/her race, education, year of birth, a quartic in age, and their interactions. This regression is given by $Y_{jit} = \mathbf{Q}_{jit}^{y'} \boldsymbol{\beta}_j^y + \varepsilon_{jit}^y$. Like before, I obtain j 's expected earnings, say, at $t+2$ ($\varsigma = 2$) setting $\mathbb{E}_t Y_{jit+2} = \mathbf{Q}_{jit+2}^{y'} \hat{\boldsymbol{\beta}}_j^y$. I repeat the same steps for each spouse's earnings separately; then $\mathbb{E}_t [\text{Lifetime Earnings}_{it}] = \sum_j^2 \mathbb{E}_t [\text{Lifetime Earnings}_{jit}]$.


Constructing π_{it} requires knowledge of beginning-of-period assets (section 2.4.1) and expected human wealth $\mathbb{E}_t [\text{Lifetime Earnings}_{it}]$. I assume assets at t are prior to any consumption-leisure choices in the same period.

Figure 2.2 plots average values of the (pre-estimated) parameters ξ_{it} , s_{it} , and π_{it} against the age of the male spouse. The grey dashed lines plot the average values; a 95% confidence interval around them appears in a grey shade.


Figure 2.2: Pre-estimated parameters



(a) ξ_{it}



(b) s_{it}



(c) π_{it}

Notes: These figures plot the average ξ_{it} , s_{it} , and π_{it} against the age of the household head. The red line depicts a 5-point moving average over the series (averages by age of household head), the grey dashed line depicts the actual means, and the grey shade is the 95% confidence interval around the actual mean.

The red solid lines plot a 5-point moving average around the average values in order to ease legibility of the graphs.

According to graph 2.2a, the share of expected lifetime public in total expenditure drops slightly over the life-cycle, from around .82 at age 25 to .72 by age 65. Graph 2.2b informs that men’s share of expected lifetime earnings is around .6 of combined lifetime earnings in the family over the first 2/3 of the life-cycle. Post age 50, however, this share drops, initially slowly but subsequently faster, to less than .5, reflecting the presence of a relatively younger spouse and, therefore, different retirement ages. Graph 2.2c illustrates that young households hold very little assets at the beginning of their life-cycle but π_{it} steadily increases in a convex way. By age 60, assets constitute approximately half of their financial and (remaining) human wealth combined. Note that graphs 2.2b-2.2c are very similar quantitatively to the ‘partial insurance’ graphs in Blundell et al. (2016).

2.4.3 Estimation Procedure

In this section I describe the steps I take to estimate the structural parameters and I discuss issues pertaining to measurement error in the data and inference. The ‘chronology’ of the model estimation is as follows. First, I clear wages, earnings, and consumption from the effect of observable characteristics. Subsequently, I estimate the parameters of the wage process. Finally, I estimate the parameters of the structural model conditional on the estimated wage parameters.

I regress $\ln W_{jit}$, $j = \{1, 2\}$, on a set of observable characteristics that include dummies for calendar year, age, education, race, state of residence, and education-year and race-year interactions, and I obtain the residual. I carry out such regression separately for men and women, married and singles. If wages suffer from measurement error, and if such error is classical (which implies, among others, that it is independent of the covariates), the wage residual is

$$\tilde{w}_{jit} = w_{jit} + e_{jit}^w$$

where w_{jit} is the error-free (economically relevant) residual and e_{jit}^w is the measurement error.

I stack together growth in residual wages \tilde{w}_{jit} from one period to another, that is $\Delta\tilde{\mathbf{w}}_{ji} = (\Delta\tilde{w}_{ji1999}; \Delta\tilde{w}_{ji2001}; \dots \Delta\tilde{w}_{ji2011})'$. Motivated by the theoretical representation (2.7), I use the second moments of $\Delta\tilde{\mathbf{w}}_{ji}$ to estimate the parameters of the gender-specific wage process, *i.e.* the second moments of wage shocks. I estimate those separately for married and singles. I use the GMM estimator and I weigh the moments by the identity matrix.

The measurement error presents a challenge as it is not possible to estimate the variance of the transitory shock separately from the variance of the error. To get around this, I remove a priori the variability in wages that is attributed to error using a well-known validation study for the PSID. Bound et al. (1994) compare interview responses and official records for a sample of workers in a single large manufacturing firm; they extrapolate their findings appropriately to representative samples and argue that measurement error is responsible for 7.2% to 16.2% of the variability in log hourly wages. I adopt an estimate in the middle of that range (13%; Blundell et al., 2016, use the same number too) and I assume that measurement error is serially uncorrelated as well as uncorrelated across partners.³²

I regress $\ln Y_{jit}$ on the same set of observables like above, as well as dummies for the number of children, number of household members, employment status at the time of the interview, for additional earners in the household other than the main two (main one for single-member households) and for outside recipients of financial support. I also include the first difference of those variables and employment status-year interactions. I carry out the regressions separately for married and singles; for married I also control for the characteristics of the spouse.

³²The major caveat using Bound et al. (1994)'s validation study is that their sample of workers comes from years 1982 and 1986, *i.e.* almost two decades before the bulk of the data I am using in this chapter. It is unclear how the importance of measurement error has changed over time or after the restructuring of the PSID in 1997. Another caveat comes from using the same estimates to correct female wages too even though the validation study sampled male workers only.

If earnings suffer from (classical) measurement error, the earnings residual is

$$\tilde{y}_{jit} = y_{jit} + e_{jit}^y$$

where y_{jit} is the error-free residual and e_{jit}^y is the measurement error. I stack together growth in residual earnings \tilde{y}_{jit} from one period to another and I obtain $\Delta\tilde{\mathbf{y}}_{ji}$. Bound et al. (1994) report that roughly 4% of the variability in log earnings is attributed to measurement error; I use this estimate to correct the second moments of $\Delta\tilde{\mathbf{y}}_{ji}$. I also remove the error variance from $E[\Delta\tilde{\mathbf{w}}_{ji}\Delta\tilde{\mathbf{y}}_{ji}]$. Given that log hourly wages are calculated as log annual earnings minus log annual hours

$$E[e_{jit}^w e_{jit}^y] = E[(e_{jit}^y - e_{jit}^h)e_{jit}^y] = E[e_{jit}^{y^2}] - E[e_{jit}^h e_{jit}^y]$$

and

$$E[e_{jit}^h e_{jit}^y] = 1/2 \left(E[e_{jit}^{y^2}] + E[e_{jit}^{h^2}] - E[e_{jit}^{w^2}] \right)$$

where e_{jit}^h is the measurement error in spouse j 's log annual working hours at t . Bound et al. (1994) report that between 17.9% and 26.6% of the variability in log annual hours is due to measurement error; I adopt an estimate in the middle of that range (23%). Again, I assume that the errors are serially uncorrelated and uncorrelated across partners.

Finally, I remove the effect of observables from private and public consumption regressing $\ln C_{it}$ and $\ln K_{it}$ respectively on the same set of covariates that I used in the earnings regressions above. If consumption suffers from (classical) measurement error, the estimated residuals are respectively

$$\tilde{c}_{it} = c_{it} + e_{it}^c$$

$$\tilde{k}_{it} = k_{it} + e_{it}^k$$

where c_{it} and k_{it} are the error-free residuals and e_{it}^c and e_{it}^k are measurement errors. I stack together growth in residual consumption from one period to

another to obtain $\Delta\tilde{\mathbf{c}}_i$ and $\Delta\tilde{\mathbf{k}}_i$ (with Δc_{i1999} and Δk_{i1999} both missing as consumption information was first collected in 1999). For singles these variables are all subscripted by $j = \{1, 2\}$ to indicate assignability of consumption. I identify the variance of the measurement error by the first-order auto-covariance of consumption given that the transmission of transitory shocks into consumption is identified by the covariance between consumption and wages.³³

I estimate the parameters of the structural model by assigning the empirical second moments of wages, earnings and consumption of couples to their theoretical counterparts, *i.e.* to the covariance of (2.9). For estimating the full version of the structural model (with nonseparable preferences and lack of commitment) I also use the empirical second moments of wages, earnings and consumption of singles.³⁴ I use GMM and the identity matrix as weights. For inference I adopt the block bootstrap (see for example Section 4 in Horowitz, 2001). I draw 1,000 random samples from the original samples (by marital status) and repeat all stages of the estimation for each bootstrap resample (*i.e.* first stage regressions for wages, earnings, consumption; GMM estimation of the parameters of the wage process; GMM estimation of the remaining parameters). I account in this way for arbitrary forms of heteroscedasticity across and serial correlation within blocks as well as the fact that I use pre-estimated residuals in the main GMM estimation.

2.5 Results

This section presents the main estimation results, namely the GMM estimates of the parameters of the wage process and of the household structure, including the gender-specific Frisch elasticities, the allocation of private consumption, and the bargaining effects due to lack of commitment.

I estimate the wage process of male and female earners imposing station-

³³I allow the variance of the measurement error in consumption to differ across couples, single men, and single women. See table B.1 in appendix B.5 for estimates of the variance.

³⁴I use the empirical variances and first order auto-covariances of the joint distribution of wages, earnings, and consumption over time. Higher-order auto-covariances are almost always insignificantly different from 0.

arity of the parameters over time. I present two sets of results: panel A of table 2.6 presents the estimates for married individuals whereas panel B presents the estimates for singles.

The variance of female shocks is effectively unchanged across states of life (across panels A and B); the variance of male shocks is, however, slightly lower among single men as opposed to men who are married. Men exhibit higher transitory variability/instability than women across both states of life possibly reflecting higher job mobility (Gottschalk and Moffitt, 2009). Overall, the relatively higher variance of women's permanent shocks may reflect wider dispersion of returns to skills and labor market experience compared to men.

The covariance of shocks between spouses is always positive, although this is estimated less precisely than the variance; it implies a correlation of $\rho_{v_1v_2} = 0.15$ for permanent shocks and $\rho_{u_1u_2} = 0.17$ for transitory shocks. These correlations suggest that, possibly due to positive assortative mating, the spouses work in similar industries or sectors, or their skills are similar, implying that they are hit by shocks that co-move.

Stationarity is not a restrictive assumption in practice. Although relaxing it renders the model more flexible, in practice the results change very little. The standard errors, however, are inflated due to the smaller sample sizes that are applicable per estimable parameter. These results are available upon request.

Subsequently, I present the estimates for the main set of model parameters. Following the illustration of identification in section 2.2.4, I present two sets of results. In section 2.5.1 the public good is additively separable in individual preferences; moreover spouses fully commit to each other for life. To estimate the (applicable) model parameters in this case I rely on wage estimates from panel A in table 2.6. In section 2.5.2 preferences are nonseparable and full commitment relaxed; but spousal preferences are assumed common among married and singles. To estimate the model parameters in this case I rely on wage estimates from both panels A and B in table 2.6.

Table 2.6: Estimates of wage parameters

	I. Men	II. Women	III. Family
		Panel A: <i>Married couples</i>	
Permanent:	$\sigma_{v_1}^2$ 0.0343 (0.0045)	$\sigma_{v_2}^2$ 0.0343 (0.0032)	$\sigma_{v_1v_2}$ 0.0051 (0.0020)
			$\rho_{v_1v_2}$ 0.1481 (0.0619)
Transitory:	$\sigma_{u_1}^2$ 0.0366 (0.0063)	$\sigma_{u_2}^2$ 0.0157 (0.0039)	$\sigma_{u_1u_2}$ 0.0040 (0.0023)
			$\rho_{u_1u_2}$ 0.1660 (0.1144)
		Panel B: <i>Single individuals</i>	
Permanent:	$\sigma_{v_1}^2$ 0.0258 (0.0081)	$\sigma_{v_2}^2$ 0.0335 (0.0087)	$\sigma_{v_1v_2}$ n/a (n/a)
			$\rho_{v_1v_2}$ n/a (n/a)
Transitory:	$\sigma_{u_1}^2$ 0.0211 (0.0080)	$\sigma_{u_2}^2$ 0.0173 (0.0094)	$\sigma_{u_1u_2}$ n/a (n/a)
			$\rho_{u_1u_2}$ n/a (n/a)

Notes: The table presents the GMM estimates of the parameters of the wage process under stationarity. Block bootstrap standard errors are in parentheses based on 1,000 bootstrap replications. Panel A presents the estimates for married spouses; panel B presents the estimates for single individuals.

2.5.1 Results with Separable Public Good and Full Commitment

The results appear in columns 1 across blocks I-III of table 2.7. The own-wage labor supply elasticity, at $\eta_{1,h,w} = 0.349$ for men and $\eta_{2,h,w} = 0.992$ for women, is within the range of other studies (and consistent with them, women's elasticity is higher than men's; for a review see Keane, 2011).

The cross-elasticity of labor supply with respect to the price of the private good, at $\eta_{1,h,p^c} = 0.016$ for men and $\eta_{2,h,p^c} = 0.039$ for women, is statistically insignificant but positive, implying that market hours and private consumption are Frisch substitutes at the intensive margin of labor supply (leisure and private consumption are Frisch complements).³⁵ The 'reciprocal' cross-elasticities, of private consumption with respect to wages, weighed by the average private consumption share of each spouse, are effectively the same between men and women. Note that Frisch symmetry *cannot* be used to infer the consumption share *given* η_{1,h,p^c} and η_{2,h,p^c} , because Frisch symmetry *is* already used to identify η_{1,h,p^c} and η_{2,h,p^c} in the first place.

The *household-level* private consumption substitution elasticity $\tilde{\eta}_{c,p^c} = \mathbb{E}[\varphi_{it-1}] \eta_{1,c,p^c} + \mathbb{E}[1 - \varphi_{it-1}] \eta_{2,c,p^c}$ is estimated at -0.331 whereas the *household-level* public consumption substitution elasticity

$$\tilde{\eta}_{k,p^k} = \mathbb{E} \left[\eta_{1,k,p^k} \eta_{2,k,p^k} / \left((1 - \nu_{it-1}) \eta_{1,k,p^k} + \nu_{it-1} \eta_{2,k,p^k} \right) \right]$$

at -0.645 . Given the average shares of public and private consumption over the period 1999-2011, these numbers imply a total consumption substitution elasticity, *i.e.* a parameter estimated in standard unitary models, approximately equal to -0.55 (similar estimate found in Blundell et al., 2016).

³⁵Blundell et al. (2016) also find evidence of Frisch complementarity between spousal leisure and consumption. Theirs is a study using the unitary model though.

Table 2.7: Estimates of Frisch elasticities

	I. Men			II. Women	
	(1) separable	(2) nonseparable		(1) separable	(2) nonseparable
$\eta_{1,h,w}$	0.3493 (0.0629)	0.3390 (0.1158)	$\eta_{2,h,w}$	0.9923 (0.1791)	0.7340 (0.2963)
η_{1,h,p^c}	0.0162 (0.0141)	0.0004 (0.0128)	η_{2,h,p^c}	0.0391 (0.0394)	0.0134 (0.0402)
η_{1,h,p^k}		0.2822 (0.1139)	η_{2,h,p^k}		0.3804 (0.4697)
$\eta_{1,c,w}$		-0.0027 (0.0874)	$\eta_{2,c,w}$		-0.0390 (0.1164)
$\mathbb{E}[\varphi_1]$ $\times \eta_{1,c,w}$	-0.0550 (0.0481)		$\mathbb{E}[\varphi_2]$ $\times \eta_{2,c,w}$	-0.0568 (0.0572)	
η_{1,c,p^c}		-0.4209 (0.2045)	η_{2,c,p^c}		-0.4209 (0.2045)
η_{1,c,p^k}		0 (n/a)	η_{2,c,p^k}		0 (n/a)
$\eta_{1,k,w}$		-0.3633 (0.1461)	$\eta_{2,k,w}$		-0.2329 (0.2885)
η_{1,k,p^c}		0 (n/a)	η_{2,k,p^c}		0 (n/a)
η_{1,k,p^k}		-1.5386 (0.8064)	η_{2,k,p^k}		-1.5386 (0.8064)
III. Family					
	(1) separable	(2) nonseparable			
$\tilde{\eta}_{k,p^k}$	-0.6454 (0.1098)	-1.5386 (n/a)			
$\bar{\eta}_{c,p^c}$	-0.3306 (0.0763)	-0.4209 (n/a)			

Notes: The table presents the GMM estimates of gender-specific Frisch elasticities when the public good is additively separable from private consumption and leisure (columns 1) and when it is nonseparable (columns 2). Block bootstrap standard errors are in parentheses based on 1,000 bootstrap replications. The full set of Frisch elasticities are defined in table 2.1. φ_1 is men's share of private consumption in the household and $\varphi_2 = 1 - \varphi_1$ is women's.

2.5.2 Results with Nonseparable Public Good and Lack of Commitment

The results appear in columns 2 across blocks I-III of table 2.7. A number of restrictions are imposed on the model parameters that deviate from (and are stronger than) the set of restrictions, described in section 2.2.4.2, that suffice for identification of the household structure. Specifically, i) I enforce a separability between the public and the private good, meaning that $\eta_{j,c,p^k} = \eta_{j,k,p^c}$, $j = \{1, 2\}$; ii) I impose equal intra-family bargaining power $\mu_{1it} = \mu_{2it} = 1/2$, $\forall i, t$, and $\nu_{it} = 1/2$ ($\forall i, t$; recall that $\nu \in [0, 1]$ reflects a mixture of preferences and bargaining power); iii) I impose equal sharing of private consumption between spouses, that is $\varphi_{it} = 1/2$, $\forall i, t$; and iv) I restrict the consumption substitution elasticities to be the same across spouses, that is $\eta_{1,c,p^c} = \eta_{2,c,p^c}$ and $\eta_{1,k,p^k} = \eta_{2,k,p^k}$. These restrictions serve the purpose of obtaining a robust but meaningful first set of estimates and will be relaxed in further versions of this work.³⁶

The own-wage labor supply elasticity, at $\eta_{1,h,w} = 0.339$ for men and $\eta_{2,h,w} = 0.734$ for women, is again within the range of other studies (Keane, 2011). Note that men's elasticity changes very little from when preferences are separable (however, the standard error is almost doubled). The cross-elasticity of hours with respect to the price of the private good, at $\eta_{1,h,p^c} = 0.0004$ for men and $\eta_{2,h,p^c} = 0.013$ for women, is statistically insignificant but, albeit marginally, positive. The cross-elasticity of hours with respect to the price of the public good, at $\eta_{1,h,p^k} = 0.282$ for men and $\eta_{2,h,p^k} = 0.380$ for women, is positive, implying that market hours and public consumption are Frisch substitutes at the intensive margin of labor supply, and, at least for men, statistically significant. Across all prices, and consistent with previous evidence, women's labor supply is more elastic than men's.

The cross-elasticity of private consumption with respect to the wage is

³⁶The results presented in this section reflect the specific restrictions currently imposed on the parameters. These results will likely change while I implement the full estimation in future versions of this work.

negative and statistically insignificant as expected by being a function of the hours elasticity with respect to the price of the private good. I estimate the private consumption substitution elasticity, that is the elasticity of private consumption with respect to its price η_{j,c,p^c} , $j = \{1, 2\}$, at -0.421 (and statistically significant).

The cross-elasticity of public consumption with respect to the wage is negative and, at least for men, statistically significant as expected by being a function of the hours elasticity with respect to the price of the public good. The public consumption substitution elasticity, that is the elasticity of public consumption with respect to its price η_{j,k,p^k} , $j = \{1, 2\}$, is estimated at -1.539 (and marginally significant).

As a means to compare the results with the estimated parameters when the public good is separable and full commitment imposed (section 2.5.1), I construct the implied *household-level* consumption substitution elasticities; $\tilde{\eta}_{c,p^c} = -0.421$ is close to its counterpart in the former case; however $\tilde{\eta}_{k,p^k} = -1.539$ is nearly 2.5 times larger (in absolute value) than its counterpart in the former case. The presence of bargaining effects, or a poor choice of restrictions (see beginning of this section), may explain this discrepancy.

Overall, men's Frisch elasticities seem to be estimated more precisely than women's. This is an interesting finding that is also confirmed by the main estimation results (again on Frisch elasticities) of Blundell et al. (2016).

Table 2.8 presents the estimates of the (average) bargaining effects induced by permanent shocks on household outcomes through the dependence of intra-family bargaining power (2.6) on wages. Column 1 presents the bargaining effects from men's permanent shock v_1 whereas column 2 presents the effects from women's permanent shock v_2 .

Two things are worth noting. First, all bargaining effects, as well as the 'surplus extraction' elasticities, are insignificant from a statistical point of view. This suggests that the spouses do not engage in reallocation of intra-family power when permanent shocks hit. This contradicts Mazzocco (2007)

Table 2.8: Bargaining effects

	(1) from v_1	(2) from v_2
on y_1	0.0681 (0.0692)	-0.0939 (0.0843)
on y_2	-0.1611 (0.1632)	0.2223 (0.1731)
on c	0.0043 (0.0173)	-0.0060 (0.0252)
on k	-0.0156 (0.0800)	0.0215 (0.1098)

Notes: The table presents the GMM estimates of the bargaining effects induced by permanent shocks. Block bootstrap standard errors are in parentheses based on 1,000 bootstrap replications. Column 1 presents the effects from men's permanent shocks; column 2 presents the effects from women's permanent shocks.

and serves as evidence against lack of commitment.

Second, the bargaining effects on earnings have the opposite signs than one would expect. One would expect a *negative* bargaining effect from one's own wage on one's own earnings and a *positive* on their partner's earnings. The idea is the following: as one's wage increases permanently ($v_j > 0$), one should, if anything, gain in bargaining power. Ceteris paribus, this should result in reducing one's market hours in order to enjoy more leisure and, therefore, it would also result in a reduction in one's earnings. The opposite would hold for their partner's market hours and earnings. As one becomes relatively more powerful due to a positive wage shock, their partner becomes relatively less powerful and, therefore, increases their market hours and reduces their leisure. In practice, however, a positive male permanent shock of 0.1 (10% permanent wage increase) increases his earnings by 0.068×0.1 (0.68% increase) whereas it reduces his wife's earnings by -0.161×0.1 (1.61% decrease). A positive female permanent shock of 0.1 increases her earnings by 0.222×0.1 (2.22% increase) whereas it reduces her husband's earnings by -0.094×0.1 (0.94%

decrease).

The bargaining effects on consumption are harder to sign. The estimates suggest that the empowerment of women in the household results in a reduction in private consumption and a boost in public consumption. The opposite holds for men.

The reason why the earnings bargaining effects have opposite signs is attributed to the signs of the ‘surplus extraction’ elasticities η_{μ,w_j} . These are estimated at $\eta_{\mu,w_1} = -0.239$ (*s.e.* = 0.290) and $\eta_{\mu,w_2} = 0.330$ (*s.e.* = 0.345). Note that the reason why I identify these parameters is because I set the Pareto weight to $1/2$. Their signs are opposite than one would expect (they imply that a positive permanent shock makes its recipient less powerful in their household); flipping their signs would mechanically reverse the signs of the bargaining effects. A possible reason why these elasticities end up having the ‘wrong’ signs is the definition I employ for the consumption substitution elasticities. I define η_{j,c,p^c} and η_{j,k,p^k} , $j = \{1, 2\}$, as negative whereas Blundell et al. (2016) define them as positive; flipping their sign would likely result in flipped signs for η_{μ,w_j} and the bargaining effects too. Another reason may pertain to the presence of cross-sectional heterogeneity in the bargaining effects that may imply that their unconditional first moment is a poor choice of parameter to estimate.

