Wages and Family Time Allocation*

Alexandros Theloudis†
LISER and UCL

February 1, 2018

Abstract

This paper examines changes in married people's allocation of time since 1980, a period in which female labor supply increased substantially, men's share of household work rose, and the gender wage gap narrowed down. I develop a life-cycle collective household model for market and non-market work, consumption and asset accumulation, which also features lack of commitment to lifetime marriage. Wages in the model shift intra-family bargaining power and induce bargaining effects on outcomes in addition to standard income and substitution effects. I estimate gender-specific preferences and how intra-family bargaining power changes with a narrowing gender gap using data from the PSID. The results suggest that a narrowing gender gap improved women's bargaining power in the family resulting in a shift of household work to their husbands. It also contributed to the increase in female labor market participation. If the gender gap is counterfactually eliminated, the proportion of women in full-time work rises throughout the lifecycle to match approximately that of men. The increase comes from women who cut down household chores and enter the labor market when they previously did not participate.

Keywords: Life-cycle collective model, home production, lack of commitment, gender wage gap, bargaining effects, equal pay, simulated method of moments, PSID

JEL classification: D12, D13, D91, J22

*This paper is based on the first chapter of my dissertation. I am deeply indebted to Richard Blundell and Ian Preston for their invaluable advice, and to Eric French, Valérie Lechene and Luigi Pistaferri for their support. I also thank Arun Advani, Monica Costa Dias, Mariacristina De Nardi, Tim Lee, Peter Levell, Jeremy Lise, Costas Meghir, Aureo de Paula, Suphanit Piyapromdee, Uta Schönberg, Jonathan Shaw, Michela Tincani, and participants at various seminars and conferences for comments. I am grateful for financial support by a graduate studentship from the ESRC and a junior fellowship from the Royal Economic Society.

†Alexandros Theloudis: Luxembourg Institute for Socio-Economic Research and Department of Economics, University College London. Email: a.theloudis@gmail.com.

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? And what for 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 (two spouses) choose jointly how to allocate their time across market work, 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, induce shifts in the bargaining power spouses have 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 the PSID. I exploit cross-sectional variation in wages and family composition as well as the sharp decline in the gender wage gap 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. I use the model to investigate the likely implications the elimination of the gender wage gap has for family time allocation. In such counterfactual environment, the proportion of women in full-time market work increases up to 18 percentage points even during the childbearing years (from a base rate of 57%) and women enter the labor market when they previously abstain. The allocation of time into home production becomes more equal between spouses (primarily with women reducing their much higher hours) and the spouses' total time into chores decreases by as much as 7 hours per week (from an average base number of 34 hours/week).

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 80's and 90's 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 chores in the household whereas men's household work remained flat. The paper revolves currently around a single cohort whose life-cycle spans the period when the biggest changes in the gender wage gap and family time allocation occurred, namely the 3 decades since 1980. The paper, therefore, focuses currently on how the narrowing gender gap affected this cohort's time allocation *over the life-cycle* and serves as a *first* only step towards

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

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

[3] This figure is based on raw data from the PSID described in section 2.

understanding how wages affect family time allocation *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 are made between spouses. The precise channels through which wages likely affect family time allocation are the following: First, an increase in one's hourly wage renders market work 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 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 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, likely affects 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 these channels. Two spouses[5] have 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 expenditure') 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 to market work, household work, and leisure. An hour of market work is compensated by a gender-specific stochastic wage which individuals take as exogenous; earnings are used to fund public expenditure 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

[4]Another potential effect of wages and the gender wage gap is on the selection of individuals into marriage and, in general, on marital patterns. This paper abstracts from this feature taking marriage as given. 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. Their paper, however, 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.

[5]I use the terms 'spouses' to refer to two decision making individuals in the model; the model applies equally to traditional nuclear families as well as more modern forms of cohabitation.

[6]The model in this paper 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 (also Manser and Brown, 1980; McElroy and Horney, 1981).

(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, thus associate it with 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.

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. To help distinguish between the two channels, I fix, essentially normalize, 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 the spouses' lifetime earnings in the hypothetical scenario of divorce. I estimate their hypothetical earnings upon divorce using reduced-form information on divorcees in the PSID (*divorcees* because I treat divorce as the relevant outside option).

I estimate the model by the method of simulated moments using data from the PSID after 1980. Currently focusing on one cohort whose lifetime spans the period 1980-2009, 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 if the two supply the same amount of household hours.

The data reveal that both the monetary reward of market work and changes in bargaining power are important when the gender wage gap narrows down. A 10% closing of the gap in favor of women (relative to the 1980 rate where the 'median' man earns an hourly wage 1.7-1.8 higher than the 'median' woman) 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. 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's household work too (an income effect). In reality, however, men keep their household work unchanged as their weakened bargaining position counterbalances the income effect. As women become more powerful thanks to a narrower gender wage gap, they shift household work to their husbands. Finally, 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.