2.6 Discussion and Conclusions

This chapter presents and estimates a life-cycle collective household model with the aim of understanding how shocks to spousal wages transmit into earnings and household consumption. The model allows for a number of important features such as intra-family correlation of wages, wage shocks of various persistence, the distinction of consumption to public and private, asset accumulation, endogenous labor supply, and lack of spousal commitment to future allocations. Specifically, permanent wage shocks enter the spouses’ outside options and, therefore, also affect their intra-family bargaining powers.

I obtain analytical expressions for earnings and household consumption applying Taylor approximations to the problem's first order conditions and the lifetime budget constraint. I show how, using those expressions and information on singles of either gender, one can identify the household structure, namely a rich set of spouse-specific Frisch elasticities for labor supply and consumption, the unobserved sharing of private consumption between spouse, as well as the bargaining effects of wages. Importantly, information on assignable or exclusive consumption is not required.

Using PSID data post 1999 I find sizeable Frisch labor supply as well as consumption substitution elasticities. Consistent with the bulk of previous empirical evidence, women's labor supply is considerably more elastic than men's. The public consumption elasticities are up to two and a half times as large as the private consumption ones. The evidence on the cross-elasticities points to market hours being Frisch substitutes to all types of consumption at the intensive margin of labor supply. The bargaining effects are economically, but not statistically, significant although they appear to have the opposite sign than one would expect.

The chapter, at its current state, is subject to a number of limitations. The signs of the bargaining effects remain a puzzle that must be better understood. At the same time, the current estimation results reflect a much stricter set of restrictions than what suffices for identification and an important parameter, the sharing of private consumption between spouses, remains fixed at an arbitrary $1/2$ in the cross-section and over time. Subsequently, robustness checks are missing such as checks involving alternative categorizations within private and public consumption, or checks involving a proper treatment of the selection into market work (which, at least for women, may be important). All these are issues that deserve further investigation.

Chapter 3

Consumption Inequality across Heterogeneous Families

3.1 Introduction

This chapter studies the transmission of income shocks into consumption across households that exhibit unobserved preference heterogeneity. The chapter develops a tractable life-cycle model for household consumption, savings, and labor supply, and presents theoretical conditions with straightforward empirical counterparts that deliver identification of the cross-sectional distribution of household preferences if panel data on consumption, hours, and earnings are available. Importantly, identification does not depend on any specific parametrization of household preferences or their distribution.

A consistent empirical finding is that consumption inequality across households appears significantly lower than income inequality (Blundell and Preston, 1998; Blundell and Etheridge, 2010; Heathcote et al., 2010). This finding holds across different definitions of income and earnings, even for earnings post benefits and transfers. This discrepancy is usually attributed to consumption insurance or consumption smoothing in the household.

A large literature in macroeconomics and labor economics focuses on the mechanisms transmitting idiosyncratic income or wage changes into consumption within the household. Krueger and Perri (2006), Blundell et al. (2008),

Kaplan and Violante (2010) are recent contributions investigating the link between exogenous household income and consumption. Attanasio et al. (2002) model a two-earner household where the participation of the second earner in the labor market follows a Markov stochastic process. Hyslop (2001) investigates the link between wage and earnings inequality; he abstracts from consumption but focuses explicitly on endogenous family labor supply that is a potentially crucial insurance mechanism against shocks. This point is also made by Blundell et al. (2016) who study the transmission of wage shocks into consumption through a model of endogenous family labor supply, savings, and external insurance.

In a general equilibrium framework, Heathcote et al. (2014) model the labor supply of one earner in the household and implement an ‘insurance dichotomy’ whereby some income shocks are perfectly insurable whereas other are uninsured. Attanasio and Pavoni (2011) model labor supply (effort of one earner) jointly with consumption and allow households to possess private information about their effort and savings. Jappelli and Pistaferri (2010) and Meghir and Pistaferri (2011) provide an overview of the extensive literature.

With the exception of Heathcote et al. (2014), a consistent feature of this literature is that households are assumed *ex ante* identical. Conditional on their observable characteristics and idiosyncratic incomes/wages, any two households are modeled to behave the same when a given shock hits them. This is a poor feature especially in the light of extensive evidence of heterogeneity across individuals or households in consumer demand analysis and labor supply studies.¹ Heathcote et al. (2014) admit that part of the dispersion in consumption and hours in the cross-section is unrelated to income or price variation and allow for unobserved preference heterogeneity. Their treatment of heterogeneity is, however, restrictive as heterogeneity is, in effect, additively separable and specific to the parametrization of household preferences

¹For example, Abowd and Card (1989) find sizeable cross-sectional dispersion in working hours even at fixed wage rates. Blundell et al. (2007) and Blundell and Stoker (2007) provide reviews of the treatment of heterogeneity in labor supply or consumer demand studies.

that they employ.

The present chapter studies the transmission of idiosyncratic earnings/wage shocks into consumption across two-member unitary households that exhibit unobserved preference heterogeneity. The treatment of heterogeneity is general for: (i) within-period household preferences are nonseparable from such heterogeneity; (ii) household preferences are semi-nonparametric (nonparametric up to a separability between consumption and leisure; the separability is not needed for any of the results in this chapter but simplifies the illustration); (iii) heterogeneity is not restricted to a single dimension (to a single preference parameter in the analog of parametric preferences); instead it is multi-dimensional meaning that *any* within-period preference parameter in the analog of parametric preferences might exhibit unobserved heterogeneity; (iv) the multi-variate distribution of preferences is nonparametric and unrestricted (up to the requirement of finite first and higher moments).

The specific workings of the household environment are the following: I model life-cycle choices over consumption and female labor supply for a household of two adult members (for example, two spouses). The household can borrow and save at a deterministic interest rate. The primary earner (suppose the male spouse) works fixed hours in the labor market and receives full-time earnings that are subject to productivity shocks. The secondary earner (suppose the female spouse) has an endogenous choice over working hours (in the intensive margin) and, for each one hour of work, receives a wage that is subject to productivity shocks. Each spouse's productivity shocks are decomposed into permanent and transitory components. Such shocks are the only source of uncertainty for the household members.

How do I solve the model? Is a model with such general heterogeneity empirically useful? I follow Blundell and Preston (1998) and a sequence of papers thereafter (Attanasio et al., 2002; Blundell et al., 2008, 2016) and I apply first- and second-order Taylor approximations to the lifetime budget constraint and the problem's first-order conditions. Such approximations en-

able me to obtain analytical expressions for consumption and labor supply as functions of the productivity shocks, household idiosyncratic preferences (namely marginal-utility-constant or Frisch elasticities), and a number of ‘partial insurance’ parameters (term due to Blundell et al., 2008). The analytical expressions are convenient because they provide a neat picture for the contribution of different components of the model to consumption and hours growth, and, therefore, a straightforward way to identify the household structure.

The second and higher moments of the empirical joint distribution of consumption, earnings, and hours growth have, thanks to the aforementioned analytical expressions, a clear theoretical counterpart. This mapping between data and theory provides restrictions that can be used to identify second and higher moments of the cross-sectional distribution of household preferences.

I show how, effectively, *any* moment of the distribution of female (Frisch, own-wage) labor supply elasticities can be identified from panel data on hours and wages.² Recovering the said moments provides a straightforward and empirically-feasible test for unobserved preference heterogeneity in labor supply. In the general case with preferences nonseparable between consumption and leisure, any moment of the distribution of consumption (Frisch) elasticities with respect to female wages can also be identified. However, lack of variation in the real price of consumption prohibits identification of the distribution of (consumption, hours) elasticities with respect to the price of consumption, including the consumption substitution elasticity.

The chapter contributes to the literature that studies the link between idiosyncratic income changes and consumption (see above). Given the precise method I follow for solving the household problem, this chapter can be seen as complementary to Blundell and Preston (1998), Attanasio et al. (2002), Blundell et al. (2008), and Blundell et al. (2016). None of these papers allow for unobserved preference heterogeneity, thus they effectively enforce a rep-

²The Michigan Panel Study of Income Dynamics is a widely used panel dataset with detailed information on wages, incomes, employment and labor supply of household heads and their spouses. Since 1999 it also provides detailed information on household expenditure and assets, albeit on a biennial basis.

representative agent assumption (conditional on observables). By contrast, the present chapter: (i) embeds a general form of unobserved preference heterogeneity across households, and (ii) provides moment restrictions that identify the distribution of and test for such heterogeneity. However, unlike Blundell et al. (2016), I impose an exogenous labor supply for the primary earner and separable preferences between consumption and leisure. This is so as to facilitate the illustration herein and both restrictions can be relaxed without jeopardizing the results of the chapter.

Finally, although this chapter is not a demand analysis study, it does share a common goal with the extensive literature on consumer demand, namely the identification of preferences from observed behavior. Lewbel (2001) studies various forms of random preferences, with and without nonseparable heterogeneity, and argues that restricting heterogeneity to additive errors comes close to imposing a representative agent assumption. In light of this, a number of recent consumer demand studies present identification results when preferences exhibit nonseparable unobserved heterogeneity (like in this chapter); examples are Matzkin (2003, 2008), Blundell et al. (2013, 2014), and Cosaert and Demuyneck (2014).

The chapter is organized as follows. Section 3.2 presents the life-cycle household model and presents analytical expressions for household consumption, hours, and earnings, as well as expressions for cross-sectional inequality in such outcomes. Section 3.3 discusses the identification of the parameters of the income process and the household structure, namely the distribution of preferences. Section 3.4 presents simple tests for unobserved preference heterogeneity and the model specifications. Section 3.5 concludes.

3.2 A Life-Cycle Household Model for Consumption and Labor Supply

This section has two parts. In the first part I develop a life-cycle model for household consumption, savings, and female labor supply, which exhibits unob-

served preference heterogeneity. In addition, I overview the method I employ to solve the model, namely the Taylor approximations to the intertemporal budget constraint and the problem's first order conditions. In the second part I present a closed-form system of equations for household consumption, earnings, and female labor supply as well as closed-form equations for the dynamics of consumption and earnings inequality in the cross-section.

3.2.1 The Model

A household consists of two earners, each one subscripted by j . To fix ideas suppose the two earners are a male ($j = 1$) and a female ($j = 2$) spouse, although the model applies equally to any modern or traditional form of cohabitation. In any period t , the spouses make choices regarding total household consumption C_t , their future assets A_{t+1} , and female hours of work in the labor market H_{2t} (intensive margin labor supply only). The male spouse is assumed to work always full-time, thus he has no labor supply choice.

I assume the spouses stay together and commit to one another for life (the length of which is T). I model the household problem as unitary, that is as the problem of a single economic agent. This facilitates the discussion of cross-sectional preference heterogeneity without confounding it with issues pertaining to *intra*-household heterogeneity, commitment, or intra-household inequality.

A household i in the cross-section chooses $\{C_{it}, A_{it+1}, H_{2it}\}$ over its life-cycle to maximize its expected discounted lifetime utility

$$\max_{\{C_{it}, A_{it+1}, H_{2it}\}_{t=0}^T} \mathbb{E}_0 \sum_{t=0}^T \beta^t [U_i(C_{it}, \mathbf{Z}_{it}) - V_i(H_{2it}, \mathbf{Z}_{it})] \quad (3.1)$$

subject to an intertemporal budget constraint, the sequential version of which at time t is

$$A_{it} + Y_{1it} + W_{2it}H_{2it} = C_{it} + \frac{A_{it+1}}{1+r}.$$

In the budget constraint, A_{it} is beginning-of-period household assets, Y_{1it}

is the male spouse's full-time earnings, W_{2it} is the female's hourly wage rate, and r is the deterministic market interest rate.

In the objective function, U_i is the household utility from consumption and V_i is the household disutility from female market work.³ β is the geometric discount factor and, for simplicity, is assumed the same across households. Vector \mathbf{Z}_{it} includes observable taste shifters such as education or age (thus, it captures *observed* preference heterogeneity). Household preferences U_i and V_i are all subscripted by i to indicate unobserved preference heterogeneity across households, strictly speaking household-specific preferences not captured by the conditioning observed taste shifters. This is a general way to model such heterogeneity and is consistent with various different sources that preference heterogeneity may stem from, such as cross-household differences in unobserved costs of work, in consumption substitution elasticities, or in intra-family bargaining powers (for example, in the context of a full commitment life-cycle collective model). I do not parameterize U_i or V_i but I do require they have continuous first- and second-order derivatives.

Male earnings Y_1 and female wages W_2 are the primitive sources of labor market uncertainty for the household and are assumed exogenous to household choices. I model them using a permanent-transitory process which is a prominent representation in the income dynamics literature (see, for example, Meghir and Pistaferri, 2004).⁴ For men, I decompose residual log earnings into

³If male labor supply was an endogenous choice, then the household (static) objective would be $U_i(C_{it}, \mathbf{Z}_{it}) - V_{1i}(H_{1it}, \mathbf{Z}_{it}) - V_{2i}(H_{2it}, \mathbf{Z}_{it})$, where H_{1it} would be the male spouse's hours of market work and V_j the corresponding disutility from work, $j = \{1, 2\}$.

⁴The permanent-transitory process for income, wages or earnings, has been used extensively in the literature, for example in MaCurdy (1982), Abowd and Card (1989), Attanasio et al. (2002), Attanasio et al. (2008), Blundell et al. (2008), or Blundell et al. (2016). The permanent-transitory process results in, so-called, 'restricted income profiles', meaning that different economic subjects have the same income profiles conditional on observables but differ with respect to the idiosyncratic shocks they are hit by. An alternative family of income processes in the literature supports individual-specific life-cycle income profiles, even conditional on observables, and results in, so-called, 'heterogeneous income profiles' (see, for example, Browning et al., 2010; Guvenen, 2007).

the sum of a permanent and a transitory earnings shock

$$\begin{aligned}\ln Y_{1it} &= X'_{Y_1,it} \boldsymbol{\alpha}_{Y_1} + \ln Y_{1it}^p + u_{1it} \\ \ln Y_{1it}^p &= \ln Y_{1it-1}^p + v_{1it}\end{aligned}$$

where $\ln Y_{1it}^p$ is the permanent component of (log) earnings, $X'_{Y_1,it}$ is a vector of covariates (such as age or education) and $\boldsymbol{\alpha}_{Y_1}$ is the corresponding coefficient, u_{1it} is the transitory shock, and v_{1it} is the permanent shock.

Similarly, I decompose residual log wages of women

$$\begin{aligned}\ln W_{2it} &= X'_{W_2,it} \boldsymbol{\alpha}_{W_2} + \ln W_{2it}^p + u_{2it} \\ \ln W_{2it}^p &= \ln W_{2it-1}^p + v_{2it}.\end{aligned}$$

The notation is similar; note that shocks subscripted by $_1$ indicate shocks to *earnings* (permanent or transitory) whereas shocks subscripted by $_2$ indicate *wage* shocks. The latter induce income and substitution effects on female labor supply whereas the former only induce income effects.

The income processes can be written more compactly

$$\Delta F_{jit} = v_{jit} + \Delta u_{jit} \tag{3.2}$$

where $\Delta F_{jit} = \Delta \ln Y_{1it} - \Delta X'_{Y_1,it} \boldsymbol{\alpha}_{Y_1} \equiv \Delta y_{1it}$ for men ($j = 1$) and $\Delta F_{jit} = \Delta \ln W_{2it} - \Delta X'_{W_2,it} \boldsymbol{\alpha}_{W_2} \equiv \Delta w_{2it}$ for women ($j = 2$).

The permanent shock reflects a permanent change in the returns to one's skills in the labor market whereas the transitory shock indicates short-lived mean reverting fluctuations in productivity such as fluctuations in effort when effort is observed and tied to one's wage. I refrain from labeling the permanent shock as technical change as that would be aggregate in nature and the model does not account for aggregate shocks.

Assumption (Properties of shocks). *Earnings and wage shocks are idiosyn-*

cratic in nature with zero cross-sectional means, second moments given by

$$\begin{aligned}\mathbb{E}(v_{jit}v_{kit+s}) &= \begin{cases} \sigma_{v_j,t}^2 & \text{if } j = k \text{ and } s = 0 \\ 0 & \text{otherwise} \end{cases} \\ \mathbb{E}(u_{jit}u_{kit+s}) &= \begin{cases} \sigma_{u_j,t}^2 & \text{if } j = k \text{ and } s = 0 \\ 0 & \text{otherwise} \end{cases} \\ \mathbb{E}(v_{jit}u_{kit+s}) &= 0 \quad \forall j, k, s\end{aligned}$$

and third moments by

$$\begin{aligned}\mathbb{E}(v_{jit}^2v_{kit+s}) &= \begin{cases} \gamma_{v_j,t} & \text{if } j = k \text{ and } s = 0 \\ 0 & \text{otherwise} \end{cases} \\ \mathbb{E}(u_{jit}^2u_{kit+s}) &= \begin{cases} \gamma_{u_j,t} & \text{if } j = k \text{ and } s = 0 \\ 0 & \text{otherwise} \end{cases} \\ \mathbb{E}(v_{jit}^2u_{kit+s}) &= 0 \quad \forall j, k, s \\ \mathbb{E}(v_{jit}u_{kit+s}^2) &= 0 \quad \forall j, k, s\end{aligned}$$

where $j = \{1, 2\}$, $k = \{1, 2\}$, $s = \{0, \dots, T\}$, and $\mathbb{E}(\cdot)$ denotes the mean over i . The spouses hold no advance information about future shocks.

I allow the second and third moments of shocks to vary with time as different time periods can be associated with different amounts of income volatility or skewness. As life-cycle effects are captured by the conditioning observables (age), the time dependence of moments indicates calendar-time effects such as a financial crisis leading to bigger amounts of income variability in certain years. I assume shocks between spouses are independent. This is a strong assumption and implies no assortative mating, positive or negative, between the two. Relaxing it will not alter the chapter's conclusions regarding the implications of unobserved preference heterogeneity; it would, however, complicate the solution of the model and the derivation of the closed-form expressions of

section 3.2.2.

I solve the model by obtaining the first-order conditions of (3.1) and by applying first- and second-order Taylor approximations to them. These approximations produce closed-form expressions for household consumption, female hours of work, and the marginal utility of wealth given by

$$\begin{aligned}\Delta c_{it} &\approx \phi_i \Delta \ln \lambda_{it} \\ \Delta h_{2it} &\approx \psi_i (\Delta \ln \lambda_{it} + \Delta w_{2it}) \\ \Delta \ln \lambda_{it} &\approx \varepsilon_{it} + \omega_{it}.\end{aligned}\tag{3.3}$$

The notation is as follows: $\Delta c_{it} = \Delta \ln C_{it}$ net of observables such as age or education; $\Delta h_{2it} = \Delta \ln H_{2it}$ net of the same observables; λ_{it} is the Lagrange multiplier on the household budget constraint (the marginal utility of wealth). ε_{it} is an innovation term that captures idiosyncratic revisions to the marginal utility of wealth when income shocks hit. ω_{it} is a function of the discount factor, the interest rate and the variance of $\Delta \ln \lambda_{it}$ and captures household precautionary motives over time. Appendix C.1 provides the details of these approximations.

Parameters ϕ_i and ψ_i load income shocks (or innovations to the marginal utility of wealth) onto household consumption and female labor supply. Specifically, $\phi_i = \frac{U'_i(C_{it})}{C_{it}U''_i(C_{it})}$ is household i 's consumption substitution elasticity (the Frisch elasticity of consumption with respect to its price); $\psi_i = -\frac{V'_i(H_{2it})}{H_{2it}V''_i(H_{2it})}$ is household i 's labour supply elasticity (the Frisch elasticity of hours with respect to own wage).⁵ It is easy to see why ϕ_i and ψ_i serve as loading factors of income shocks. For instance, female hours of work respond to a transitory wage shock so long as female labor supply is elastic and $\psi_i > 0$. In this case the response of female labor supply is a proportion ψ_i of the wage shock. If women's labor supply is perfectly inelastic, *i.e.* if $\psi_i = 0$, then $\Delta h_{2it} = 0$ and a wage shock will leave female labor supply unchanged.

There is a distribution of consumption substitution and labor supply elas-

⁵ U'_i and V'_i denote first-order derivatives; U''_i and V''_i denote second-order derivatives.

ticities in the cross-section as preferences in this model are household-specific due to unobserved heterogeneity. I am interested in recovering moments of the cross-sectional distribution of preferences, such as the mean $\mathbb{E}(\cdot)$ and the variance $\text{Var}(\cdot)$ of both ϕ_i and ψ_i . For simplicity, I assume that the distribution of preferences does not change with time but note that this assumption is not actually needed for the identification of such distributions.

Although Δc_{it} , Δh_{2it} and spousal incomes Δy_{1it} and Δw_{2it} are all directly observed in standard household panel surveys, such as the PSID after 1999, $\Delta \ln \lambda_{it}$ is not. That renders the system of equations (3.3) empirically unattractive. Following Blundell and Preston (1998), Attanasio et al. (2002) and a sequence of papers thereafter, I apply a first-order Taylor approximation to the intertemporal budget constraint around the path the household would follow if income shocks were absent and I take expectations before and after the income shocks are realized. This way I relate the innovation ε_{it} to spousal earnings and wage shocks. This approach draws on Campbell (1993)'s log-linear approximation to the intertemporal budget constraint.⁶

The details of this derivation appear in appendix C.2. There I show how I approximate the innovation to the marginal utility of wealth by a linear function of permanent earnings and wage shocks, namely by

$$\varepsilon_{it} \approx \frac{1}{\phi_i - (1 - s_{it})\psi_i} (s_{it}v_{1it} + (1 - s_{it})(1 + \psi_i)v_{2it}). \quad (3.4)$$

Here s_{it} is approximately equal to the average expected share of *male* earnings in *family total* earnings over the period between t and T (end of lifetime). It follows that $1 - s_{it}$ is the average expected share of female earnings over the same horizon. Standard panel surveys usually provide adequate information to infer s_{it} ; for example, Blundell et al. (2016) use the PSID to construct measures of s_{it} per household and time period.

Two important assumptions help get expression (3.4). These are:

⁶Also see Blundell et al. (2013) for a detailed illustration of this approach.

Assumption (No wealth effects from transitory shocks). *Transitory income (earnings, wage) shocks are mean-reverting and do not shift the intertemporal budget constraint. As a result, transitory income shocks do not affect the innovation ε_{it} to the marginal utility of wealth and, therefore, do not induce wealth effects on consumption or female labor supply.*

Assumption (Negligible assets). *Household assets at any time t are negligible compared to household expected lifetime resources (lifetime financial and human wealth).*

The first assumption implies that transitory income shocks induce standard income and substitution effects on choices, but not wealth effects through the intertemporal budget constraint. Although this is not an unreasonable assumption, at least for the most common transitory income shocks, relaxing it would likely complicate, if not jeopardize, identification of the cross-sectional distribution of preferences. The second assumption is likely to be true if households are sufficiently young and have not accumulated significant assets yet. Relaxing it does not jeopardize identification or tractability of the approximation; instead it introduces a ‘partial insurance parameter’ π_{it} in (3.4) (term due to Blundell et al. (2008); more on this in appendix C.2) which can be inferred directly from the data (as in Blundell et al., 2016).

Relation (3.4) is informative about how permanent shocks impact the revisions to the marginal utility of wealth. A permanent shock v_{1it} to male earnings is transmitted into ε_{it} weighed by s_{it} because male earnings are only a fraction s_{it} of total family earnings. In the extreme when $s_{it} \approx 0$ an earnings shock v_{1it} has no effect. Similarly, a permanent shock v_{2it} to female wages is transmitted into ε_{it} weighed by $1 - s_{it}$. This wage shock, however, is further mitigated by $1 + \psi_i$ as the female spouse can adjust her labour supply in response to such a shock. For example, in the case of a positive permanent wage shock, she can take advantage of the higher wage by increasing her working hours and thus raise her permanent earnings by even more than just the magnitude of the shock.

The transmission of shocks is further mitigated by $1/\phi_i - (1-s_{it})\psi_i$; the larger the absolute value of ϕ_i or ψ_i is, the smaller the effect of shocks on the revisions to the marginal utility of wealth. Ceteris paribus, a high absolute value for ϕ_i (high consumption substitution elasticity) implies the household can substitute future for current consumption less reluctantly and thus attenuate the impact of shocks on current consumption; a high ψ_i (high labor supply elasticity) implies the female can adjust her labour supply flexibly enough to smooth fluctuations in contemporaneous labor supply or consumption induced, otherwise, by shifts in the intertemporal budget constraint.⁷

3.2.2 Closed-Form Expressions for Inequality

Combining the spousal income processes in (3.2) with the approximated first-order conditions of the household problem in (3.3) and the approximate expression for the marginal utility of wealth in (3.4), one can write compact, closed-form, analytical expressions for the evolution of (growth in) earnings, labor supply, and consumption in the household given by⁸

$$\begin{pmatrix} \Delta y_{1it} \\ \Delta w_{2it} \\ \Delta h_{2it} \\ \Delta y_{2it} \\ \Delta y_{it} \\ \Delta c_{it} \end{pmatrix} \approx \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ k_{h_2,v_1} & k_{h_2,v_2} & 0 & k_{h_2,u_2} \\ k_{y_2,v_1} & k_{y_2,v_2} & 0 & k_{y_2,u_2} \\ k_{y,v_1} & k_{y,v_2} & k_{y,u_1} & k_{y,u_2} \\ k_{c,v_1} & k_{c,v_2} & 0 & 0 \end{pmatrix} \begin{pmatrix} v_{1it} \\ v_{2it} \\ \Delta u_{1it} \\ \Delta u_{2it} \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ k_{h_2,\omega} \\ k_{y_2,\omega} \\ k_{y,\omega} \\ k_{c,\omega} \end{pmatrix} \omega_{it}. \quad (3.5)$$

Here $\Delta y_{2it} = \Delta w_{2it} + \Delta h_{2it}$ is growth in female (log) earnings net of observables and $\Delta y_{it} \approx \rho_{it-1} \Delta y_{1it} + (1 - \rho_{it-1}) \Delta y_{2it}$ is growth in total (log) earnings of the family net of observables (with ρ_{it} the ratio of male over family earnings at t).

The k 's are transmission parameters of shocks (and of ω_{it}) into endogenous

⁷This is the *indirect* effect of wage shocks on labor supply; whereas the direct effect works towards increasing female hours of work after a *positive* shock, the indirect one works oppositely: it attenuates fluctuations in ε_{it} and, therefore, further fluctuations in labor supply to which ε_{it} feeds back.

⁸The first two equations in the system are exact, not approximate, equalities.

household outcomes. They all vary in the cross-section (with i) and over time (with t) but such subscripts are removed to ease the notation. Defining $\zeta_{it} = 1/\phi_i - (1-s_{it})\psi_i$, the transmission parameters are

$$\begin{aligned}
 k_{h_2,v_1} &= \psi_i \zeta_{it} s_{it} \\
 k_{h_2,v_2} &= \psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i \\
 k_{h_2,u_2} &= \psi_i \\
 k_{h_2,\omega} &= \psi_i \\
 \\
 k_{y_2,v_1} &= \psi_i \zeta_{it} s_{it} \\
 k_{y_2,v_2} &= \psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i + 1 \\
 k_{y_2,u_2} &= \psi_i + 1 \\
 k_{y_2,\omega} &= \psi_i \\
 \\
 k_{y,v_1} &= \rho_{it-1} + (1 - \rho_{it-1})\psi_i \zeta_{it} s_{it} \\
 k_{y,v_2} &= (1 - \rho_{it-1}) (\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i + 1) \\
 k_{y,u_1} &= \rho_{it-1} \\
 k_{y,u_2} &= (1 - \rho_{it-1})(\psi_i + 1) \\
 k_{y,\omega} &= (1 - \rho_{it-1})\psi_i \\
 \\
 k_{c,v_1} &= \phi_i \zeta_{it} s_{it} \\
 k_{c,v_2} &= \phi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) \\
 k_{c,\omega} &= \phi_i
 \end{aligned}$$

Note that the expressions for Δy_{2it} and Δy_{it} in (3.5) are, in practice, transformations of Δy_{1it} , Δw_{2it} and Δh_{2it} .

System (3.5) is appealing because it is empirically tractable and provides a neat picture for how shocks, preferences, and other factors contribute to earnings and consumption growth. It also offers a theoretical interpretation to the dynamics of earnings and consumption inequality. Blundell and Preston (1998) use the covariance matrix of (a simplified) system (3.5) to study the

transmission of income uncertainty to consumption inequality. In a follow-up paper, Blundell et al. (2008) use the covariance matrix of (a simplified) system (3.5) to motivate and empirically estimate the degree of insurance against income shocks that is available to households. Finally, Blundell et al. (2016) use a more sophisticated system (3.5), one that stems from a household problem where they relax many of the structural simplifications imposed herein, in order to identify household preferences when there is no unobserved preference heterogeneity.

Inspecting (3.5), there are two distinct parts contributing to earnings or consumption growth. One is due to shocks to the primitive incomes, namely to male earnings and female wages. The other is due to household precautionary motives over the life, namely due to ω_{it} (for details see appendix C.1). Blundell et al. (2008, p. 1897) note that factors other than shocks also contribute to earnings and consumption growth and they refer to factors pertaining to “measurement error in consumption, preference shocks, innovations to higher moments of the income process, etc”. As ω_{it} in the present model is a function of the variance of (the growth in) the marginal utility of wealth, it comes closer to reflecting preference shocks and innovations to higher moments of the income process.

The transmission parameters k , which load the income shocks onto labor supply, earnings, and consumption, vary in the cross-section through their dependence on preferences ϕ_i and ψ_i , and the shares s_{it} and ρ_{it} . Thus when a given *same* shock hits two different households, the households will not necessarily respond similarly to it. Crucially, even if the two households share the same initial conditions (i.e. same s_{it} and ρ_{it}), their heterogeneous preferences will likely lead them to divergent labor supply, earnings, and consumption paths. As a consequence, the variability of (inequality in) earnings or consumption growth in the cross-section is partly due to the variance of shocks and partly due to preference heterogeneity.

Before I present closed-form expressions for cross-sectional inequality, it

is useful to model the second moments of ω_{it} . Moreover, the presence of cross-sectional preference heterogeneity and precautionary motives ω_{it} , necessitates a discussion of how income shocks relate to preferences and ω_{it} .

Assumption (Second moments of ω_{it}).

$$\mathbb{E}(\omega_{it}\omega_{it+s}) = \begin{cases} \sigma_{\omega,t}^2 & \text{if } s = 0 \\ 0 & \text{otherwise} \end{cases}$$

where $s = \{0, \dots, T\}$.

I allow the variance of ω_{it} to differ with time but I assume ω_{it} is not serially correlated. Note that if households have heterogenous discount factors β then ω_{it} is serially correlated. In that case, identification of the distribution of household preferences (section 3.3) would still be possible but more involved.

Assumption (Independence of income shocks and preferences). *Income shocks are independent of household preferences, namely*

$$\begin{aligned} v_{jit} &\perp \phi_i \text{ or } \psi_i \\ u_{jit} &\perp \phi_i \text{ or } \psi_i \end{aligned}$$

for all $j = \{1, 2\}, i, t$. *Income shocks are mean independent of ω_{it} , namely*

$$\begin{aligned} \mathbb{E}(v_{jit}\omega_{it}) &= \mathbb{E}(v_{jit})\mathbb{E}(\omega_{it}) = 0 \\ \mathbb{E}(u_{jit}\omega_{it}) &= \mathbb{E}(u_{jit})\mathbb{E}(\omega_{it}) = 0 \end{aligned}$$

for all $j = \{1, 2\}, i, t$ and, consequently, $Cov(v_{jit}, \omega_{it}) = Cov(u_{jit}, \omega_{it}) = 0$.

Preferences are mean independent of ω_{it} , namely

$$\begin{aligned} \mathbb{E}(\phi_i\omega_{it}) &= \mathbb{E}(\phi_i)\mathbb{E}(\omega_{it}) = 0 \\ \mathbb{E}(\psi_i\omega_{it}) &= \mathbb{E}(\psi_i)\mathbb{E}(\omega_{it}) = 0 \end{aligned}$$

for all i, t and, consequently, $Cov(\phi_i, \omega_{it}) = Cov(\psi_i, \omega_{it}) = 0$.

This assumption postulates the independence of income shocks and preferences; this is a standard assumption in the literature as, for example, in Cossaert and Demuynck (2014). It implies that, conditional on observables, unobserved preference heterogeneity is independent of primitive incomes. The assumption also postulates the mean independence between income shocks and precautionary motives ω_{it} as well as between preferences and ω_{it} . I cannot impose strong stochastic independence between shocks and ω_{it} as higher moments of income may affect the variance of the growth in the marginal utility of wealth (of which ω_{it} is a function).

I rely on (3.5) to obtain closed-form expressions for the variance of female earnings and consumption. These are:

$$\begin{aligned}\text{Var}_t(\Delta y_{2it}) &= \mathbb{E}((\psi_i \zeta_{it} s_{it})^2) \sigma_{v_1,t}^2 + \mathbb{E}((\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i + 1)^2) \sigma_{v_2,t}^2 \\ &\quad + 2\mathbb{E}((\psi_i + 1)^2) \sigma_{u_2,t}^2 + \mathbb{E}(\psi_i^2) \sigma_{\omega,t}^2 \\ \text{Var}_t(\Delta c_{it}) &= \mathbb{E}((\phi_i \zeta_{it} s_{it})^2) \sigma_{v_1,t}^2 + \mathbb{E}((\phi_i \zeta_{it} (1 - s_{it})(1 + \psi_i))^2) \sigma_{v_2,t}^2 \\ &\quad + \mathbb{E}(\phi_i^2) \sigma_{\omega,t}^2\end{aligned}$$

where $\text{Var}_t(\cdot)$ indicates the cross-sectional variance at time t . Goodman (1960) and Bohrnstedt and Goldberger (1969) provide tools for the second moments of *products* of random variables that are not necessarily mutually independent. It is straightforward to obtain similar closed-form expressions for the variance of labor supply, household earnings, as well as all possible covariances among them. Such expressions describe fully the dynamics of consumption, earnings, and hours inequality across households.

3.3 Identification

This section has three parts. First, I show which conditions in the data can identify the parameters of the primitive income processes, namely the second or higher moments of income shocks. Second, I present the moments that identify the transmission parameters of income shocks (the k 's) into labor

supply, earnings, and consumption. Third, I discuss what can be learned about the distribution of preferences in the cross section, namely the distribution of the consumption substitution elasticity ϕ_i and the labor supply elasticity ψ_i . I will show that, although the mean and the variance of the latter *are* identifiable, there is very little to be learned about the distribution of the former without variation in the real price of consumption. Throughout this section I assume that panel data for incomes, hours of work, and consumption are available (for example the Panel Study of Income Dynamics after 1999).

3.3.1 Income Shocks

Identification of the parameters of the income processes follows Meghir and Pistaferri (2004) and earlier studies. Specifically,

$$\text{Cov}(\Delta y_{1it}, \Delta y_{1it-1} + \Delta y_{1it} + \Delta y_{1it+1})$$

identifies $\sigma_{v_1,t}^2$ and

$$\text{Cov}(\Delta w_{2it}, \Delta w_{2it-1} + \Delta w_{2it} + \Delta w_{2it+1})$$

identifies $\sigma_{v_2,t}^2$ with $\text{Cov}(\cdot, \cdot)$ indicating the covariance. The intuition is as follows: $\sum_{s=-1}^{s=1} \Delta y_{1it+s}$ (or $\sum_{s=-1}^{s=1} \Delta w_{2it+s}$) strips Δy_{1it} (Δw_{2it}) of its mean-reverting transitory shock at t and, therefore, the covariance between the sum and Δy_{1it} (Δw_{2it}) identifies the variance of the permanent shock.

The covariance of immediately consecutive income / wage growths identifies (minus) the variance of the transitory shock because the transitory shock is the only mean-reverting component in (3.2). Specifically

$$-\text{Cov}(\Delta y_{1it}, \Delta y_{1it+1})$$

identifies $\sigma_{u_1,t}^2$ and

$$-\text{Cov}(\Delta w_{2it}, \Delta w_{2it+1})$$

identifies $\sigma_{u_2,t}^2$.

Higher moments of the distribution of shocks can be identified by similar arguments. As an illustration for the third moments,

$$\text{Cov}((\Delta y_{1it})^2, \Delta y_{1it-1} + \Delta y_{1it} + \Delta y_{1it+1})$$

identifies $\gamma_{v_1,t}$ and

$$\text{Cov}((\Delta w_{2it})^2, \Delta w_{2it-1} + \Delta w_{2it} + \Delta w_{2it+1})$$

identifies $\gamma_{v_2,t}$. Moreover

$$-\text{Cov}((\Delta y_{1it})^2, \Delta y_{1it+1})$$

identifies $\gamma_{u_1,t}$ and

$$-\text{Cov}((\Delta w_{2it})^2, \Delta w_{2it+1})$$

identifies $\gamma_{u_2,t}$.

3.3.2 Transmission Parameters of Income Shocks

Identification of the transmission parameters k relies on the logic of Blundell et al. (2016). For $x = \{y_1, w_2, h_2, y_2, y, c\}$ the following conditions identify the *average* transmission parameter k of shocks into x

$$\begin{aligned} \mathbb{E}(k_{x,v_1}) &= \frac{\text{Cov}(\Delta x_{it}, \sum_{s=-1}^{s=1} \Delta y_{1it+s})}{\text{Cov}(\Delta y_{1it}, \sum_{s=-1}^{s=1} \Delta y_{1it+s})} \\ \mathbb{E}(k_{x,v_2}) &= \frac{\text{Cov}(\Delta x_{it}, \sum_{s=-1}^{s=1} \Delta w_{2it+s})}{\text{Cov}(\Delta w_{2it}, \sum_{s=-1}^{s=1} \Delta w_{2it+s})} \\ \mathbb{E}(k_{x,u_1}) &= \frac{\text{Cov}(\Delta x_{it}, \Delta y_{1it+1})}{\text{Cov}(\Delta y_{1it}, \Delta y_{1it+1})} \\ \mathbb{E}(k_{x,u_2}) &= \frac{\text{Cov}(\Delta x_{it}, \Delta w_{2it+1})}{\text{Cov}(\Delta w_{2it}, \Delta w_{2it+1})}. \end{aligned} \tag{3.6}$$

As there are 6 elements in x 's set and 4 identifying conditions above, in total 24 transmission parameters are identified. Note that these may vary with time

through their dependence on s_{it} or ρ_{it-1} and (3.6) permits a period-by-period identification.

The intuition behind the identification result is the following. The covariance between Δx_{it} and $\sum_{s=-1}^{s=1} \Delta y_{1it+s}$ in the numerator of the first condition in (3.6) identifies the variance of men's permanent income shock weighed by its average loading factor (transmission parameter) onto Δx_{it} . The denominator adjusts for the variance of the permanent shock, thus the ratio identifies the average loading factor. Identification of the remaining transmission parameters follows a similar logic.

3.3.3 Distribution of Preferences

Now I turn to the first two moments of the cross-sectional distribution of preferences, namely of ψ_i and ϕ_i . Household-specific parameter ψ_i is the female labor supply elasticity with respect to wage w_{2it} . Note that ψ_i is 1.) an *intensive* margin elasticity, as the decision variable h_{2it} is hours of work conditional on working; 2.) a *Frisch* elasticity, as the marginal utility of wealth λ is held constant in expected terms. Household-specific parameter ϕ_i is the elasticity of consumption with respect to its own price p . Again, this is a *Frisch* elasticity, as the marginal utility of wealth λ is held constant in expected terms (Browning et al., 1999, section 3.2).

3.3.3.1 Labor Supply Elasticity

The average ψ_i in the cross-section is identified through the transmission parameter of transitory wage shocks to female labor supply, namely

$$\mathbb{E}(\psi_i) = \mathbb{E}(k_{h_2, u_2}).$$

Alternative conditions also identify $\mathbb{E}(\psi_i)$ or serve as overidentifying restrictions. For example, $\mathbb{E}(k_{y_2, u_2})$ identifies $\mathbb{E}(\psi_i) + 1$ whereas $\mathbb{E}(k_{y, u_2})$ identifies $(\mathbb{E}(\psi_i) + 1)(1 - \mathbb{E}(\rho_{it-1}))$ while $\mathbb{E}(\rho_{it-1})$ is obtainable directly from the data.⁹

⁹I assume preferences are independent of incomes, therefore $\mathbb{E}(\psi_i \rho_{it-1}) = \mathbb{E}(\psi_i) \mathbb{E}(\rho_{it-1})$.

Identification of the cross-sectional variance of ψ_i hinges on the idea that the covariance between consecutive growths in labor supply must be due to the variance of the mean-reverting transitory shock as well as the variability in the labor supply response, i.e. the variance of the labor supply elasticity. Specifically

$$\text{Cov}(\Delta h_{2it}, \Delta h_{2it+1}) = \text{Cov}(\psi_i u_{2it}, -\psi_i u_{2it}) = -\mathbb{E}(\psi_i^2) \sigma_{u_{2,t}}^2$$

(see appendix C.3 for details) and, consequently, $\mathbb{E}(\psi_i^2)$ is identified as

$$\mathbb{E}(\psi_i^2) = \frac{\text{Cov}(\Delta h_{2it}, \Delta h_{2it+1})}{\text{Cov}(\Delta w_{2it}, \Delta w_{2it+1})}.$$

Then it is straightforward to identify the cross-sectional variance of the labor supply elasticity as

$$\text{Var}(\psi_i) = \mathbb{E}(\psi_i^2) - (\mathbb{E}(\psi_i))^2.$$

Again, a number of alternative moments can be implicated in the identification of $\mathbb{E}(\psi_i^2)$ such as the covariance between consecutive growths in female earnings.

Finally, identification of higher moments of the cross-sectional distribution of ψ_i is possible following the same logic illustrated above. In all cases, a starting point for the identification is the covariance between $(\Delta h_{2it})^m$ and Δh_{2it} (or between $(\Delta y_{2it})^m$ and Δy_{2it}) where m is the order of the moment of ψ_i minus 1.

3.3.3.2 Consumption Substitution Elasticity

The mean and the variance of the consumption substitution elasticity are not separately identified when there is unobserved preference heterogeneity. The fundamental underlying reason is the lack of cross-sectional variation in the real price of consumption.

To see this point take the transmission parameter of men's permanent

shocks k_{c,v_1} into consumption. This identifies

$$\mathbb{E}(k_{c,v_1}) = \mathbb{E}(\phi_i \zeta_{it} s_{it}) = \mathbb{E}(\phi_i) \mathbb{E}(\zeta_{it} s_{it}) + \text{Cov}(\phi_i, \zeta_{it} s_{it}) \quad (3.7)$$

with the covariance being due to the variance of ϕ_i or the likely covariance between ϕ_i and ψ_i ($\zeta_{it} = 1/\phi_i - (1-s_{it})\psi_i$ is a function of both ϕ_i and ψ_i). I cannot separate the moments of ζ_{it} and s_{it} in (3.7) as the former is a function of the latter.

Expression (3.7) illustrates that the transmission of shocks into consumption involves both the first and second moments of ϕ_i . Unlike ψ_i , this is true across all transmission parameters of *permanent* shocks in (3.5). ϕ_i always appears multiplying ζ_{it} ; in turn, ζ_{it} , which is a transformation of ϕ_i and ψ_i , always appears multiplying either ϕ_i or ψ_i . It is this multiplicity that prohibits recovering distinct moments of ϕ_i .

Finally, even without unobserved heterogeneity in ϕ_i or ψ_i , the transmission parameter k_{c,v_1} (or any transmission parameter of a *permanent* shock) fails to identify ϕ . Notice that the second term in (3.7), the covariance, drops out when ϕ does not vary, whereas the first term, the product of expectations, becomes $\phi \mathbb{E}\left(\frac{1}{\phi - (1-s_{it})\psi} s_{it}\right)$. As ϕ is weighed by the expectation of a nonlinear function of ϕ , ψ , and s_{it} , Jensen's inequality prohibits the identification of ϕ itself. This result applies more generally to all papers that rely on approximations to the intertemporal budget constraint and the first order conditions of a household model (for example Blundell et al., 2015, 2016). Without variation in the price of consumption these papers rely on the approximated budget constraint to identify the consumption substitution elasticity. In practice, however, they only identify the said elasticity *up to scale*, where the scale depends on the specificities of the household problem in each case, as is shown above.

3.4 Tests

In this section I present a number of restrictions that must be satisfied in the data for the model of section 3.2 to be consistent with observed household behavior (specification tests) and I discuss two particular specification features. I also present a test for unobserved preference heterogeneity across households.

3.4.1 Specification Test and Further Discussion

The model presented in section 3.2, with preferences additively separable between consumption and labor supply and with the assumption that transitory shocks do not induce wealth effects, implies that transitory income shocks do not transmit into household consumption. In addition, men's fixed labor supply implies that men's transitory shocks do not transmit into female labor supply or earnings. The corresponding theoretical transmission parameters k are 0 and, as a consequence, their empirical counterparts must be statistically insignificantly different from 0. This provides the basis for a specification test.

In addition, the model postulates *same* transmission parameters of men's permanent shocks into female labor supply and earnings, and that $k_{y_2, u_2} - k_{h_2, u_2} = 1$ and $k_{y, u_1} = s_{it}$, both directly testable in the data.

Bringing the empirical counterparts of these conditions together, I propose the following specification test.

Proposition (Model specification test). *If the following moment restrictions are (jointly or separately) rejected in the data*

$$(a) \mathbb{E}(k_{h_2, u_1}) = 0$$

$$(e) \mathbb{E}(k_{y_2, v_1}) - \mathbb{E}(k_{h_2, v_1}) = 0$$

$$(b) \mathbb{E}(k_{y_2, u_1}) = 0$$

$$(f) \mathbb{E}(k_{y_2, u_2}) - \mathbb{E}(k_{h_2, u_2}) = 1$$

$$(c) \mathbb{E}(k_{c, u_1}) = 0$$

$$(d) \mathbb{E}(k_{c, u_2}) = 0$$

$$(g) \mathbb{E}(k_{y, u_1}) - \mathbb{E}(s_{it}) = 0$$

then the model of section 3.2 is not consistent with observed household behaviour.

Obviously the model of section 3.2 is rather stylized and that makes its empirical rejection likely. Relaxing the additive separability between consumption and labor supply will likely improve the chances that the model passes a (modified) specification test. Indeed there is ample empirical evidence against the separability of consumption and leisure (eg. Browning and Meghir, 1991).

Relaxing such strong separability will introduce two additional preference parameters, the (Frisch) elasticity of consumption with respect to female wage and the (Frisch) elasticity of labor supply with respect to the price of consumption. Moments of the cross-sectional distribution of the former parameter can be identified by the covariation between consumption and wages like in section 3.3.3.1 for ψ_i . Lack of variation in the price of consumption impedes identification of the distribution of the latter, although symmetry of the Frisch matrix of substitution effects may provide additional identifying equations.

Finally, endowing men with an endogenous labor choice, and allowing that to be nonseparable from consumption and women's labor supply, will introduce a number of additional parameters, such as men's own-wage labor supply elasticity or their labor supply elasticity with respect to women's wages. The cross-sectional distribution of most such elasticities is identified through variation in men's and women's wages following the logic developed in section 3.3.3.1.

3.4.2 Unobserved Preference Heterogeneity

The model specification test above does not provide any indication as to whether the data provide evidence for unobserved preference heterogeneity or not. A simple intuitive way to test for such heterogeneity is based on the variance of the labor supply elasticity ψ_i whose identification was presented in section 3.3.3.1.

Proposition (Preference heterogeneity test). *If the cross-sectional variance of the own-wage labor supply elasticity is statistically insignificantly different*

from 0, namely if the data cannot reject

$$\text{Var}(\psi_i) = 0,$$

then there is evidence against cross-household heterogeneity in ψ_i .

The preference heterogeneity test only involves the labor supply elasticity ψ_i as the distribution of ϕ_i is not identified without variation in the price of consumption (section 3.3.3.2). This obviously renders the test partial. A possible failure to reject $\text{Var}(\psi_i) = 0$ is not informative about unobserved preference heterogeneity in ϕ_i and, similarly, a rejection of $\text{Var}(\psi_i) = 0$ does not necessarily also imply heterogeneity in ϕ_i .

3.5 Conclusions

This chapter investigates the transmission of income shocks into consumption across households that exhibit unobserved preference heterogeneity. The chapter develops a tractable life-cycle model for household consumption, savings, and labor supply. I introduce a general form of unobserved heterogeneity through allowing preferences to be household-specific (and nonseparable from heterogeneity) even conditional on observable characteristics. I obtain analytical expressions for consumption and earnings as functions of, among other things, income shocks and household-specific consumption and labor supply elasticities. To do so, I rely on Taylor approximations to the household lifetime budget constraint and first-order conditions as in Blundell and Preston (1998).

I present theoretical conditions with straightforward empirical counterparts that deliver identification of the cross-sectional distribution of (consumption, labor supply) elasticities with respect to wages if panel data on consumption, earnings, and hours are available. Identification does not depend on any specific parametrization of household preferences or their distribution. I show that identification of (consumption, labor supply) elasticities with respect to

the price of consumption, including the consumption substitution elasticity, is not feasible from the afoementioned analytical expressions without variation in the real price of consumption.

Appendix A

Appendix to Wages and Family Time Allocation

A.1 Data: Sample Selection and Variables

Chapter 1 uses information from the Panel Study of Income Dynamics (PSID). I select men and women aged 25 to 65 from the core sample ('Survey Research Center') between years 1980 and 2009. I split this into two distinct and non-overlapping samples: (i) a major sample of households of continuously married men and women throughout the years they're observed, and (ii) a minor sample of singles of both genders. Below I describe the two samples in detail.

Major PSID sample I follow households headed by a married or permanently cohabiting opposite-sex couple. I require that these households are stable in that they do not experience any compositional changes in the head couple such as divorce or remarriage. Compositional changes regarding children are permitted. Currently I follow one cohort of households only. I define this cohort as those households whose male spouse (male head of the household in the PSID) is born between years 1943 and 1955 (implying he is between 25 and 37 years old in 1980). Given that the age difference between him and his spouse in approximately two thirds of households in this cohort

does not exceed ± 3 years, I do not explicitly condition on similar years of birth for the female spouse. I drop a few households for which information on their state of residence is ambiguous (these may be households that reside outside the US for part of the survey year). I also drop households with one or more spouses being farmers (hard to trust their earnings), disabled or students (because their time allocations may be constrained by their circumstances), or households for which labor earnings of a *working* spouse fall below 1% or above 99% of the (gender- and time-specific) distribution. The resulting dataset is an unbalanced panel of 1279 households observed over at least two consecutive years. More than 55% of households are observed for at least 10 years and more than 30% for at least 20.¹

Hourly wages are calculated as annual labor earnings over annual hours of work for those working. To account partly for measurement error in wages I only use the central 96% of the wage distribution for each gender after excluding those who work less than 10 hours per year. Figure A.1 in this appendix plots median and mean wages by gender. *Annual labor earnings* are self-reported gross earnings from all jobs and include salaries, bonuses, overtime, tips etc. Around 1993 the definition of earnings changes slightly and the available measure excludes some previously included minor components of earnings (such as the labor part of business income).² I remove inflation from all monetary values using the All-Urban-Consumers CPI.

Annual hours of work are defined as total work hours across all jobs in a given year including overtime. I assume that hours reported at one point in the year are evenly allocated over the year. I discretise the amount of time women put into market work (see table 1.2) using a 3-point distribution: not working (0-10 annual hours modeled as 0 hours), working part-time (10-1000 annual hours modeled as 4 daily hours in a 5-day 50-week annual schedule),

¹The proportion of households that are ‘stable’, among all households satisfying the other selection criteria laid out in this paragraph, is approximately 81%.


²Despite this, the PSID officially treats men’s earnings series as consistent over time. For female earnings two different series are provided (one prior to 1993 and one after). I combine the two into a single female series.

and working full time (more than 1000 annual hours modeled as 8 daily hours). There is sufficient bunching of hours in the data to justify the above discrete approximation.

Household work is defined on a weekly basis as time spent on cooking, cleaning, and “doing other work around the household”. I discretise the amount of time put into household work using a separate 2-point distribution for each gender: for men, ‘low’ hours (up to 2 hours weekly modeled as 0.4 hours/day in a 5-day week) or ‘low middle’ hours (more than 2 hours weekly modeled as 1.6 daily hours); for women ‘high middle’ hours (up to 15 hours weekly modeled as 3 daily hours in a 5-day week) or ‘maximum’ hours (more than 15 modeled as 6 daily hours). Again, there is sufficient bunching of household hours in the data to justify these discrete approximations and the precise choice of modeling hours.

Age of the youngest child is classified in four groups to reflect the way stochastic fertility is modeled in section 1.3.3: an age 0 in the data indicates the absence of a child younger than 18 years old (modeled as $N_t = 1$), ages 1 – 4 indicate a child less than 5 years ($N_t = 2$), ages 5 – 9 indicate a child at least 5 but less than 10 years old ($N_t = 3$), and ages 10 – 17 indicate older children up to 18 years old ($N_t = 4$).

Minor PSID sample This sample consists of single men and women and mimics many of the selection criteria applied to the major PSID sample above. I select individuals who report having been divorced or separated, work in the labor market (as I require information on their earnings), and whose earnings do not fall below 1% or above 99% of the (gender- and time-specific) distribution. I drop a few individuals for which information on their state of residence is ambiguous (these may reside outside the US for part of the survey year), farmers (hard to trust their earnings), or those with missing information on their education (required for the projections of earnings). The resulting dataset consists of 4561 divorced male-year and 7614 divorced female-year observations. I define and deflate *annual labor earnings* like above.

Figure A.1: Evolution of wages with mean age of household head

Notes: This figure plots median and mean hourly wages by gender for married people over their life-cycle. One cohort only is depicted; mean age on the horizontal axis coincides with calendar time (1980-2009) and this is partly responsible for the continuous growth of wages in the graph. Only the central 96% of the wage distribution by gender and mean age (year) is used. A 95% confidence interval appears in gray shade.

Source: Major PSID sample.

A.2 Model: Public and Private Consumption

This appendix extends the model in section 2.2 to allow spouses to consume public *as well as* private consumption goods. In this case, individual j has preferences \tilde{U}_j given by

$$\tilde{U}_j(Q, q_j, l_j; \mathbf{z}_j)$$

where q_j is the private (rival) consumption good. The rest of the notation remains like in the main text. One can think of the private good as, for example, own clothing and the public good as food at home or children's expenditure.

The household problem during the working period of life is given by (1.1)-(1.7) after replacing individual preferences with \tilde{U}_j and the budget constraint (1.3) with

$$A_t + \sum_{j=1}^2 w_{jt} h_{jt} = K_t + p_t \sum_{j=1}^2 q_{jt} + CC_t(h_{2t}, N_t) + \frac{A_{t+1}}{1+r}.$$

Here p_t is the relative price of the private good at t after normalizing the price of the public good to 1 in every period. The set of state variables is unaffected but the set of choice variables \mathcal{C}_t is augmented to include q_{1t} and q_{2t} .

Preferences can be parameterized by

$$\tilde{U}_j(Q_t, q_{jt}, l_{jt}; \mathbf{z}_{jt}) = \frac{1}{1-\gamma} \left(\alpha_j (Q_t/s(N_t))^{1-\gamma} + (1-\alpha_j) q_{jt}^{1-\gamma} \right) \times \exp(g_j(l_{jt}; \mathbf{z}_{jt}))$$

which is a straightforward extension of (1.8). The leisure sub-utility $g_j(\cdot)$ remains unchanged. Here α_j serves as the utility share of public consumption which may depend on observables such as the presence or age of the youngest child in the family.

The extension to private consumption does not alter the fundamentals of the problem: the problem still is a typical mixture of discrete (time-use) and continuous choices (public and private consumption, savings). The solution algorithm is not complicated significantly: for each optimal public consumption-savings bundle, and conditional on a time-use choice, the marginal rates of substitution between the private consumption goods *and* between public-private consumption deliver the optimal quantities for q_1 and q_2 . The separability between public and private consumption facilitates the solution. However, the algorithm is more time-consuming as one now has to search for the best Q and (with the use of the marginal rates of substitution) for the optimal q_1 and q_2 given some future assets and then repeat this along a grid of future assets (*i.e.* two-dimensional instead of one-dimensional ‘table look-up’).

For identification of α_j one needs information on private goods *for each spouse* as well as public consumption goods. The Consumer Expenditure Survey in the US provides information on clothing expenditure by gender. However this tends to be a tiny proportion of total household expenditure and it is unclear which other goods reported therein could serve as private.

A.3 Estimation: Exogenous Blocks

Wages To obtain estimates of the second moments of shocks I run a Minimum Distance estimation matching the empirical covariance matrix of log wages to its theoretical counterpart. I illustrate the main points of this estimation referring to time t as calendar time but recall that calendar time coincides with mean age of the household head given the focus on one cohort only. From the major PSID sample I collate the vector $\widetilde{\mathbf{W}} = (\widetilde{\mathbf{W}}_{1980}, \widetilde{\mathbf{W}}_{1981}, \dots, \widetilde{\mathbf{W}}_{2009})'$ where

$$\widetilde{\mathbf{W}}_t = \begin{pmatrix} \mathbb{E}[(\Delta \ln w_{1t})^2] \\ \mathbb{E}[\Delta \ln w_{1t} \Delta \ln w_{1t+1}] \\ \mathbb{E}[(\Delta \ln w_{2t})^2] \\ \mathbb{E}[\Delta \ln w_{2t} \Delta \ln w_{2t+1}] \\ \mathbb{E}[\Delta \ln w_{1t} \Delta \ln w_{2t}] \\ \mathbb{E}[\Delta \ln w_{1t} \Delta \ln w_{2t+1}] \\ \mathbb{E}[\Delta \ln w_{2t} \Delta \ln w_{1t+1}] \end{pmatrix}, \quad t \in [1980, 2009]$$

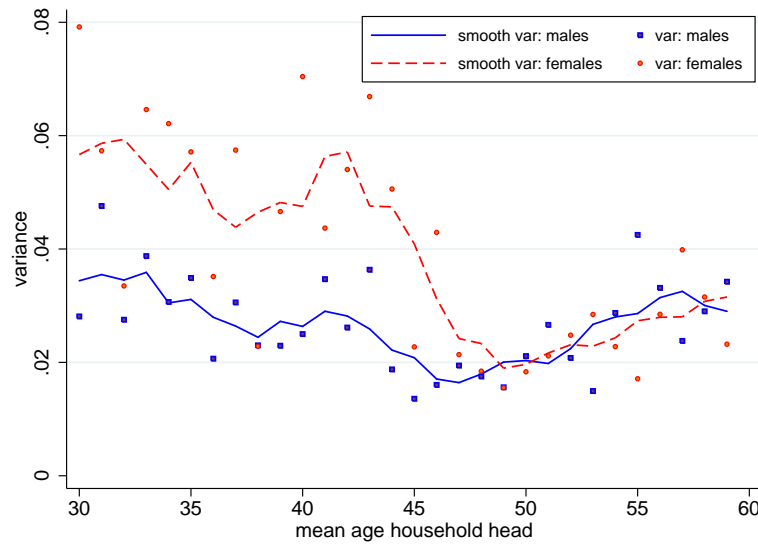
Due to the parametrization in (1.10), I ignore any auto-covariances of order higher than 1. In the PSID these are insignificantly different from 0 anyway.

The theoretical counterpart of $\widetilde{\mathbf{W}}_t$ is $\mathbf{W}(\boldsymbol{\vartheta})$ and is a function of the second moments of shocks over the life-cycle (parameter $\boldsymbol{\vartheta}$). An estimate of $\boldsymbol{\vartheta}$ is

$$\hat{\boldsymbol{\vartheta}} = \arg \min_{\boldsymbol{\vartheta}} (\widetilde{\mathbf{W}} - \mathbf{W}(\boldsymbol{\vartheta}))' I (\widetilde{\mathbf{W}} - \mathbf{W}(\boldsymbol{\vartheta}))$$

where I is the identity matrix. The estimates of $\boldsymbol{\vartheta}$ appear in table A.1 alongside their standard errors; to calculate those I adopt the block bootstrap with 500 replications.

The point-estimates in table A.1 are not used directly as inputs to the structural model. To reduce the effect of measurement error, I replace the point-estimates with 5-point two-sided moving averages (suitably adapted to deal with corners); a similar approach is taken by French (2005). The original point-estimates of the variances of men's and women's permanent shocks, as

Figure A.2: Actual and smoothed variances of permanent shocks

Notes: This figure plots the estimates of the variance of permanent shocks for each spouse (scatter points) as well as 5-point two-sided moving averages that pass through the scatters. The central 96% only of the wage distribution by gender and year is used for estimation.

Source: Major PSID sample and own calculations.

well as the refined ones, appear graphically in figure A.2 in this appendix.

Child care costs I calibrate $cchr$ at a constant \$6.59 (expressed in 2010 dollars) throughout the 1980-2009 period; see section 1.5.1 in the main text for details. Whenever this rate is below the real federal minimum wage, I update $cchr$ to reflect this. Essentially the hourly wage of child care workers in the model decreases relative to that of the general population (of both men and women) reflecting -what seems to be- a consensus that child care has steadily become less expensive in the last 3 decades. Finally, I calculate the probability of a family facing positive child care costs by counting the number of families in a given fertility state who report non-zero such costs (the PSID collects information on child care expenditure after 1988). This is done separately by calendar year. In years when child care expenditure is missing from the PSID I use the probabilities estimated in the closest period when data are available.

Table A.2 in this appendix reports the hourly rate of child care over time as well as the estimated probabilities of positive child care costs for the relevant

fertility states $N_t = 2$ (youngest child younger than 5 years) and $N_t = 3$ (youngest child between 5 and 10 years). Households in fertility states $N_t = 1$ and $N_t = 4$ are modeled to not require formal child care. This is confirmed by the data (but not reported in table A.2).

Normalization of the initial Pareto weight To project lifetime earnings if spouses get divorce, I first pool earnings of divorcees for all years and ages; I do so separately by gender. The data come from the minor PSID sample described above. I regress earnings on race, education, a quadratic polynomial in age and their interactions. This is regression (1.15) in the main text and the results appear in table A.3 in this appendix. These results use information on divorcees between years 1980-1989 because I normalize intra-family bargaining power in the first 10 years of the family life-cycle only. Using the estimates from (1.15) I project lifetime earnings for each married spouse in the event of divorce and I use these projections to form a proxy for intra-family bargaining power (see section 1.5.3 in the main text for details). Table A.4 below reports how the derived intra-family bargaining power of the male spouse correlates with a number of characteristics of each partner.

A.4 Estimation: Model Fit and Overidentification

The following two tables present the full set of targeted moments used in the structural estimation (see section 1.5.2 in main text). The tables report the values of the moments in the data as well as their counterparts from the model.

In addition, figure A.3 reports the values of 64 non-targeted joint dynamic moments of time allocation. These are transition probabilities between time allocations and periods of time, namely probabilities that an individual engages in a given time allocation conditional on what they or their partner did one or two periods in the past. These moments, even though not targeted in the simulated method of moments estimation, are well aligned along or around the 45-degree line (this line indicates a perfect match).

Table A.1: Second moments of wage shocks

Year	Mean age head	I. Permanent shocks			II. Transitory shocks		
		Men	Women	Covariance	Men	Women	Covariance
1980	30	0.0281 (0.0080)	0.0792 (0.0226)	0.0081 (0.0080)	0.0214 (0.0058)	0.0546 (0.0164)	0.0039 (0.0052)
1981	31	0.0476 (0.0107)	0.0573 (0.0200)	-0.0009 (0.0079)	0.0253 (0.0055)	0.0710 (0.0211)	0.0060 (0.0051)
1982	32	0.0275 (0.0079)	0.0335 (0.0211)	0.0097 (0.0073)	0.0226 (0.0046)	0.0697 (0.0228)	0.0037 (0.0063)
1983	33	0.0388 (0.0071)	0.0646 (0.0194)	0.0093 (0.0075)	0.0176 (0.0045)	0.0602 (0.0155)	-0.0019 (0.0047)
1984	34	0.0307 (0.0070)	0.0621 (0.0146)	0.0062 (0.0064)	0.0254 (0.0052)	0.0380 (0.0127)	-0.0010 (0.0049)
1985	35	0.0349 (0.0062)	0.0571 (0.0169)	0.0121 (0.0081)	0.0239 (0.0051)	0.0725 (0.0161)	-0.0045 (0.0049)
1986	36	0.0207 (0.0061)	0.0351 (0.0146)	0.0102 (0.0057)	0.0305 (0.0059)	0.0737 (0.0155)	-0.0032 (0.0051)
1987	37	0.0306 (0.0065)	0.0575 (0.0249)	0.0172 (0.0059)	0.0253 (0.0048)	0.0789 (0.0251)	-0.0017 (0.0038)
1988	38	0.0230 (0.0062)	0.0229 (0.0153)	0.0108 (0.0060)	0.0290 (0.0057)	0.0638 (0.0123)	-0.0053 (0.0045)
1989	39	0.0229 (0.0060)	0.0466 (0.0122)	0.0099 (0.0050)	0.0337 (0.0073)	0.0199 (0.0071)	0.0027 (0.0044)
1990	40	0.0250 (0.0064)	0.0704 (0.0177)	0.0043 (0.0058)	0.0178 (0.0049)	0.0408 (0.0089)	0.0055 (0.0041)
1991	41	0.0347 (0.0059)	0.0437 (0.0116)	0.0084 (0.0064)	0.0259 (0.0060)	0.0466 (0.0092)	-0.0010 (0.0039)
1992	42	0.0261 (0.0097)	0.0540 (0.0147)	0.0165 (0.0070)	0.0469 (0.0101)	0.0343 (0.0097)	-0.0118 (0.0053)
1993	43	0.0363 (0.0123)	0.0669 (0.0143)	0.0243 (0.0072)	0.0673 (0.0116)	0.0566 (0.0123)	-0.0032 (0.0059)
1994	44	0.0188 (0.0085)	0.0506 (0.0175)	0.0173 (0.0069)	0.0595 (0.0125)	0.0885 (0.0165)	0.0035 (0.0055)
1995	45	0.0136 (0.0077)	0.0227 (0.0135)	0.0071 (0.0063)	0.0337 (0.0057)	0.0514 (0.0099)	0.0011 (0.0044)
1996	46	0.0160 (0.0055)	0.0429 (0.0128)	0.0128 (0.0071)	0.0106 (0.0032)	0.0589 (0.0139)	-0.0017 (0.0058)
1997	47	0.0194 (0.0059)	0.0214 (0.0069)	-0.0004 (0.0040)	0.0413 (0.0101)	0.0628 (0.0118)	0.0083 (0.0051)
1999	49	0.0156 (0.0046)	0.0155 (0.0061)	0.0024 (0.0028)	0.0436 (0.0090)	0.0494 (0.0092)	-0.0048 (0.0042)
2001	51	0.0266 (0.0083)	0.0212 (0.0062)	0.0057 (0.0033)	0.0522 (0.0102)	0.0359 (0.0065)	0.0112 (0.0063)
2003	53	0.0149 (0.0052)	0.0285 (0.0087)	0.0034 (0.0036)	0.0571 (0.0126)	0.0415 (0.0091)	-0.0049 (0.0059)
2005	55	0.0425 (0.0103)	0.0171 (0.0045)	0.0019 (0.0042)	0.0621 (0.0155)	0.0411 (0.0067)	-0.0003 (0.0051)
2007	57	0.0238 (0.0075)	0.0399 (0.0078)	0.0032 (0.0038)	0.0430 (0.0129)	0.0390 (0.0105)	0.0054 (0.0059)
2009	59	0.0342 (0.0091)	0.0232 (0.0069)	0.0008 (0.0053)	0.0431 (0.0098)	0.0304 (0.0123)	-0.0004 (0.0102)

Notes: The table presents minimum distance estimates of the variances of permanent and transitory shocks over time, as well as their covariances between spouses. Block I refers to permanent shocks; block II refers to transitory shocks. Within each block the first column is men's variance of the shock, the second column is women's variance, and the third column is the covariance between the two. Block bootstrap standard errors from 500 replications are reported in parentheses. *Source:* Major PSID sample and own calculations.

Table A.2: Child care costs: price and probabilities

Year	Hourly rate (in \$2010)	Probability child care expenditure > 0	
		Fertility state $N_t = 2$	Fertility state $N_t = 3$
1980	8.71	59.07%	40.32%
1981	8.20	59.07%	40.32%
1982	8.04	59.07%	40.32%
1983	7.57	59.07%	40.32%
1984	7.33	59.07%	40.32%
1985	7.03	59.07%	40.32%
1986	6.79	59.07%	40.32%
1987	6.67	59.07%	40.32%
1988	6.59	59.07%	40.32%
1989	6.59	56.47%	45.56%
1990	6.59	60.34%	43.17%
1991	6.59	57.66%	42.11%
1992	6.80	60.93%	42.97%
1993	6.61	51.30%	41.74%
1994	6.59	55.85%	48.70%
1995	6.59	59.69%	47.13%
1996	6.59	59.13%	47.89%
1997	6.60	58.58%	45.39%
1998*	7.00	56.05%	45.25%
1999	6.89	53.52%	45.10%
2000*	6.74	52.28%	43.34%
2001	6.59	51.04%	41.58%
2002*	6.59	49.99%	43.01%
2003	6.59	48.95%	44.44%
2004*	6.59	49.25%	41.44%
2005	6.59	49.56%	38.43%
2006*	6.59	50.19%	40.07%
2007	6.59	50.82%	41.70%
2008*	6.59	50.64%	40.14%
2009	6.63	50.45%	38.58%

Notes: This table presents the hourly rate of child care in 2010 dollars (column 2) alongside the probability of a family reporting positive child care expenditure by fertility state (columns 3 and 4). Only the relevant fertility states are reported. In years when child care expenditure is missing from the PSID (prior to 1988) I use the probabilities estimated in the closest period when data are available. *In even years after 1997, when the PSID did not collect data, I use the mid-point of probabilities in adjacent years.

Table A.3: Earnings regressions: male and female divorcees

<i>Regressors:</i>	Dependent variable: Annual Earnings			
	I. Male divorcees		II. Female divorcees	
	Coeff.	p-value	Coeff.	p-value
Age	4269.56	[0.005]	807.36	[0.292]
Age ²	-50.71	[0.006]	-8.44	[0.330]
Race (black)	45811.67	[0.717]	-25464.03	[0.123]
Race (other)	287579.56	[0.018]	19119.09	[0.667]
Educ. (high school)	24278.46	[0.557]	-30085.92	[0.103]
Educ. (some college)	54787.30	[0.240]	-43205.38	[0.029]
Educ. (college)	80539.48	[0.124]	42461.88	[0.130]
Educ. (post-college)	-13021.28	[0.826]	-88158.37	[0.005]
Race (black) × Age	-3309.25	[0.657]	874.94	[0.282]
Race (other) × Age	-17150.95	[0.006]	-1014.91	[0.658]
Race (black) × Age ²	41.97	[0.698]	-5.98	[0.533]
Race (other) × Age ²	245.77	[0.001]	9.80	[0.729]
Educ. (high school) × Age	-718.35	[0.727]	1939.59	[0.028]
Educ. (some college) × Age	-2747.19	[0.254]	2654.67	[0.006]
Educ. (college) × Age	-3009.72	[0.235]	-1759.51	[0.193]
Educ. (post-college) × Age	1491.24	[0.622]	5652.83	[0.000]
Educ. (high school) × Age ²	9.57	[0.695]	-22.70	[0.023]
Educ. (some college) × Age ²	47.91	[0.110]	-29.94	[0.007]
Educ. (college) × Age ²	40.81	[0.164]	27.26	[0.082]
Educ. (post-college) × Age ²	-13.05	[0.727]	-66.13	[0.000]
Cons.	-50450.89	[0.096]	1033.31	[0.949]
R-Square	0.201		0.204	
Regression p value	0.000		0.000	
Obs. (singles in 1980-1989)	1201		1915	

Notes: This table presents OLS estimates and p-values from linear regressions of divorcees' earnings on a set of individual characteristics. These include: a quadratic polynomial in age, race and education dummies, and their interactions with the age polynomial. Race takes on three values for: 'white' (omitted), 'black', and 'other'. Education takes on five values for 'less than high school' (omitted), 'high school', 'some (less than) college', 'college', and 'post college'. The regressions are carried out separately by gender using years 1980-1989 of the minor PSID sample described in Appendix A.1. The number of observations reflects the number of male/female divorcees-year observations in 1980-1989.

Table A.4: Married men's initialized bargaining power: correlation with spousal attributes

<i>Regressors:</i>	Dependent variable: men's bargaining power	
	Coeff.	p-value
Educ. Male (high school)	0.080	[0.000]
Educ. Male (some college)	0.179	[0.000]
Educ. Male (college)	0.168	[0.000]
Educ. Male (post-college)	0.144	[0.000]
Educ. Female (high school)	-0.082	[0.000]
Educ. Female (some college)	-0.110	[0.000]
Educ. Female (college)	-0.176	[0.000]
Educ. Female (post-college)	-0.187	[0.000]
Race Male (black)	-0.031	[0.000]
Race Male (other)	0.143	[0.000]
Race Female (black)	-0.053	[0.000]
Race Female (other)	0.055	[0.000]
Age Male	0.006	[0.000]
Age Male ²	-0.000	[0.000]
Age Female	-0.007	[0.000]
Age Female ²	0.000	[0.000]
Cons.	0.622	[0.000]
Obs. (couples in 1980-1989)	5970	

Notes: This table presents estimates and p-values of the correlations (linear regressions) between the derived intra-family bargaining power of *married men* and a number of characteristics of each spouse. These include: education dummies, race dummies, and a quadratic polynomial in age. Education of either spouse takes on five values for 'less than high school' (omitted), 'high school', 'some (less than) college', 'college', and 'post college'. Race of either spouse takes on three values for: 'white' (omitted), 'black', and 'other'. The number of observations reflects the number of married household-year observations in 1980-1989.

Table A.5: Targeted moments: proportions men


	fertility state 1		fertility state 2	
	Data	Model	Data	Model
<i>Mean age head: 30-39</i>				
‘Low middle’ household work	0.732	0.704	0.734	0.713
<i>Mean age head: 40-49</i>				
‘Low middle’ household work	0.680	0.681	0.784	0.749
<i>Mean age head: 50-59</i>				
‘Low middle’ household work	0.705	0.705	0.679	0.737
	fertility state 3		fertility state 4	
	Data	Model	Data	Model
<i>Mean age head: 30-39</i>				
‘Low middle’ household work	0.686	0.716	0.620	0.585
<i>Mean age head: 40-49</i>				
‘Low middle’ household work	0.742	0.694	0.693	0.752
<i>Mean age head: 50-59</i>				
‘Low middle’ household work	0.673	0.728	0.786	0.738

Notes: This table reports the values of men’s targeted moments in the data and the model simulations. These moments are proportions of married men engaging in ‘low-middle’ household work by men’s mean age and family composition. For the definition of ‘low middle’ household work refer to table 1.2 in the main text.

Table A.6: Targeted moments: proportions women

	fertility state 1		fertility state 2		fertility state 3		fertility state 4	
	Data	Model	Data	Model	Data	Model	Data	Model
<i>Mean age head: 30-39</i>								
FT market work & 'max' household work	0.282	0.333	0.236	0.302	0.341	0.309	0.400	0.339
FT market work & 'high middle' household work	0.537	0.434	0.154	0.164	0.209	0.258	0.283	0.311
PT market work & 'max' household work	0.063	0.079	0.217	0.090	0.184	0.180	0.124	0.138
PT market work & 'high middle' household work	0.035	0.029	0.037	0.034	0.034	0.033	0.036	0.047
No market work & 'max' household work	0.060	0.096	0.328	0.312	0.214	0.219	0.144	0.129
<i>Mean age head: 40-49</i>								
FT market work & 'max' household work	0.296	0.312	0.271	0.308	0.283	0.319	0.345	0.350
FT market work & 'high middle' household work	0.468	0.479	0.229	0.249	0.295	0.293	0.394	0.348
PT market work & 'max' household work	0.050	0.062	0.195	0.095	0.165	0.181	0.103	0.100
PT market work & 'high middle' household work	0.037	0.044	0.042	0.035	0.069	0.021	0.034	0.039
No market work & 'max' household work	0.120	0.097	0.225	0.239	0.156	0.186	0.106	0.110
<i>Mean age head: 50-59</i>								
FT market work & 'max' household work	0.213	0.303	0.256	0.320	0.218	0.348	0.280	0.376
FT market work & 'high middle' household work	0.460	0.505	0.385	0.259	0.300	0.299	0.448	0.354
PT market work & 'max' household work	0.064	0.063	0.038	0.115	0.109	0.159	0.085	0.092
PT market work & 'high middle' household work	0.047	0.037	0.077	0.037	0.009	0.017	0.037	0.034
No market work & 'max' household work	0.144	0.086	0.205	0.213	0.227	0.177	0.122	0.095

Notes: This table reports the values of women's targeted moments in the data and the model simulations. These moments are proportions of married women engaging in various time allocations by the mean age of their husband and their family composition. For the definitions of 'maximum' and 'high middle' household work refer to table 1.2 in the main text.

Figure A.3: Non-targeted moments: transition probabilities

Notes: This figure plots 64 non-targeted joint dynamic moments of time allocation in the data (horizontal axis) against their model counterparts (vertical axis). These moments are transition probabilities, namely probabilities $\text{Prob}[\text{spouse}_j \text{ time allocation}_t \mid \text{spouse}_k \text{ time allocation}_{t-s}]$ that an individual of gender j engages in a given time allocation conditional on what they ($j = \{1,2\}$) or their partner ($k = \{1,2\}$) did $s = 1$ or $s = 2$ periods in the past. For men: $\text{spouse}_1 \text{ time allocation}_t = \{\text{'low middle' household work}\}$; for women: $\text{spouse}_2 \text{ time allocation}_t = \{\text{FT market work, PT market work, 'maximum' household work}\}$.

Appendix B

Appendix to Consumption Dynamics and Allocation in the Family

B.1 Frisch Elasticities

There are 9 Frisch (λ -constant) elasticities for each spouse j in the family (see Table 2.1 in the main text for verbal descriptions). The analytical expressions for these elasticities are

$$\begin{aligned}\eta_{j,h,w} &\equiv \left. \frac{\partial H_j}{\partial W_j} \frac{W_j}{H_j} \right|_{\lambda\text{-const.}} = D_j^{-1} \frac{U_{j,H_j}}{H_j} (U_{j,KK} U_{j,C_j C_j} - U_{j,KC_j}^2) \\ \eta_{j,h,p^c} &\equiv \left. \frac{\partial H_j}{\partial P^c} \frac{P^c}{H_j} \right|_{\lambda\text{-const.}} = D_j^{-1} \frac{U_{j,C_j}}{H_j} (U_{j,KC_j} U_{j,KH_j} - U_{j,KK} U_{j,C_j H_j}) \\ \eta_{j,h,p^k} &\equiv \left. \frac{\partial H_j}{\partial P^k} \frac{P^k}{H_j} \right|_{\lambda\text{-const.}} = D_j^{-1} \frac{U_{j,K}}{H_j} (U_{j,KC_j} U_{j,C_j H_j} - U_{j,C_j C_j} U_{j,KH_j}) \\ \eta_{j,c,w} &\equiv \left. \frac{\partial C_j}{\partial W_j} \frac{W_j}{C_j} \right|_{\lambda\text{-const.}} = D_j^{-1} \frac{U_{j,H_j}}{C_j} (U_{j,KH_j} U_{j,KC_j} - U_{j,KK} U_{j,C_j H_j}) \\ \eta_{j,c,p^c} &\equiv \left. \frac{\partial C_j}{\partial P^c} \frac{P^c}{C_j} \right|_{\lambda\text{-const.}} = D_j^{-1} \frac{U_{j,C_j}}{C_j} (U_{j,KK} U_{j,H_j H_j} - U_{j,KH_j}^2) \\ \eta_{j,c,p^k} &\equiv \left. \frac{\partial C_j}{\partial P^k} \frac{P^k}{C_j} \right|_{\lambda\text{-const.}} = D_j^{-1} \frac{U_{j,K}}{C_j} (U_{j,KH_j} U_{j,C_j H_j} - U_{j,KC_j} U_{j,H_j H_j}) \\ \eta_{j,k,w} &\equiv \left. \frac{\partial K}{\partial W_j} \frac{W_j}{K} \right|_{\lambda\text{-const.}}^j = D_j^{-1} \frac{U_{j,H_j}}{K} (U_{j,KC_j} U_{j,C_j H_j} - U_{j,KH_j} U_{j,C_j C_j})\end{aligned}$$

$$\begin{aligned}\eta_{j,k,p^c} &\equiv \left. \frac{\partial K}{\partial P^c} \frac{P^c}{K} \right|_{\lambda\text{-const.}}^j = D_j^{-1} \frac{U_{j,C_j}}{K} (U_{j,KH_j} U_{j,C_j H_j} - U_{j,KC_j} U_{j,H_j H_j}) \\ \eta_{j,k,p^k} &\equiv \left. \frac{\partial K}{\partial P^k} \frac{P^k}{K} \right|_{\lambda\text{-const.}}^j = D_j^{-1} \frac{U_{j,K}}{K} (U_{j,C_j C_j} U_{j,H_j H_j} - U_{j,C_j H_j}^2) \end{aligned}$$

The partial effects are calculated at the individual level (*not* the household level) keeping everything else fixed and λ constant in expectation. D_j is the determinant of the hessian matrix of U_j and is given by

$$\begin{aligned} D_j &= U_{j,KK} \left(U_{j,C_j C_j} U_{j,H_j H_j} - U_{j,C_j H_j}^2 \right) \\ &\quad - U_{j,KC_j}^2 U_{j,H_j H_j} - U_{j,KH_j}^2 U_{j,C_j C_j} + 2U_{j,KH_j} U_{j,KC_j} U_{j,C_j H_j}. \end{aligned}$$

U_{j,x_j} is the first-order and $U_{j,x_j \chi_j}$ is the second-order partial derivative of U_j with respect to x_j , $\chi_j = \{K, C_j, H_j\}$, $j = \{1, 2\}$. The partial derivatives are i - and t -specific but I suppress such subscripts to ease the notation.

From Phelps (1974, section 2.4) the matrix of substitution effects after a marginal-utility-of-wealth-compensated price change is

$$\begin{pmatrix} -\frac{dH_j}{dW_j} & -\frac{dC_j}{dW_j} & -\frac{dK}{dW_j} \\ \frac{dH_j}{dP^c} & \frac{dC_j}{dP^c} & \frac{dK}{dP^c} \\ \frac{dH_j}{dP^k} & \frac{dC_j}{dP^k} & \frac{dK}{dP^k} \end{pmatrix} = \lambda_j \mathbf{I}_3 \mathbf{H}_j^{-1} \quad (\text{B.1})$$

where λ_j is the Lagrange multiplier on the sequential budget constraint from the budget-constrained maximization of individual j 's utility function U_j ('individual problem'). \mathbf{I}_3 is a 3×3 identity matrix and \mathbf{H}_j is the hessian matrix of U_j . One can obtain the matrix of substitution effects by totally differentiating the intra-temporal first order conditions of the individual problem with respect to prices and noting that $\Delta\lambda = 0$ in expectation.¹

As the right hand side of (B.1) is a 3×3 symmetric matrix (the Hessian is symmetric by Young's theorem and standard regularity assumptions on U_j), it follows that $\frac{dH_j}{dP^c} = -\frac{dC_j}{dW_j}$, $\frac{dH_j}{dP^k} = -\frac{dK}{dW_j}$, and $\frac{dC_j}{dP^k} = \frac{dK}{dP^c}$. Simple manipulations

¹The first order conditions of the individual problem are $U_{j,K} = \lambda_j P^k$, $U_{j,C_j} = \lambda_j P^c$, $U_{j,H_j} = -\lambda_j W_j$, and the Euler equation.

yield the restrictions on the Frisch elasticities that are used for identification of the household structure (sections 2.2.4.1-2.2.4.2).

B.2 Approximations to First Order Conditions

In this appendix I show how I approximate the first order conditions of the household problem, given by (2.8). For convenience suppose that $U_j(K_{it}, C_{jit}, 1 - H_{jit}; \mathbf{z}_{jit}) = U_j(\tilde{K}_{it}, \tilde{C}_{jit}, 1 - \tilde{H}_{jit})$ where $\tilde{K}_{it} = K_{it} \exp(-\sum_j \mathbf{z}_{jit}^K \zeta_{jt}^K)$, $\tilde{C}_{jit} = C_{jit} \exp(-\mathbf{z}_{jit}^C \zeta_{jt}^C)$, and $\tilde{H}_{jit} = H_{jit} \exp(-\mathbf{z}_{jit}^H \zeta_{jt}^H)$. \mathbf{z}_{jit}^K , \mathbf{z}_{jit}^C , and \mathbf{z}_{jit}^H are spouse-specific covariates, such as race, age, or education.

Consider the first order condition for H_{1it} ; applying logs and taking a first difference in time yields

$$\Delta \ln \left(-U_{1,H_1}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}) \right) = \Delta \ln \lambda_{it} + \Delta \ln W_{1it} - \Delta \ln \mu_{1it}.$$

A first order Taylor approximation to $\ln(-U_{1,H_1}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}))$ about \tilde{K}_{it-1} , \tilde{C}_{1it-1} , \tilde{H}_{1it-1} yields

$$\Delta \ln \left(-U_{1,H_1}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}) \right) \approx \frac{U_{1,H_1K}}{U_{1,H_1}} \Delta \tilde{K}_{it} + \frac{U_{1,H_1C_1}}{U_{1,H_1}} \Delta \tilde{C}_{1it} + \frac{U_{1,H_1H_1}}{U_{1,H_1}} \Delta \tilde{H}_{1it}. \quad (\text{B.2})$$

All partial derivatives in (B.2) are evaluated at \tilde{K}_{it-1} , \tilde{C}_{1it-1} , and \tilde{H}_{1it-1} .

The first order conditions for H_{2it} , C_{1it} , and C_{2it} are all approximated in a similar way. All four of them together are

$$\frac{U_{1,H_1K}}{U_{1,H_1}} \Delta \tilde{K}_{it} + \frac{U_{1,H_1C_1}}{U_{1,H_1}} \Delta \tilde{C}_{1it} + \frac{U_{1,H_1H_1}}{U_{1,H_1}} \Delta \tilde{H}_{1it} \approx \Delta \ln \lambda_{it} + \Delta \ln W_{1it} - \Delta \ln \mu_{1it} \quad (\text{B.3})$$

$$\frac{U_{2,H_2K}}{U_{2,H_2}} \Delta \tilde{K}_{it} + \frac{U_{2,H_2C_2}}{U_{2,H_2}} \Delta \tilde{C}_{2it} + \frac{U_{2,H_2H_2}}{U_{2,H_2}} \Delta \tilde{H}_{2it} \approx \Delta \ln \lambda_{it} + \Delta \ln W_{2it} - \Delta \ln \mu_{2it} \quad (\text{B.4})$$

$$\frac{U_{1,C_1K}}{U_{1,C_1}} \Delta \tilde{K}_{it} + \frac{U_{1,C_1C_1}}{U_{1,C_1}} \Delta \tilde{C}_{1it} + \frac{U_{1,C_1H_1}}{U_{1,C_1}} \Delta \tilde{H}_{1it} \approx \Delta \ln \lambda_{it} + \Delta \ln P_t - \Delta \ln \mu_{1it} \quad (\text{B.5})$$

$$\frac{U_{2,C_2K}}{U_{2,C_2}} \Delta \tilde{K}_{it} + \frac{U_{2,C_2C_2}}{U_{2,C_2}} \Delta \tilde{C}_{2it} + \frac{U_{2,C_2H_2}}{U_{2,C_2}} \Delta \tilde{H}_{2it} \approx \Delta \ln \lambda_{it} + \Delta \ln P_t - \Delta \ln \mu_{2it}. \quad (\text{B.6})$$

As I do not observe cross-sectional variation in P_t , I assume $\Delta \ln P_t = 0$ (or equivalently, aggregate shocks are captured by conditioning year dummies).

The approximation to the first order condition for K_{it} is trickier as it involves both partners' marginal utilities. Applying logs and taking a first difference in time yields

$$\Delta \ln (\mu_{1it} U_{1,K}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}) + \mu_{2it} U_{2,K}(\tilde{K}_{it}, \tilde{C}_{2it}, \tilde{H}_{2it})) = \Delta \ln \lambda_{it}.$$

I approximate $\ln (\mu_{1it} U_{1,K}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}) + \mu_{2it} U_{2,K}(\tilde{K}_{it}, \tilde{C}_{2it}, \tilde{H}_{2it}))$ about \tilde{K}_{it-1} , \tilde{C}_{jit-1} , \tilde{H}_{jit-1} , $j = \{1, 2\}$, and I get

$$\begin{aligned} & \Delta \ln \left(\mu_{1it} U_{1,K}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}) + \mu_{2it} U_{2,K}(\tilde{K}_{it}, \tilde{C}_{2it}, \tilde{H}_{2it}) \right) \\ & \approx (\mu_{1it-1} U_{1,K} + \mu_{2it-1} U_{2,K})^{-1} \\ & \quad \times \left(\Delta (\mu_{1it} U_{1,K}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it})) + \Delta (\mu_{2it} U_{2,K}(\tilde{K}_{it}, \tilde{C}_{2it}, \tilde{H}_{2it})) \right) \\ & \approx \nu_{it-1} (\mu_{1it-1} U_{1,K})^{-1} \Delta (\mu_{1it} U_{1,K}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it})) \\ & \quad + (1 - \nu_{it-1}) (\mu_{2it-1} U_{2,K})^{-1} \Delta (\mu_{2it} U_{2,K}(\tilde{K}_{it}, \tilde{C}_{2it}, \tilde{H}_{2it})) \\ & \approx \nu_{it-1} \left(\Delta \ln \mu_{1it} + \Delta \ln U_{1,K}(\tilde{K}_{it}, \tilde{C}_{1it}, \tilde{H}_{1it}) \right) \\ & \quad + (1 - \nu_{it-1}) \left(\Delta \ln \mu_{2it} + \Delta \ln U_{2,K}(\tilde{K}_{it}, \tilde{C}_{2it}, \tilde{H}_{2it}) \right) \end{aligned}$$

where $\nu_{it-1} = \mu_{1it-1} U_{1,K} (\mu_{1it-1} U_{1,K} + \mu_{2it-1} U_{2,K})^{-1}$ is a mixture of preferences (marginal utilities of the public good) and bargaining power. Again, all partial derivatives are evaluated at $t - 1$ values of the relevant variables. Expanding $\ln U_{j,K}(\tilde{K}_{it}, \tilde{C}_{jit}, \tilde{H}_{jit})$, $j = \{1, 2\}$, follows the same logic as in (B.2). Combining

the results I get

$$\begin{aligned}
& \left(\nu_{it-1} \frac{U_{1,KK}}{U_{1,K}} + (1 - \nu_{it-1}) \frac{U_{2,KK}}{U_{2,K}} \right) \Delta \tilde{K}_{it} + \\
& \nu_{it-1} \left(\frac{U_{1,KC_1}}{U_{1,K}} \Delta \tilde{C}_{1it} + \frac{U_{1,KH_1}}{U_{1,K}} \Delta \tilde{H}_{1it} \right) + \\
& (1 - \nu_{it-1}) \left(\frac{U_{2,KC_2}}{U_{2,K}} \Delta \tilde{C}_{2it} + \frac{U_{2,KH_2}}{U_{2,K}} \Delta \tilde{H}_{2it} \right) \approx \\
& \Delta \ln \lambda_{it} - \nu_{it-1} \Delta \ln \mu_{1it} - (1 - \nu_{it-1}) \Delta \ln \mu_{2it}.
\end{aligned} \tag{B.7}$$

I solve the system of equations (B.3)-(B.7) for $\Delta \tilde{K}_{it}$, $\Delta \tilde{C}_{jit}$, and $\Delta \tilde{H}_{jit}$ ($j = \{1, 2\}$, 5 equations in 5 choice variables) and I get approximate analytical expressions for the growth of labor supply and consumption as functions of the growth rates of wages, bargaining power, and the marginal utility of wealth. Given that $\mu_{1it} + \mu_{2it} = 1$ (see section 2.2.1.1), I approximate $\Delta \ln \mu_{2it} \approx -\tilde{\mu}_{it-1} \Delta \ln \mu_{1it}$ where $\tilde{\mu}_{it} = \frac{\mu_{1it}}{\mu_{2it}}$, so the aforementioned analytical expressions become

$$\begin{pmatrix} \Delta \tilde{K}_{it} \\ \Delta \tilde{C}_{1it} \\ \Delta \tilde{C}_{2it} \\ \Delta \tilde{H}_{1it} \\ \Delta \tilde{H}_{2it} \end{pmatrix} \approx \mathbf{M}_{it} \begin{pmatrix} \Delta \ln W_{1it} \\ \Delta \ln W_{2it} \\ \Delta \ln \mu_{1it} \\ \Delta \ln \lambda_{it} \end{pmatrix} \tag{B.8}$$

where \mathbf{M}_{it} is a 5×4 matrix of loading factors.²

Equation (B.8) is still not very useful empirically because it involves the unobserved growth in the marginal utility of wealth $\Delta \ln \lambda_{it}$. To help characterize this object I apply a second order Taylor approximation to the Euler equation in (2.8). Let $\exp(\rho) = \frac{1}{\beta(1+r)}$ for an appropriate ρ . I approximate the exponential function evaluated at λ_{it+1} about $\ln \lambda_{it} + \rho$ and I get

$$\begin{aligned}
\exp(\ln \lambda_{it+1}) & \approx \exp(\ln \lambda_{it} + \rho) + \exp(\ln \lambda_{it} + \rho) (\Delta \ln \lambda_{it+1} - \rho) \\
& + \exp(\ln \lambda_{it} + \rho) \frac{1}{2} (\Delta \ln \lambda_{it+1} - \rho)^2
\end{aligned}$$

²As matrix \mathbf{M}_{it} constitutes an intermediate output towards the solution of the household problem, I do not present its elements in detail but I can make them available upon request.

$$= \lambda_{it} \exp(\rho) [1 + \Delta \ln \lambda_{it+1} - \rho + \frac{1}{2} (\Delta \ln \lambda_{it+1} - \rho)^2].$$

Taking expectations at time t and noting that $\lambda_{it} \exp(\rho) = \mathbb{E}_t \lambda_{it+1}$ (the Euler equation) yields $\mathbb{E}_t \Delta \ln \lambda_{it+1} \approx \rho - \frac{1}{2} \mathbb{E}_t (\Delta \ln \lambda_{it+1} - \rho)^2$ which in turn can be written as

$$\Delta \ln \lambda_{it+1} \approx \omega_{it+1} + \epsilon_{it+1} \quad (\text{B.9})$$

with $\omega_{it+1} = \rho - \frac{1}{2} \mathbb{E}_t (\Delta \ln \lambda_{it+1} - \rho)^2$ and ϵ_{it+1} an expectations error with $\mathbb{E}_t \epsilon_{it+1} = 0$. A log-linearization of the intertemporal budget constraint, presented in appendix B.3, will help relate the components of $\Delta \ln \lambda_{it+1}$ to spousal wage shocks.

Using equations (2.10) in the main text and (B.9) above, together with the assumption that transitory wage shocks do not shift intra-family bargaining power, I write equation (B.8) as

$$\begin{pmatrix} \Delta \tilde{K}_{it} \\ \Delta \tilde{C}_{1it} \\ \Delta \tilde{C}_{2it} \\ \Delta \tilde{Y}_{1it} \\ \Delta \tilde{Y}_{2it} \end{pmatrix} \approx \begin{pmatrix} \alpha_{k,w_1} + \beta_{k,w_1} & \alpha_{k,w_2} + \beta_{k,w_2} & \alpha_{k,w_1} & \alpha_{k,w_2} & \alpha_{k,\lambda} \\ \alpha_{c_1,w_1} + \beta_{c_1,w_1} & \alpha_{c_1,w_2} + \beta_{c_1,w_2} & \alpha_{c_1,w_1} & \alpha_{c_1,w_2} & \alpha_{c_1,\lambda} \\ \alpha_{c_2,w_1} + \beta_{c_2,w_1} & \alpha_{c_2,w_2} + \beta_{c_2,w_2} & \alpha_{c_2,w_1} & \alpha_{c_2,w_2} & \alpha_{c_2,\lambda} \\ \alpha_{y_1,w_1} + \beta_{y_1,w_1} & \alpha_{y_1,w_2} + \beta_{y_1,w_2} & \alpha_{y_1,w_1} & \alpha_{y_1,w_2} & \alpha_{y_1,\lambda} \\ \alpha_{y_2,w_1} + \beta_{y_2,w_1} & \alpha_{y_2,w_2} + \beta_{y_2,w_2} & \alpha_{y_2,w_1} & \alpha_{y_2,w_2} & \alpha_{y_2,\lambda} \end{pmatrix} \begin{pmatrix} v_{1it} \\ v_{2it} \\ \Delta u_{1it} \\ \Delta u_{2it} \\ \omega_{it} + \epsilon_{it} \end{pmatrix} \quad (\text{B.10})$$

where I replace hours of market work by earnings using the identity $Y_{jit} = W_{jit} H_{jit}$, $j = \{1, 2\}$. To obtain (B.10) I have assumed away variation in any distribution factor in (2.10). The transmission parameters β above, which reflect the difference between the impact of permanent and transitory wage shocks on household outcomes (conditional on $\Delta \ln \lambda_{it}$), capture the *bargaining* effects of permanent wages shocks, *i.e.* the effects through shifting intra-family bargaining power. The transmission parameters are reported analytically in appendix B.4.

B.3 Approximation to Intertemporal Budget Constraint

Let $F(\boldsymbol{\psi}) = \ln \sum_{s=0}^J \exp \psi_s$ with $\boldsymbol{\psi} = (\psi_0, \psi_1, \dots, \psi_J)'$; a first order Taylor approximation of $F(\boldsymbol{\psi})$ around $\boldsymbol{\psi}^0$ yields

$$F(\boldsymbol{\psi}) \approx \ln \sum_{s=0}^J \exp \psi_s^0 + \sum_{s=0}^J \frac{\exp \psi_s^0}{\sum_{s=0}^J \exp \psi_s^0} (\psi_s - \psi_s^0).$$

Consider the left hand side of the intertemporal budget constraint implied by (2.5); that is given by $A_{i0} + \mathbb{E}_0 \sum_{t=0}^T \frac{W_{1it}H_{1it}}{(1+r)^t} + \mathbb{E}_0 \sum_{t=0}^T \frac{W_{2it}H_{2it}}{(1+r)^t}$. Making this look similar to $F(\boldsymbol{\psi})$ above, I define

$$F^{LH}(\boldsymbol{\psi}) = \ln \left[\exp(\ln A_{i0}) + \sum_{t=0}^T \exp \left(\ln \frac{W_{1it}H_{1it}}{(1+r)^t} \right) + \sum_{t=0}^T \exp \left(\ln \frac{W_{2it}H_{2it}}{(1+r)^t} \right) \right]$$

where, in this case, $J = 2T + 2$ and

$$\psi_s = \begin{cases} \ln A_{is} & s = 0 \\ \ln W_{1is-1}H_{1is-1} - (s-1) \ln(1+r) & s = 1, \dots, T+1 \\ \ln W_{2is-(T+2)}H_{2is-(T+2)} - (s-(T+2)) \ln(1+r) & s = T+2, \dots, 2T+2. \end{cases}$$

Also define

$$\psi_s^0 = \begin{cases} \mathbb{E}_{-1} \ln A_{is} & s = 0 \\ \mathbb{E}_{-1} \ln W_{1is-1}H_{1is-1} - (s-1) \ln(1+r) & s = 1, \dots, T+1 \\ \mathbb{E}_{-1} \ln W_{2is-(T+2)}H_{2is-(T+2)} - (s-(T+2)) \ln(1+r) & s = T+2, \dots, 2T+2 \end{cases}$$

$$D_{i0} = \exp(\mathbb{E}_{-1} \ln A_{i0}) + \sum_{k=0}^T \exp \left(\mathbb{E}_{-1} \ln \frac{W_{1ik}H_{1ik}}{(1+r)^k} \right) + \sum_{k=0}^T \exp \left(\mathbb{E}_{-1} \ln \frac{W_{2ik}H_{2ik}}{(1+r)^k} \right)$$

$$\pi_{i0} = \frac{\exp(\mathbb{E}_{-1} \ln A_{i0})}{D_{i0}}$$

$$s_{i0} = \frac{\sum_{k=0}^T \exp(\mathbb{E}_{-1} \ln W_{1ik}H_{1ik} - k \ln(1+r))}{D_{i0} - \exp(\mathbb{E}_{-1} \ln A_{i0})}$$

$$\theta_{jit} = \frac{\exp(\mathbb{E}_{-1} \ln W_{jit}H_{jit} - t \ln(1+r))}{\sum_{k=0}^T \exp(\mathbb{E}_{-1} \ln W_{jik}H_{jik} - k \ln(1+r))}$$

where \mathbb{E}_{-1} defines expectations at time $t = -1$ (while the intertemporal budget constraint covers periods 0 through T).

Parameter π_{i0} is approximately equal to the ($t = -1$)-expected ratio of financial wealth at $t = 0$ over total human and financial wealth in the household over the period between $t = 0$ and $t = T$.³ s_{i0} is the ($t = -1$)-expected ratio of the male spouse's lifetime human wealth over total lifetime human wealth in the household. θ_{jit} , $j = \{1, 2\}$, is the ($t = -1$)-expected ratio of spouse j 's earnings at t over the same person's lifetime human wealth.

Expanding the left hand side of the intertemporal budget constraint around $\boldsymbol{\psi}^0$ and taking expectations conditional on an information set I , where such information set includes information known at $t = -1$, I get

$$\begin{aligned} \mathbb{E}_I F^{LH}(\boldsymbol{\psi}) &\approx \mathbb{E}_I F^{LH}(\boldsymbol{\psi}^0) + \frac{\exp(\mathbb{E}_{-1} \ln A_{i0})}{D_{i0}} (\mathbb{E}_I - \mathbb{E}_{-1}) \ln A_{i0} \\ &\quad + \sum_{t=0}^T \frac{\exp(\mathbb{E}_{-1} \ln W_{1it} H_{1it} - t \ln(1+r))}{D_{i0}} (\mathbb{E}_I - \mathbb{E}_{-1}) \ln W_{1it} H_{1it} \\ &\quad + \sum_{t=0}^T \frac{\exp(\mathbb{E}_{-1} \ln W_{2it} H_{2it} - t \ln(1+r))}{D_{i0}} (\mathbb{E}_I - \mathbb{E}_{-1}) \ln W_{2it} H_{2it} \\ \mathbb{E}_I F^{LH}(\boldsymbol{\psi}) &\approx \mathbb{E}_I F^{LH}(\boldsymbol{\psi}^0) + \pi_{i0} (\mathbb{E}_I - \mathbb{E}_{-1}) \ln A_{i0} \\ &\quad + (1 - \pi_{i0}) s_{i0} \sum_{t=0}^T \theta_{1it} (\mathbb{E}_I - \mathbb{E}_{-1}) (\Delta \ln Y_{1it} + \ln Y_{1it-1}) \\ &\quad + (1 - \pi_{i0}) (1 - s_{i0}) \sum_{t=0}^T \theta_{2it} (\mathbb{E}_I - \mathbb{E}_{-1}) (\Delta \ln Y_{2it} + \ln Y_{2it-1}), \end{aligned}$$

where $(\mathbb{E}_I - \mathbb{E}_{-1})\chi = \mathbb{E}_I(\chi) - \mathbb{E}_{-1}(\chi)$ for a generic variable χ .

Replacing $\Delta \ln Y_{jit}$ by its analytical expression in (B.10) (note that $(\mathbb{E}_I - \mathbb{E}_{-1})\Delta \ln Y_{jit} = (\mathbb{E}_I - \mathbb{E}_{-1})\Delta \tilde{Y}_{jit}$ because changes in the observable characteristics are in the spouses' information sets and therefore canceled out by the first difference $\mathbb{E}_I - \mathbb{E}_{-1}$), assuming that each spouse's current earnings are negligible compared to their expected lifetime earnings,⁴ placing ω_{it} in the spouses'

³By 'lifetime human wealth' I refer to expected lifetime earnings over $t = 0, \dots, T$.

⁴This implies that $\theta_{jit} \approx 0$, $j = \{1, 2\}$, and transitory shocks to current earnings do not shift the intertemporal budget constraint.

information sets, and taking a first difference of $\mathbb{E}_I F^{LH}(\boldsymbol{\psi})$ across $I : t = 0$ and $I : t = -1$, I get

$$\begin{aligned} & (\mathbb{E}_0 - \mathbb{E}_{-1})F^{LH}(\boldsymbol{\psi}) \\ & \approx (1 - \pi_{i0})s_{i0}((\alpha_{y_1,w_1} + \beta_{y_1,w_1})v_{1i0} + (\alpha_{y_1,w_2} + \beta_{y_1,w_2})v_{2i0} + \alpha_{y_1,\lambda}\epsilon_{i0}) \\ & + (1 - \pi_{i0})(1 - s_{i0})((\alpha_{y_2,w_1} + \beta_{y_2,w_1})v_{1i0} + (\alpha_{y_2,w_2} + \beta_{y_2,w_2})v_{2i0} + \alpha_{y_2,\lambda}\epsilon_{i0}). \end{aligned}$$

Similarly for the right hand side of the intertemporal budget constraint, define

$$\begin{aligned} F^{RH}(\boldsymbol{\psi}) &= \ln \left[\sum_{t=0}^T \exp \left(\ln \frac{K_{it}}{(1+r)^t} \right) + \sum_{t=0}^T \exp \left(\ln \frac{P_t C_{it}}{(1+r)^t} \right) \right] \\ \psi_s &= \begin{cases} \ln K_{is} - s \ln(1+r) & s = 0, \dots, T \\ \ln P_{s-(T+1)} C_{is-(T+1)} - (s - (T+1)) \ln(1+r) & s = T+1, \dots, 2T+1. \end{cases} \end{aligned}$$

Also define

$$\xi_{it} = \frac{\sum_{k=0}^T \exp(\mathbb{E}_{-1} \ln \frac{K_{ik}}{(1+r)^k})}{\sum_{k=0}^T \exp(\mathbb{E}_{-1} \ln \frac{K_{ik}}{(1+r)^k}) + \sum_{k=0}^T \exp(\mathbb{E}_{-1} \ln \frac{P_k C_{ik}}{(1+r)^k})}$$

as the ($t = -1$)-expected share of lifetime family expenditure on the public good over total lifetime family expenditure on all goods between $t = 0$ and $t = T$.

I follow the same procedure as for the left-hand side and I get

$$\begin{aligned} & (\mathbb{E}_0 - \mathbb{E}_{-1})F^{RH}(\boldsymbol{\psi}) \\ & \approx \xi_{i0}((\alpha_{k,w_1} + \beta_{k,w_1})v_{1i0} + (\alpha_{k,w_2} + \beta_{k,w_2})v_{2i0} + \alpha_{k,\lambda}\epsilon_{i0}) \\ & + (1 - \xi_{i0})\varphi_{i,-1}((\alpha_{c_1,w_1} + \beta_{c_1,w_1})v_{1i0} + (\alpha_{c_1,w_2} + \beta_{c_1,w_2})v_{2i0} + \alpha_{c_1,\lambda}\epsilon_{i0}) \\ & + (1 - \xi_{i0})(1 - \varphi_{i,-1})((\alpha_{c_2,w_1} + \beta_{c_2,w_1})v_{1i0} + (\alpha_{c_2,w_2} + \beta_{c_2,w_2})v_{2i0} + \alpha_{c_2,\lambda}\epsilon_{i0}). \end{aligned}$$

Here $\varphi_{it} = \frac{C_{1it}}{C_{it}}$ is the male spouse's share of private consumption at t . φ_{it} enters the result because $\Delta \ln C_{it} \approx \varphi_{it-1} \Delta \ln C_{1it} + (1 - \varphi_{it-1}) \Delta \ln C_{2it}$.

Noting that the budget constraint must balance, I bring the two sides

together and solve for ϵ_{i0} to get

$$\epsilon_{i0} \approx \ell_{i0}^{w_1} v_{1i0} + \ell_{i0}^{w_2} v_{2i0} \quad (\text{B.11})$$

where

$$\begin{aligned} \ell_{i0}^{w_1} &= (1/\ell_{i0}^\lambda) \left(\xi_{i0}(\alpha_{k,w_1} + \beta_{k,w_1}) \right. \\ &\quad + (1 - \xi_{i0})(\varphi_{i,-1}(\alpha_{c_1,w_1} + \beta_{c_1,w_1}) + (1 - \varphi_{i,-1})(\alpha_{c_2,w_1} + \beta_{c_2,w_1})) \\ &\quad \left. - (1 - \pi_{i0})s_{i0}(\alpha_{y_1,w_1} + \beta_{y_1,w_1}) - (1 - \pi_{i0})(1 - s_{i0})(\alpha_{y_2,w_1} + \beta_{y_2,w_1}) \right) \\ \ell_{i0}^{w_2} &= (1/\ell_{i0}^\lambda) \left(\xi_{i0}(\alpha_{k,w_2} + \beta_{k,w_2}) \right. \\ &\quad + (1 - \xi_{i0})(\varphi_{i,-1}(\alpha_{c_1,w_2} + \beta_{c_1,w_2}) + (1 - \varphi_{i,-1})(\alpha_{c_2,w_2} + \beta_{c_2,w_2})) \\ &\quad \left. - (1 - \pi_{i0})s_{i0}(\alpha_{y_1,w_2} + \beta_{y_1,w_2}) - (1 - \pi_{i0})(1 - s_{i0})(\alpha_{y_2,w_2} + \beta_{y_2,w_2}) \right) \\ \ell_{i0}^\lambda &= (1 - \pi_{i0})(s_{i0}\alpha_{y_1,\lambda} + (1 - s_{i0})\alpha_{y_2,\lambda}) \\ &\quad - \xi_{i0}\alpha_{k,\lambda} - (1 - \xi_{i0})(\varphi_{i,-1}\alpha_{c_1,\lambda} + (1 - \varphi_{i,-1})\alpha_{c_2,\lambda}). \end{aligned}$$

For a generic period t the mapping between ϵ_{it} and the permanent shocks v_{1it} and v_{2it} looks alike. One has to follow the same steps, the only difference being that the budget constraint must start counting at t (rather than 0) and the difference in expectations must be $\mathbb{E}_t - \mathbb{E}_{t-1}$.

Similar arguments can be used to show that an approximation to single j 's intertemporal budget constraint, the sequential version of which appears in (2.15), results in a mapping from permanent shock v_{jis} to the innovation to the single's marginal utility of wealth ϵ_{jis} , given by $\epsilon_{jis} = \ell_{jit}^s v_{jis}$. Here

$$\ell_{jit}^s = \frac{\xi_{jis}\eta_{j,k,w} + (1 - \xi_{jis})\eta_{j,c,w} - (1 - \pi_{jis})(1 + \eta_{j,h,w})}{(1 - \pi_{jis}) \sum \eta_{j,h} - \xi_{jis} \sum \eta_{j,k} - (1 - \xi_{jis}) \sum \eta_{j,c}} \quad (\text{B.12})$$

and the notation is as follows: $\sum \eta_{j,h} = \eta_{j,h,w} + \eta_{j,h,p^c} + \eta_{j,h,p^k}$; $\sum \eta_{j,c} = \eta_{j,c,w} + \eta_{j,c,p^c} + \eta_{j,c,p^k}$; $\sum \eta_{j,k} = \eta_{j,k,w} + \eta_{j,k,p^c} + \eta_{j,k,p^k}$; ξ_{jis} is j 's ratio of public to total expected expenditure as single; $\pi_{jis} \approx \text{Assets}_{jis}/(\text{Assets}_{jis} + \text{Earnings as Single}_{jis})$ is j 's financial wealth relative to his/her total financial

and human wealth as single.

B.4 Transmission Parameters

In this appendix I report the transmission parameters of wage shocks into household outcomes, namely the elements of matrix \mathbf{T}_{it} in (2.9). Less compactly, (2.9) can be written as

$$\begin{pmatrix} \Delta k_{it} \\ \Delta c_{it} \\ \Delta y_{1it} \\ \Delta y_{2it} \end{pmatrix} \approx \begin{pmatrix} \alpha_{k,w_1} + \beta_{k,w_1} + \gamma_{k,w_1} & \alpha_{k,w_2} + \beta_{k,w_2} + \gamma_{k,w_2} & \alpha_{k,w_1} & \alpha_{k,w_2} \\ \alpha_{c,w_1} + \beta_{c,w_1} + \gamma_{c,w_1} & \alpha_{c,w_2} + \beta_{c,w_2} + \gamma_{c,w_2} & \alpha_{c,w_1} & \alpha_{c,w_2} \\ \alpha_{y_1,w_1} + \beta_{y_1,w_1} + \gamma_{y_1,w_1} & \alpha_{y_1,w_2} + \beta_{y_1,w_2} + \gamma_{y_1,w_2} & \alpha_{y_1,w_1} & \alpha_{y_1,w_2} \\ \alpha_{y_2,w_1} + \beta_{y_2,w_1} + \gamma_{y_2,w_1} & \alpha_{y_2,w_2} + \beta_{y_2,w_2} + \gamma_{y_2,w_2} & \alpha_{y_2,w_1} & \alpha_{y_2,w_2} \end{pmatrix} \begin{pmatrix} v_{1it} \\ v_{2it} \\ \Delta u_{1it} \\ \Delta u_{2it} \end{pmatrix} \quad (\text{B.13})$$

where \mathbf{T}_{it} is the 4×4 matrix on the right-hand side.

I obtain (B.13) directly from (B.10) after replacing ϵ_{it} by (B.11) ($\omega_{it} \equiv \omega_t$ is absorbed by the conditioning observables) and collapsing individual private consumption into *total* private consumption using (2.11). The left-hand side vectors of outcome variables in (B.10) and (B.13) are equivalent because $\Delta k_{it} = \Delta \tilde{K}_{it} = \Delta \ln K_{it} - \sum_j \Delta(\mathbf{z}_{jit}^K \boldsymbol{\zeta}_{jt}^K) - \sum_j \Delta(\alpha_{k,w_j} \mathbf{x}_{jit}^{W'} \boldsymbol{\zeta}_{jt}^W)$ (and similarly for the other variables).

Permanent and transitory wage shocks induce static effects on outcomes; these effects are captured by the α 's. The static effects summarize the standard substitution effects that a price change (in this case a wage change) induces through tilting the *static* budget constraint. In addition, permanent shocks induce bargaining effects through their impact on intra-family bargaining power; these effects are captured by the β 's. Finally, permanent shocks also induce wealth/income effects through shifting the *intertemporal*, and therefore also the *static*, budget constraint and impacting on the revisions to the marginal utility of wealth λ ; such effects are captured by the γ 's.

Static effects

$$\begin{aligned}
\text{On } \Delta k_{it} &: \alpha_{k,w_1} = \nu_{it-1} \eta_{2,k,p^k} \eta_{1,k,w} (1/\bar{\eta}_{k,p^k}) \\
&: \alpha_{k,w_2} = (1 - \nu_{it-1}) \eta_{1,k,p^k} \eta_{2,k,w} (1/\bar{\eta}_{k,p^k}) \\
\text{On } \Delta c_{1it} &: \alpha_{c_1,w_1} = \eta_{1,c,w} - (1 - \nu_{it-1}) \eta_{1,c,p^k} \eta_{1,k,w} (1/\bar{\eta}_{k,p^k}) \\
&: \alpha_{c_1,w_2} = (1 - \nu_{it-1}) \eta_{1,c,p^k} \eta_{2,k,w} (1/\bar{\eta}_{k,p^k}) \\
\text{On } \Delta c_{2it} &: \alpha_{c_2,w_1} = \nu_{it-1} \eta_{2,c,p^k} \eta_{1,k,w} (1/\bar{\eta}_{k,p^k}) \\
&: \alpha_{c_2,w_2} = \eta_{2,c,w} - \nu_{it-1} \eta_{2,c,p^k} \eta_{2,k,w} (1/\bar{\eta}_{k,p^k}) \\
\text{On } \Delta c_{it} &: \alpha_{c,w_1} = \varphi_{it-1} \alpha_{c_1,w_1} + (1 - \varphi_{it}) \alpha_{c_2,w_1} \\
&: \alpha_{c,w_2} = \varphi_{it-1} \alpha_{c_1,w_2} + (1 - \varphi_{it}) \alpha_{c_2,w_2} \\
\text{On } \Delta y_{1it} &: \alpha_{y_1,w_1} = 1 + \eta_{1,h,w} - (1 - \nu_{it-1}) \eta_{1,h,p^k} \eta_{1,k,w} (1/\bar{\eta}_{k,p^k}) \\
&: \alpha_{y_1,w_2} = (1 - \nu_{it-1}) \eta_{1,h,p^k} \eta_{2,k,w} (1/\bar{\eta}_{k,p^k}) \\
\text{On } \Delta y_{2it} &: \alpha_{y_2,w_1} = \nu_{it-1} \eta_{2,h,p^k} \eta_{1,k,w} (1/\bar{\eta}_{k,p^k}) \\
&: \alpha_{y_2,w_2} = 1 + \eta_{2,h,w} - \nu_{it-1} \eta_{2,h,p^k} \eta_{2,k,w} (1/\bar{\eta}_{k,p^k})
\end{aligned}$$

Bargaining effects

$$\begin{aligned}
\text{On } \Delta k_{it} &: \beta_{k,w_1} = \eta_{\mu,w_1} \beta_k \\
&: \beta_{k,w_2} = \eta_{\mu,w_2} \beta_k \\
\text{On } \Delta c_{1it} &: \beta_{c_1,w_1} = \eta_{\mu,w_1} \beta_{c_1} \\
&: \beta_{c_1,w_2} = \eta_{\mu,w_2} \beta_{c_1} \\
\text{On } \Delta c_{2it} &: \beta_{c_2,w_1} = \eta_{\mu,w_1} \beta_{c_2} \\
&: \beta_{c_2,w_2} = \eta_{\mu,w_2} \beta_{c_2} \\
\text{On } \Delta c_{it} &: \beta_{c,w_1} = \varphi_{it-1} \beta_{c_1,w_1} + (1 - \varphi_{it}) \beta_{c_2,w_1} \\
&: \beta_{c,w_2} = \varphi_{it-1} \beta_{c_1,w_2} + (1 - \varphi_{it}) \beta_{c_2,w_2} \\
\text{On } \Delta y_{1it} &: \beta_{y_1,w_1} = \eta_{\mu,w_1} \beta_{y_1} \\
&: \beta_{y_1,w_2} = \eta_{\mu,w_2} \beta_{y_1} \\
\text{On } \Delta y_{2it} &: \beta_{y_2,w_1} = \eta_{\mu,w_1} \beta_{y_2}
\end{aligned}$$

$$: \beta_{y_2, w_2} = \eta_{\mu, w_2} \beta_{y_2}$$

Wealth effects

$$\begin{aligned} \text{On } \Delta k_{it} &: \gamma_{k, w_1} = \alpha_{k, \lambda} \ell_{it}^{w_1} \\ &: \gamma_{k, w_2} = \alpha_{k, \lambda} \ell_{it}^{w_2} \\ \text{On } \Delta c_{1it} &: \gamma_{c_1, w_1} = \alpha_{c_1, \lambda} \ell_{it}^{w_1} \\ &: \gamma_{c_1, w_2} = \alpha_{c_1, \lambda} \ell_{it}^{w_2} \\ \text{On } \Delta c_{2it} &: \gamma_{c_2, w_1} = \alpha_{c_2, \lambda} \ell_{it}^{w_1} \\ &: \gamma_{c_2, w_2} = \alpha_{c_2, \lambda} \ell_{it}^{w_2} \\ \text{On } \Delta c_{it} &: \gamma_{c, w_1} = \varphi_{it-1} \gamma_{c_1, w_1} + (1 - \varphi_{it}) \gamma_{c_2, w_1} \\ &: \gamma_{c, w_2} = \varphi_{it-1} \gamma_{c_1, w_2} + (1 - \varphi_{it}) \gamma_{c_2, w_2} \\ \text{On } \Delta y_{1it} &: \gamma_{y_1, w_1} = \alpha_{y_1, \lambda} \ell_{it}^{w_1} \\ &: \gamma_{y_1, w_2} = \alpha_{y_1, \lambda} \ell_{it}^{w_2} \\ \text{On } \Delta y_{2it} &: \gamma_{y_2, w_1} = \alpha_{y_2, \lambda} \ell_{it}^{w_1} \\ &: \gamma_{y_2, w_2} = \alpha_{y_2, \lambda} \ell_{it}^{w_2} \end{aligned}$$

where

$$\begin{aligned} \bar{\eta}_{k, p^k} &= (1 - \nu_{it-1}) \eta_{1, k, p^k} + \nu_{it-1} \eta_{2, k, p^k} \\ \alpha_{k, \lambda} &= (1/\bar{\eta}_{k, p^k}) (\eta_{1, k, p^k} \eta_{2, k, p^k} + \nu_{it-1} \eta_{2, k, p^k} (\eta_{1, k, p^c} + \eta_{1, k, w})) \\ &\quad + (1 - \nu_{it-1}) \eta_{1, k, p^k} (\eta_{2, k, p^c} + \eta_{2, k, w}) \\ \alpha_{c_1, \lambda} &= \eta_{1, c, p^c} + \eta_{1, c, w} \\ &\quad + (1/\bar{\eta}_{k, p^k}) (\eta_{1, c, p^k} \eta_{2, k, p^k} + (1 - \nu_{it-1}) \eta_{1, c, p^k} (-\eta_{1, k, p^c} + \eta_{2, k, p^c} - \eta_{1, k, w} + \eta_{2, k, w})) \\ \alpha_{c_2, \lambda} &= \eta_{2, c, p^c} + \eta_{2, c, w} \\ &\quad + (1/\bar{\eta}_{k, p^k}) (\eta_{2, c, p^k} \eta_{1, k, p^k} - \nu_{it-1} \eta_{2, c, p^k} (-\eta_{1, k, p^c} + \eta_{2, k, p^c} - \eta_{1, k, w} + \eta_{2, k, w})) \\ \alpha_{y_1, \lambda} &= \eta_{1, h, w} + \eta_{1, h, p^c} \\ &\quad + (1/\bar{\eta}_{k, p^k}) (\eta_{1, h, p^k} \eta_{2, k, p^k} + (1 - \nu_{it-1}) \eta_{1, h, p^k} (-\eta_{1, k, p^c} + \eta_{2, k, p^c} - \eta_{1, k, w} + \eta_{2, k, w})) \end{aligned}$$

$$\begin{aligned}
\alpha_{y_2, \lambda} &= \eta_{2,h,w} + \eta_{2,h,p^c} \\
&\quad + (1/\bar{\eta}_{k,p^k}) (\eta_{2,h,p^k} \eta_{1,k,p^k} - \nu_{it-1} \eta_{2,h,p^k} (-\eta_{1,k,p^c} + \eta_{2,k,p^c} - \eta_{1,k,w} + \eta_{2,k,w})) \\
\beta_k &= \tilde{\nu}_{it-1} \eta_{1,k,p^k} \eta_{2,k,p^k} (1/\bar{\eta}_{k,p^k}) - \nu_{it-1} \eta_{2,k,p^k} (1/\bar{\eta}_{k,p^k}) (\eta_{1,k,p^c} + \eta_{1,k,w}) \\
&\quad + \tilde{\mu}_{it-1} (1 - \nu_{it-1}) \eta_{1,k,p^k} (1/\bar{\eta}_{k,p^k}) (\eta_{2,k,p^c} + \eta_{2,k,w}) \\
\beta_{c_1} &= \tilde{\nu}_{it-1} \eta_{1,c,p^k} \eta_{2,k,p^k} (1/\bar{\eta}_{k,p^k}) - \eta_{1,c,p^c} - \eta_{1,c,w} \\
&\quad + (1 - \nu_{it-1}) \eta_{1,c,p^k} (1/\bar{\eta}_{k,p^k}) (\eta_{1,k,p^c} + \tilde{\mu}_{it-1} \eta_{2,k,p^c} + \eta_{1,k,w} + \tilde{\mu}_{it-1} \eta_{2,k,w}) \\
\beta_{c_2} &= \tilde{\nu}_{it-1} \eta_{2,c,p^k} \eta_{1,k,p^k} (1/\bar{\eta}_{k,p^k}) + \tilde{\mu}_{it-1} \eta_{2,c,p^c} + \tilde{\mu}_{it-1} \eta_{2,c,w} \\
&\quad - \nu_{it-1} \eta_{2,c,p^k} (1/\bar{\eta}_{k,p^k}) (\eta_{1,k,p^c} + \tilde{\mu}_{it-1} \eta_{2,k,p^c} + \eta_{1,k,w} + \tilde{\mu}_{it-1} \eta_{2,k,w}) \\
\beta_{y_1} &= \tilde{\nu}_{it-1} \eta_{1,h,p^k} \eta_{2,k,p^k} (1/\bar{\eta}_{k,p^k}) - \eta_{1,h,w} - \eta_{1,h,p^c} \\
&\quad + (1 - \nu_{it-1}) \eta_{1,h,p^k} (1/\bar{\eta}_{k,p^k}) (\eta_{1,k,p^c} + \tilde{\mu}_{it-1} \eta_{2,k,p^c} + \eta_{1,k,w} + \tilde{\mu}_{it-1} \eta_{2,k,w}) \\
\beta_{y_2} &= \tilde{\nu}_{it-1} \eta_{2,h,p^k} \eta_{1,k,p^k} (1/\bar{\eta}_{k,p^k}) + \tilde{\mu}_{it-1} \eta_{2,h,w} + \tilde{\mu}_{it-1} \eta_{2,h,p^c} \\
&\quad - \nu_{it-1} \eta_{2,h,p^k} (1/\bar{\eta}_{k,p^k}) (\eta_{1,k,p^c} + \tilde{\mu}_{it-1} \eta_{2,k,p^c} + \eta_{1,k,w} + \tilde{\mu}_{it-1} \eta_{2,k,w}) \\
\tilde{\nu}_{it} &= (1 - \nu_{it}) \tilde{\mu}_{it} - \nu_{it}.
\end{aligned}$$

Recall that the Frisch elasticities are presented in table 2.1 and defined in appendix B.1; also $\nu_{it} = \mu_{1it} U_{1,K} (\mu_{1it} U_{1,K} + \mu_{2it} U_{2,K})^{-1}$, $\tilde{\mu}_{it} = \frac{\mu_{1it}}{\mu_{2it}}$, and $\ell_{it}^{w_1}$ & $\ell_{it}^{w_2}$ are defined in appendix B.3.

Whenever a transmission parameter involves ν , μ , ℓ^{w_1} , ℓ^{w_2} , or functions of them, then such parameter must be i - and t -specific. Such subscripts are removed here for simplicity of the notation; they are, however, implied.

B.5 Measurement Error in Consumption

Table B.1: Estimates of consumption measurement error

	(1) separable	(2) nonseparable
<i>Couples</i>		
c	0.0808 (0.0032)	0.0827 (0.0030)
k	0.0404 (0.0023)	0.0339 (0.0038)
<i>Single males</i>		
c_1		0.0557 (0.0126)
k_1		0.0561 (0.0093)
<i>Single females</i>		
c_2		0.0554 (0.0144)
k_2		0.0474 (0.0097)

Notes: The table presents the GMM estimates of the variance of measurement error in private and public consumption. Block bootstrap standard errors are in parentheses based on 1,000 bootstrap replications. Column 1 presents the estimates when the public good is additively separable from private consumption and leisure; column 2 presents the estimates when it is nonseparable.

Appendix C

Appendix to Consumption Inequality across Heterogeneous Families

C.1 Approximations to First Order Conditions

Utility from consumption and disutility from labor are affected by observed taste shifters \mathbf{Z}_{it} such as education or age of spouses. Suppose the effect of such taste shifters enters utility as $U_i(C_{it}, \mathbf{Z}_{it}) = \tilde{U}_i(C_{it} \exp(-\mathbf{Z}'_{it}\boldsymbol{\eta}^C))$ and $V_i(H_{2it}, \mathbf{Z}_{it}) = \tilde{V}_i(H_{2it} \exp(-\mathbf{Z}'_{it}\boldsymbol{\eta}^H))$. For simplicity, denote $\tilde{C}_{it} = C_{it} \exp(-\mathbf{Z}'_{it}\boldsymbol{\eta}^C)$ and $\tilde{H}_{2it} = H_{2it} \exp(-\mathbf{Z}'_{it}\boldsymbol{\eta}^H)$.

Assuming an internal optimal, the first-order conditions of household problem (3.1) are

$$\begin{aligned} [C_{it}] : \quad & \tilde{U}'_i(\tilde{C}_{it}) \exp(-\mathbf{Z}'_{it}\boldsymbol{\eta}^C) = \lambda_{it} \\ [H_{2it}] : \quad & \tilde{V}'_i(\tilde{H}_{2it}) \exp(-\mathbf{Z}'_{it}\boldsymbol{\eta}^H) = \lambda_{it} W_{2it} \\ [A_{it+1}] : \quad & \beta(1+r)\mathbb{E}_t \lambda_{it+1} = \lambda_{it} \end{aligned}$$

where \tilde{U}'_i and \tilde{V}'_i denote first-order derivatives.

Applying logs to the first-order condition for consumption and taking a first difference in time yields $\Delta \ln \tilde{U}'_i(\tilde{C}_{it}) - \Delta \mathbf{Z}'_{it}\boldsymbol{\eta}^C = \Delta \ln \lambda_{it}$. A first-order

Taylor approximation of $\tilde{U}'_i(\tilde{C}_{it})$ around \tilde{C}_{it-1} yields $\Delta \ln \tilde{U}'_i(\tilde{C}_{it}) \approx \frac{1}{\phi_i} \Delta \ln \tilde{C}_{it}$ where $\phi_i = \tilde{U}'_i(\tilde{C}_{it-1})/\tilde{U}''_i(\tilde{C}_{it-1})\tilde{C}_{it-1}$ is the consumption substitution elasticity, *i.e.* the (Frisch) elasticity of consumption with respect to consumption's own price. Combining the two I get

$$\begin{aligned}\Delta \ln \tilde{C}_{it} - \phi_i \Delta \mathbf{Z}'_{it} \boldsymbol{\eta}^C &\approx \phi_i \Delta \ln \lambda_{it} \\ \Delta c_{it} &\approx \phi_i \Delta \ln \lambda_{it}\end{aligned}$$

where $\Delta c_{it} = \Delta \ln \tilde{C}_{it} - \phi_i \Delta \mathbf{Z}'_{it} \boldsymbol{\eta}^C$ (consumption net of observables).

The approximation to the first-order condition for hours of work follows a similar procedure.

The approximation to the first-order condition for assets (the Euler equation) is more involved as it pertains to future expectations. Let $\exp(\varrho) = 1/\beta(1+r)$ for an appropriate ϱ . I apply a second-order approximation to the exponential function $\exp(\cdot)$, evaluated at $\ln \lambda_{it+1}$, around $\ln \lambda_{it} + \varrho$ to get

$$\exp(\ln \lambda_{it+1}) \approx \exp(\ln \lambda_{it} + \varrho) \left(1 + (\Delta \ln \lambda_{it+1} - \varrho) + \frac{1}{2} (\Delta \ln \lambda_{it+1} - \varrho)^2 \right).$$

Taking expectations at time t and noting that $\mathbb{E}_t \lambda_{it+1} = \lambda_{it} \exp(\varrho)$ (the Euler equation) implies

$$\mathbb{E}_t \left(\Delta \ln \lambda_{it+1} - \varrho + \frac{1}{2} (\Delta \ln \lambda_{it+1} - \varrho)^2 \right) \approx 0$$

which in turn implies

$$\begin{aligned}\mathbb{E}_t \Delta \ln \lambda_{it+1} &\approx \varrho - \frac{1}{2} \mathbb{E}_t (\Delta \ln \lambda_{i,t+1} - \varrho)^2 \\ \Delta \ln \lambda_{it+1} &\approx \varrho - \frac{1}{2} \mathbb{E}_t (\Delta \ln \lambda_{it+1} - \varrho)^2 + \varepsilon_{it+1} \\ \Delta \ln \lambda_{it+1} &\approx \omega_{it+1} + \varepsilon_{it+1}.\end{aligned}$$

The first term $\omega_{it+1} = \varrho - \frac{1}{2} \mathbb{E}_t (\Delta \ln \lambda_{it+1} - \varrho)^2$ captures household precautionary motives; the second term ε_{it+1} is the expectation error at $t + 1$ and captures

idiosyncratic revisions to λ_{it} made by the household when income shocks hit.

C.2 Approximation to Intertemporal Budget Constraint

Let $F(\boldsymbol{\xi}) = \ln \sum_{s=0}^J \exp \xi_s$ for $\boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_J)'$. A first-order Taylor approximation around $\boldsymbol{\xi}^0$ is

$$F(\boldsymbol{\xi}) \approx \ln \sum_{s=0}^J \exp \xi_s^0 + \sum_{s=0}^J \frac{\exp \xi_s^0}{\sum_{s=0}^J \exp \xi_s^0} (\xi_s - \xi_s^0).$$

The intertemporal budget constraint of household i in problem (3.1) is

$$A_{it} + \sum_{s=0}^{T-t} \frac{Y_{1it+s} + W_{2it+s}H_{2it+s}}{(1+r)^s} = \sum_{s=0}^{T-t} \frac{C_{it+s}}{(1+r)^s}.$$

To ease the notation I will drop cross-sectional subscript i temporarily.

Left hand side: The logarithm of the left hand side is

$$F^{LH}(\boldsymbol{\xi}) = \ln \left[\exp(\ln A_t) + \sum_{s=0}^{T-t} \exp \left(\ln \frac{Y_{1t+s} + W_{2t+s}H_{2t+s}}{(1+r)^s} \right) \right].$$

Let

$$\xi_s = \begin{cases} \ln(Y_{1t+s} + W_{2t+s}H_{2t+s}) - s \ln(1+r) & \text{for } s = 0, \dots, T-t \\ \ln A_{t+s-(T-t+1)} & \text{for } s = T-t+1 \end{cases}$$

and

$$\xi_s^0 = \begin{cases} \mathbb{E}_{t-1} \ln(Y_{1t+s} + W_{2t+s}H_{2t+s}) - s \ln(1+r) & \text{for } s = 0, \dots, T-t \\ \mathbb{E}_{t-1} \ln A_{t+s-(T-t+1)} & \text{for } s = T-t+1. \end{cases}$$

Applying the first-order Taylor approximation to $F^{LH}(\boldsymbol{\xi})$ and taking ex-

pectations at t yields

$$\begin{aligned}\mathbb{E}_t F^{LH}(\boldsymbol{\xi}) &\approx F^{LH}(\boldsymbol{\xi}^0) + \frac{\exp(\mathbb{E}_{t-1} \ln A_t)}{\sum(\cdot)} (\mathbb{E}_t - \mathbb{E}_{t-1}) \ln A_t \\ &\quad + \sum_{s=0}^{T-t} \frac{\exp(\mathbb{E}_{t-1} \ln(Y_{1t+s} + W_{2t+s} H_{2t+s}) - s \ln(1+r))}{\sum(\cdot)} \\ &\quad \times (\mathbb{E}_t - \mathbb{E}_{t-1}) \ln(Y_{1t+s} + W_{2t+s} H_{2t+s}) \\ \mathbb{E}_t F^{LH}(\boldsymbol{\xi}) &\approx F^{LH}(\boldsymbol{\xi}^0) + \pi_t (\mathbb{E}_t - \mathbb{E}_{t-1}) \ln A_t \\ &\quad + (1 - \pi_t) \sum_{s=0}^{T-t} \delta_{t+s} (\mathbb{E}_t - \mathbb{E}_{t-1}) \ln(Y_{1t+s} + W_{2t+s} H_{2t+s}) \\ \mathbb{E}_t F^{LH}(\boldsymbol{\xi}) &\approx F^{LH}(\boldsymbol{\xi}^0) + \pi_t (\mathbb{E}_t - \mathbb{E}_{t-1}) \ln A_t + (1 - \pi_t) \\ &\quad \times \sum_{s=0}^{T-t} \delta_{t+s} (\mathbb{E}_t - \mathbb{E}_{t-1}) (\rho_{t+s-1} \Delta \ln Y_{1t+s} + (1 - \rho_{t+s-1}) \Delta \ln W_{2t+s} H_{2t+s})\end{aligned}$$

where $(\mathbb{E}_t - \mathbb{E}_{t-1}) \ln A_t = \mathbb{E}_t \ln A_t - \mathbb{E}_{t-1} \ln A_t$, and similarly for whenever $\mathbb{E}_t - \mathbb{E}_{t-1}$ appears multiplying any variable, and

$$\begin{aligned}\sum(\cdot) &= \exp(\mathbb{E}_{t-1} \ln A_t) + \sum_{k=0}^{T-t} \exp\left(\mathbb{E}_{t-1} \ln \frac{Y_{1t+k} + W_{2t+k} H_{2t+k}}{(1+r)^k}\right) \\ \pi_t &= \frac{\exp(\mathbb{E}_{t-1} \ln A_t)}{\sum(\cdot)} \\ \delta_{t+s} &= \frac{\exp(\mathbb{E}_{t-1} \ln(Y_{1t+s} + W_{2t+s} H_{2t+s}) - s \ln(1+r))}{\sum_{k=0}^{T-t} \exp(\mathbb{E}_{t-1} \ln(Y_{1t+k} + W_{2t+k} H_{2t+k}) - k \ln(1+r))} \\ \rho_{t+s} &= \frac{Y_{1t+s}}{Y_{1t+s} + W_{2t+s} H_{2t+s}}.\end{aligned}$$

The notation is as follows: π_t is the ‘partial insurance’ parameter (term due to Blundell et al., 2008) and is the share of initial financial assets in the household’s expected lifetime financial and human wealth; δ_t is the share of family current total earnings in their expected lifetime earnings (lifetime human wealth); ρ_t is the contemporaneous ratio of male earnings over family total earnings at t .

A first difference between $\mathbb{E}_t F^{LH}(\boldsymbol{\xi})$ and $\mathbb{E}_{t-1} F^{LH}(\boldsymbol{\xi})$ eliminates $F^{LH}(\boldsymbol{\xi}^0)$

and $\pi_t(\mathbb{E}_t - \mathbb{E}_{t-1}) \ln A_t$ rendering $\mathbb{E}_t F^{LH}(\boldsymbol{\xi}) - \mathbb{E}_{t-1} F^{LH}(\boldsymbol{\xi}) \approx$

$$(1 - \pi_t) \sum_{s=0}^{T-t} \delta_{t+s} (\mathbb{E}_t - \mathbb{E}_{t-1}) (\rho_{t+s-1} \Delta \ln Y_{1t+s} + (1 - \rho_{t+s-1}) \Delta \ln W_{2t+s} H_{2t+s})$$

Right hand side: The logarithm of the right hand side is

$$F^{RH}(\boldsymbol{\xi}) = \ln \sum_{s=0}^{T-t} \exp \left(\ln \frac{C_{t+s}}{(1+r)^s} \right)$$

for an appropriately defined $\boldsymbol{\xi}$. Following the same steps as for the left hand side I get

$$\mathbb{E}_t F^{RH}(\boldsymbol{\xi}) - \mathbb{E}_{t-1} F^{RH}(\boldsymbol{\xi}) \approx \sum_{s=0}^{T-t} \theta_{t+s} (\mathbb{E}_t - \mathbb{E}_{t-1}) \Delta \ln C_{t+s}$$

where

$$\theta_{t+s} = \frac{\exp(\mathbb{E}_{t-1} \ln C_{t+s} - s \ln(1+r))}{\sum_{k=0}^{T-t} \exp(\mathbb{E}_{t-1} \ln C_{t+k} - k \ln(1+r))}$$

is the share of family current consumption in the household expected lifetime consumption.

Two sides together: I bring the two sides together following Blundell et al. (2013, p. 34) who point out that “the realized budget must balance” and, therefore, the objects on the two sides of the log-linearized budget constraint “have the same distribution”:

$$\begin{aligned} (1 - \pi_t) \sum_{s=0}^{T-t} \delta_{t+s} (\mathbb{E}_t - \mathbb{E}_{t-1}) (\rho_{t+s-1} \Delta \ln Y_{1t+s} + (1 - \rho_{t+s-1}) \Delta \ln W_{2t+s} H_{2t+s}) \\ \approx \sum_{s=0}^{T-t} \theta_{t+s} (\mathbb{E}_t - \mathbb{E}_{t-1}) \Delta \ln C_{t+s} \end{aligned} \quad (\text{C.1})$$

The approximation up to this point does not relate to the first-order conditions and the specific economic problem under study. I retrieve the first-order

conditions (3.3) and the income processes (3.2) of problem (3.1) and I substitute them into equation (C.1). I bring back cross-sectional i to get

$$(\phi_i - (1 - \pi_{it})(1 - s_{it})\psi_i) \varepsilon_{it} \approx (1 - \pi_{it})(s_{it}v_{1it} + (1 - s_{it})(1 + \psi_i)v_{2it}) \quad (C.2)$$

where s_{it} is approximately equal to the average expected ratio of male earnings over family total earnings between time t and the end of the horizon, i.e. $s_{it} \approx \mathbb{E}_t \overline{\rho_{it}}$ and the average is taken over $t, t+1, \dots, T$ for a given household i .

To obtain expression (C.2), I assume: (i) ω_{it+s} is in the information set of households at times t and $t-1$ for all $s > 0$, thus $(\mathbb{E}_t - \mathbb{E}_{t-1})\omega_{it+s} = 0, \forall s > 0$, and (ii) $\delta_{t+s} \approx 0$, meaning that household earnings in any single time period are negligible when compared to lifetime earnings. The last assumption implies that transitory shocks to earnings or wages do not shift the intertemporal budget constraint and, therefore, do not induce wealth effects.

If one further assumes that the households have negligible accumulated assets compared to their lifetime financial and human wealth, as is likely the case for young households, then $\pi_{it} \approx 0$ and expression (C.2) can be written

$$\varepsilon_{it} \approx \frac{1}{\phi_i - (1 - s_{it})\psi_i} (s_{it}v_{1it} + (1 - s_{it})(1 + \psi_i)v_{2it})$$

which is expression (3.4) in the main text.

C.3 Derivation of $Cov(\Delta h_{2it}, \Delta h_{2it+1})$

From (3.5) I can write

$$\Delta h_{2it} \approx \psi_i \zeta_{it} s_{it} v_{1it} + (\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i) v_{2it} + \psi_i \Delta u_{2it} + \psi_i \omega_{it}.$$

The covariance $Cov(\Delta h_{2it}, \Delta h_{2it+1})$ is equal to

$$\begin{aligned} \text{line 1} & \quad Cov(\psi_i \zeta_{it} s_{it} v_{1it}, \psi_i \zeta_{it+1} s_{it+1} v_{1it+1}) \\ \text{line 2} & \quad + Cov(\psi_i \zeta_{it} s_{it} v_{1it}, (\psi_i \zeta_{it+1} (1 - s_{it+1})(1 + \psi_i) + \psi_i) v_{2it+1}) \end{aligned}$$

$$\begin{aligned}
\text{line 3} &+ Cov(\psi_i \zeta_{it} s_{it} v_{1it}, \psi_i \Delta u_{2it+1}) \\
\text{line 4} &+ Cov(\psi_i \zeta_{it} s_{it} v_{1it}, \psi_i \omega_{it+1}) \\
\text{line 5} &+ Cov(\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i) v_{2it}, \psi_i \zeta_{it+1} s_{it+1} v_{1it+1}) \\
\text{line 6} &+ Cov(\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i) v_{2it}, (\psi_i \zeta_{it+1} (1 - s_{it+1})(1 + \psi_i) + \psi_i) v_{2it+1}) \\
\text{line 7} &+ Cov(\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i) v_{2it}, \psi_i \Delta u_{2it+1}) \\
\text{line 8} &+ Cov(\psi_i \zeta_{it} (1 - s_{it})(1 + \psi_i) + \psi_i) v_{2it}, \psi_i \omega_{it+1}) \\
\text{line 9} &+ Cov(\psi_i \Delta u_{2it}, \psi_i \zeta_{it+1} s_{it+1} v_{1it+1}) \\
\text{line 10} &+ Cov(\psi_i \Delta u_{2it}, (\psi_i \zeta_{it+1} (1 - s_{it+1})(1 + \psi_i) + \psi_i) v_{2it+1}) \\
\text{line 11} &+ Cov(\psi_i \Delta u_{2it}, \psi_i \Delta u_{2it+1}) \\
\text{line 12} &+ Cov(\psi_i \Delta u_{2it}, \psi_i \omega_{it+1}) \\
\text{line 13} &+ Cov(\psi_i \omega_{it}, \psi_i \zeta_{it+1} s_{it+1} v_{1it+1}) \\
\text{line 14} &+ Cov(\psi_i \omega_{it}, (\psi_i \zeta_{it+1} (1 - s_{it+1})(1 + \psi_i) + \psi_i) v_{2it+1}) \\
\text{line 15} &+ Cov(\psi_i \omega_{it}, \psi_i \Delta u_{2it+1}) \\
\text{line 16} &+ Cov(\psi_i \omega_{it}, \psi_i \omega_{it+1}).
\end{aligned}$$

Lines 1 and 6 are 0 because permanent shocks are serially uncorrelated; lines 2 and 5 are 0 because permanent shocks between spouses and over time do not co-vary, lines 3, 7, 9 and 10 are 0 because permanent and transitory shocks do not co-vary; lines 4, 8, and 12-15 are 0 because income shocks are mean independent of ω ; finally line 16 is 0 because ω does not exhibit serial correlation. To characterize these covariances I use results from Goodman (1960) and Bohrnstedt and Goldberger (1969).

The only remaining non-0 line is 11 which is equal to

$$Cov(\psi_i u_{2it}, -\psi_i u_{2it})$$

and is the same expression that appears in the main text (Section 3.3.3.1).

Bibliography

- Abowd, J. M. and D. Card (1989). On the covariance structure of earnings and hours changes. *Econometrica* 57(2), 411–45.
- Adda, J. and R. W. Cooper (2003). *Dynamic Economics: Quantitative Methods and Applications*. The MIT Press.
- 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.
- Apps, P. F. and R. Rees (1988). Taxation and the household. *Journal of Public Economics* 35(3), 355–369.
- Attanasio, O., G. Berloffa, R. Blundell, and I. Preston (2002). From earnings inequality to consumption inequality. *The Economic Journal* 112(478), C52–C59.
- Attanasio, O., H. Low, and V. Sanchez-Marcos (2008). Explaining changes in female labor supply in a life-cycle model. *The American Economic Review* 98(4), 1517–1552.
- Attanasio, O. and N. Pavoni (2011). Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption. *Econometrica* 79(4), 1027–1068.
- Barmby, T. and N. Smith (2001). Household labour supply in Britain and Denmark: some interpretations using a model of Pareto optimal behaviour. *Applied Economics* 33(9), 1109–1116.

- Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal* 75(299), 493–517.
- Biddle, J. E. and D. S. Hamermesh (1990). Sleep and the allocation of time. *Journal of Political Economy* 98(5), 922–943.
- Blau, D. M. (1992). The child care labor market. *Journal of Human Resources* 27(1), 9–39.
- Blau, F. D. and L. M. Kahn (1997). Swimming upstream: Trends in the gender wage differential in 1980s. *Journal of Labor Economics* 15(1), 1–42.
- Blau, F. D. and L. M. Kahn (2006). The U.S. gender pay gap in the 1990s: Slowing convergence. *Industrial and Labor Relations Review* 60(1), 45–66.
- Blundell, R., P.-A. Chiappori, T. Magnac, and C. Meghir (2007). Collective labour supply: Heterogeneity and non-participation. *The Review of Economic Studies* 74(2), 417–445.
- Blundell, R., P.-A. Chiappori, and C. Meghir (2005). Collective labor supply with children. *Journal of Political Economy* 113(6), 1277–1306.
- Blundell, R., M. C. Dias, C. Meghir, and J. M. Shaw (2013). Female labour supply, human capital and welfare reform. NBER Working Papers 19007, National Bureau of Economic Research, Inc.
- Blundell, R. and B. Etheridge (2010). Consumption, income and earnings inequality in Britain. *Review of Economic Dynamics* 13(1), 76–102.
- Blundell, R., M. Graber, and M. Mogstad (2015). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics* 127(C), 58–73.
- Blundell, R., J. Horowitz, and M. Parey (2013). Nonparametric estimation of a heterogeneous demand function under the Slutsky inequality restriction.

- CeMMAP working papers CWP54/13, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.
- Blundell, R., D. Kristensen, and R. Matzkin (2014). Bounding quantile demand functions using revealed preference inequalities. *Journal of Econometrics* 179(2), 112–127.
- Blundell, R., H. Low, and I. Preston (2013). Decomposing changes in income risk using consumption data. *Quantitative Economics* 4(1), 1–37.
- Blundell, R., T. MaCurdy, and C. Meghir (2007). Labor supply models: Unobserved heterogeneity, nonparticipation and dynamics. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6 of *Handbook of Econometrics*, pp. 4667–4775. Elsevier.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *The American Economic Review* 98(5), 1887–1921.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2015). Children, time allocation and consumption insurance.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Blundell, R. and I. Preston (1998). Consumption inequality and income uncertainty. *The Quarterly Journal of Economics* 113(2), 603–640.
- Blundell, R. and T. M. Stoker (2007). Models of aggregate economic relationships that account for heterogeneity. In J. J. Heckman and E. E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6 of *Handbook of Econometrics*, pp. 4609–4666. Elsevier.
- Bohrnstedt, G. and A. Goldberger (1969). On the exact covariance of products of random variables. *Journal of the American Statistical Association* 64(328), 1439–1442.

- Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers (1994). Evidence on the validity of cross-sectional and longitudinal labor market data. *Journal of Labor Economics* 12(3), pp. 345–368.
- Bourguignon, F., M. Browning, and P.-A. Chiappori (2009). Efficient intra-household allocations and distribution factors: implications and identification. *The Review of Economic Studies* 76(2), 503–528.
- Bourguignon, F., M. Browning, P.-A. Chiappori, and V. Lechene (1993). Intra-household allocation of consumption: A model and some evidence from french data. *Annales d’Economie et de Statistique* (29), 137–156.
- Bronson, M. A. (2014). Degrees are forever: Marriage, educational investment, and lifecycle labor decisions of men and women.
- Browning, M., F. Bourguignon, P.-A. Chiappori, and V. Lechene (1994). Income and outcomes—a structural model of intrahousehold allocation. *Journal of Political Economy* 102(6), 1067–1096.
- 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 A. Lewbel (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies* 80(4), 1267–1303.
- Browning, M., P.-A. Chiappori, and Y. Weiss (2014). *Economics of the Family*. Cambridge Books. Cambridge University Press.
- Browning, M., M. Ejrnæs, and J. Alvarez (2010). Modelling income processes with lots of heterogeneity. *Review of Economic Studies* 77(4), 1353–1381.
- Browning, M., L. P. Hansen, and J. J. Heckman (1999). Micro data and general equilibrium models. In J. B. Taylor and M. Woodford (Eds.), *Handbook of*

- Macroeconomics* (1 ed.), Volume 1 of *Handbook of Macroeconomics*, pp. 543–633. Elsevier.
- Browning, M. and C. Meghir (1991). The effects of male and female labor supply on commodity demands. *Econometrica* 59(4), 925–951.
- Campbell, J. Y. (1993). Intertemporal asset pricing without consumption data. *The American Economic Review* 83(3), pp. 487–512.
- Cherchye, L., S. Cosaert, T. Demuynck, and B. De Rock (2015). Noncooperative household consumption with caring.
- Cherchye, L., B. De Rock, and F. Vermeulen (2012). Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information. *The American Economic Review* 102(7), 3377–3405.
- Cherchye, L., B. D. Rock, A. Lewbel, and F. Vermeulen (2015). Sharing rule identification for general collective consumption models. *Econometrica* 83(5), 2001–2041.
- Cherchye, L., B. D. Rock, and F. Vermeulen (2009). Opening the black box of intrahousehold decision making: Theory and nonparametric empirical tests of general collective consumption models. *Journal of Political Economy* 117(6), 1074–1104.
- Cherchye, L., B. D. Rock, and F. Vermeulen (2011). The revealed preference approach to collective consumption behaviour: Testing and sharing rule recovery. *Review of Economic Studies* 78(1), 176–198.
- Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica* 56(1), 63–90.
- Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy* 100(3), 437–467.

- Chiappori, P.-A. (1997). Introducing household production in collective models of labor supply. *Journal of Political Economy* 105(1), 191–209.
- Chiappori, P.-A. (2011). Collective labor supply with many consumption goods. *Review of Economics of the Household* 9(2), 207–220.
- Chiappori, P.-A., M. C. Dias, and C. Meghir (2015). The marriage market, labor supply and education choice. NBER Working Papers 21004, National Bureau of Economic Research, Inc.
- Chiappori, P.-A. and I. Ekeland (2006). The micro economics of group behavior: general characterization. *Journal of Economic Theory* 130(1), 1–26.
- Chiappori, P.-A. and I. Ekeland (2009). The microeconomics of efficient group behavior: Identification. *Econometrica* 77(3), 763–799.
- Chiappori, P.-A., B. Fortin, and G. Lacroix (2002). Marriage market, divorce legislation, and household labor supply. *Journal of political Economy* 110(1), 37–72.
- Chiappori, P.-A. and M. Mazzocco (2014). Static and intertemporal household decisions.
- Chiappori, P.-A. and C. Meghir (2014). Intrahousehold inequality. NBER Working Papers 20191, National Bureau of Economic Research, Inc.
- Cosaert, S. and T. Demuynck (2014). Nonparametric welfare and demand analysis with unobserved individual heterogeneity. Research Memorandum 010, Maastricht University, Graduate School of Business and Economics (GSBE).
- Dauphin, A., A.-R. El Lahga, B. Fortin, and G. Lacroix (2011). Are children decision-makers within the household? *The Economic Journal* 121(553), 871–903.

- Del Boca, D., C. Flinn, and M. Wiswall (2014). Household Choices and Child Development. *Review of Economic Studies* 81(1), 137–185.
- Donni, O. (2003). Collective household labor supply: nonparticipation and income taxation. *Journal of Public Economics* 87(5), 1179–1198.
- Dunbar, G. R., A. Lewbel, and K. Pendakur (2013). Children’s resources in collective households: Identification, estimation, and an application to child poverty in Malawi. *The American Economic Review* 103(1), 438–471.
- Eckstein, Z. and O. Lifshitz (2011). Dynamic female labor supply. *Econometrica* 79(6), 1675–1726.
- Eckstein, Z. and K. I. Wolpin (1989). Dynamic labour force participation of married women and endogenous work experience. *The Review of Economic Studies* 56(3), pp. 375–390.
- Fernández, R. and J. C. Wong (2014). Divorce risk, wages and working wives: A quantitative life-cycle analysis of female labour force participation. *The Economic Journal* 124(576), 319–358.
- Fong, Y.-F. and J. Zhang (2001). The identification of unobservable independent and spousal leisure. *Journal of Political Economy* 109(1), 191–202.
- 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.
- Francesconi, M. (2002). A joint dynamic model of fertility and work of married women. *Journal of Labor Economics* 20(2), 336–380.
- French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies* 72(2), 395–427.
- Gemici, A. (2011). Family migration and labor market outcomes.
- Goodman, L. (1960). On the exact variance of products. *Journal of the American Statistical Association* 55(292), 708–713.

- Gottschalk, P. and R. Moffitt (2009). The rising instability of U.S. earnings. *Journal of Economic Perspectives* 23(4), 3–24.
- Gourieroux, C., A. Monfort, and E. Renault (1993). Indirect inference. *Journal of Applied Econometrics* 8, S85–S85.
- Greenwood, J., A. Seshadri, and M. Yorukoglu (2005). Engines of liberation. *The Review of Economic Studies* 72(1), 109–133.
- Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent? *American Economic Review* 97(3), 687–712.
- Heathcote, J., F. Perri, and G. 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 (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review* 104(7), 2075–2126.
- Horowitz, J. L. (2001). The bootstrap. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 5 of *Handbook of Econometrics*, pp. 3159 – 3228. Elsevier.
- Hyslop, D. R. (2001). Rising U.S. earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *American Economic Review* 91(4), 755–777.
- Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics* 2(1), 479–506.
- Johnson, S. G. The NLOpt nonlinear-optimization package.
- Kaelo, P. and M. Ali (2006). Some variants of the controlled random search algorithm for global optimization. *Journal of Optimization Theory and Applications* 130(2), 253–264.

- Kaplan, G. and G. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2(4), 53–87.
- Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature* 49(4), 961–1075.
- Keane, M. P. and K. I. Wolpin (2010). The role of labor and marriage markets, preference heterogeneity, and the welfare system in the life cycle decisions of black, hispanic, and white women. *International Economic Review* 51(3), 851–892.
- Kiefer, N. M. (1984). Microeconometric evidence on the neoclassical model of demand. *Journal of Econometrics* 25(3), 285–302.
- Knowles, J. A. (2013). Why are married men working so much? An aggregate analysis of intra-household bargaining and labour supply. *The Review of Economic Studies* 80(3), 1055–1085.
- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *Review of Economic Studies* 63(4), 595–609.
- Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? evidence and theory. *Review of Economic Studies* 73(1), 163–193.
- Lechene, V. and I. Preston (2011). Noncooperative household demand. *Journal of Economic Theory* 146(2), 504–527.
- Lewbel, A. (2001). Demand systems with and without errors. *American Economic Review* 91(3), 611–618.
- Lewbel, A. and K. Pendakur (2008). Estimation of collective household models with engel curves. *Journal of Econometrics* 147(2), 350–358.
- Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies* 69(1), 209–44.

- Lise, J. and S. Seitz (2011). Consumption inequality and intra-household allocations. *The Review of Economic Studies* 78(1), 328–355.
- Lise, J. and K. Yamada (2014). Household sharing and commitment: Evidence from panel data on individual expenditures and time use. In *2014 Meeting Papers*, Number 152. Society for Economic Dynamics.
- Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics* 3(1), 11–37.
- 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.
- Matzkin, R. L. (2003). Nonparametric estimation of nonadditive random functions. *Econometrica* 71(5), 1339–1375.
- Matzkin, R. L. (2008). Identification in nonparametric simultaneous equations models. *Econometrica* 76(5), 945–978.
- 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 (2014). Labor Supply, Wealth Dynamics and Marriage Decisions. 2014 meeting papers, Society for Economic Dynamics.
- 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.

- Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica* 72(1), 1–32.
- Meghir, C. and L. Pistaferri (2011). Earnings, consumption and life cycle choices. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4 of *Handbook of Labor Economics*, pp. 773–854. Elsevier.
- O’Neill, G. and M. O’Connell (2001). State estimates of child care establishments: 1977–1997. Population Division Working Papers 55, U.S. Bureau of the Census.
- Phlips, L. (1974). *Applied Consumption Analysis*. Advanced textbooks in economics. North-Holland.
- Rowan, T. H. (1990). *Functional stability analysis of numerical algorithms*. Ph. D. thesis, Department of Computer Sciences, University of Texas at Austin.
- van der Klaauw, W. (1996). Female labour supply and marital status decisions: A life-cycle model. *Review of Economic Studies* 63(2), 199–235.
- Vermeulen, F. (2006). A collective model for female labour supply with non-participation and taxation. *Journal of Population Economics* 19(1), 99–118.
- Vermeulen, F., O. Bargain, M. Beblo, D. Beninger, R. Blundell, R. Carrasco, M.-C. Chiuri, F. Laisney, V. Lechene, N. Moreau, M. Myck, and J. Ruiz-Castillo (2006). Collective models of labor supply with nonconvex budget sets and nonparticipation: A calibration approach. *Review of Economics of the Household* 4(2), 113–127.
- Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review* 105(8), 2295–2332.
- Whitebook, M., D. Phillips, and C. Howes (1993). National child care staffing study revisited: Four years in the life of center-based child care.