These findings are suggestive of the likely implications gender wage equality has for family time allocation. If women are paid on average their husbands' wage, female market participation increases strongly throughout the life-cycle. The most striking effects occur in the childbearing years when the rate of female full-time market work rises to approximately 75% compared to 57% in the data (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 do not participate. Gender wage equality renders the allocation of spousal time into home production more equal between spouses but it also decreases the total household time into chores by as much as 7 hours per week in 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 paper 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 outcomes such as consumption or retirement.

The papers in the first strand of literature to which this article is mostly related 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 serves as evidence against full 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 allocation 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 bargaining power depends on a marriage market equilibrium; he abstracts from life-cycle features such as savings or fertility. He finds intra-family bargaining to be critical for explaining trends in gender-specific labor supply but he does not separate changes occurring *across* cohorts from changes occurring *within* co-

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

horts. The present paper, 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 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 the gender wage gap to enter intra-family bargaining power.

Fernández and Wong (2014) use a life-cycle collective model to study the increase in female labor supply in the second half of the 20$^{\text{th}}$ 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, 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 (2011) also study women's employment rates, paying particular attention to the differential patterns that married and singles have experienced. Blundell et al. (2016) 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, my paper reserves 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

[8] In 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 (2017). Lise and Seitz (2011) and Chiappori and Meghir (2014) argue why the intra-household allocation of resources should not be ignored.

[9] Important 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.

household structure and which are discussed in section 3.2.

In relation to the literature, my paper 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 paper is arranged as follows. Section 2 describes the empirical facts behind this research. Section 3 develops the model of household decision making. Section 4 discusses technical aspects of the model while section 5 discusses identification and estimation. Section 6 presents the results. Section 7 discusses the implications of the model for behavior and section 8 describes the policy experiment. Section 9 concludes.

2 An Empirical Overview

Section 2.1 overviews the data used in this paper, section 2.2 establishes the family time allocation facts this study aims to explain, and section 2.3 discusses the evolution of the gender wage gap over time.

2.1 Data

This paper uses data from the Panel Study of Income Dynamics (PSID).[10] 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 PSID 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 the 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 PSID sample ('Survey Research Center'). I impose this age restriction because the model in this paper does not deal with early-life (education) or late-life (retirement) decisions although it does model retirement in a stylized way. 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

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

I describe it in more detail below. I postpone a discussion of the latter sample until section 5.3.

The major sample consists of married opposite-sex couples.[11] The analysis in this paper 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 (the years where I draw data from). 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 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 values[12] 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 age-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 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 while appendix A provides further details.

Table 1: Major sample descriptive statistics: married couples

	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 12^{th} 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. All descriptive statistics are calculated across all household-year observations. *Source:* PSID 1980-2009.

I concentrate on continuously married (stable) couples because I do not solve explicitly for the problem of divorce (see sections 3 and 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 6.

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

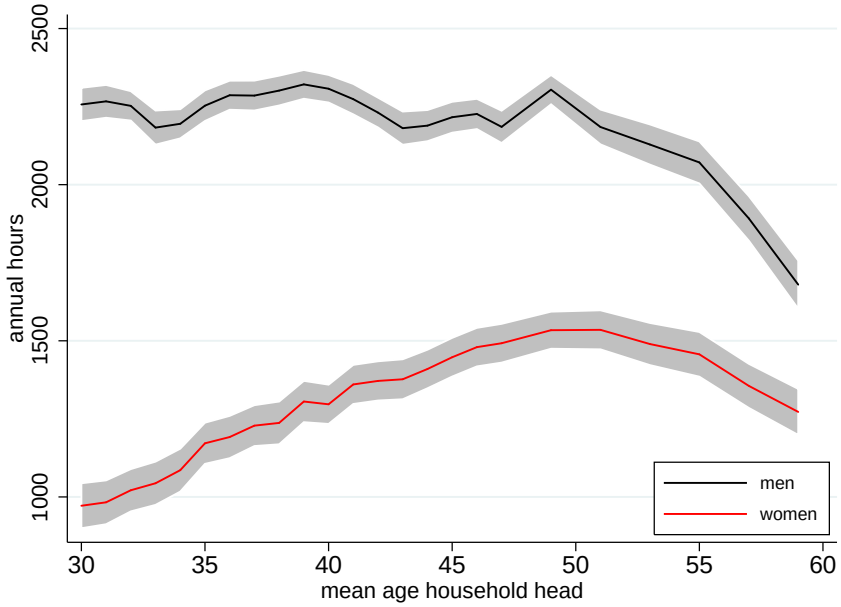
[12] I express all monetary amounts in 2010 dollars. To deflate I use the All-Urban-Consumers CPI available by the BLS at www.bls.gov/cpi.

2.2 Time Allocation Facts

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

Figure 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 1,000 annual hours 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.

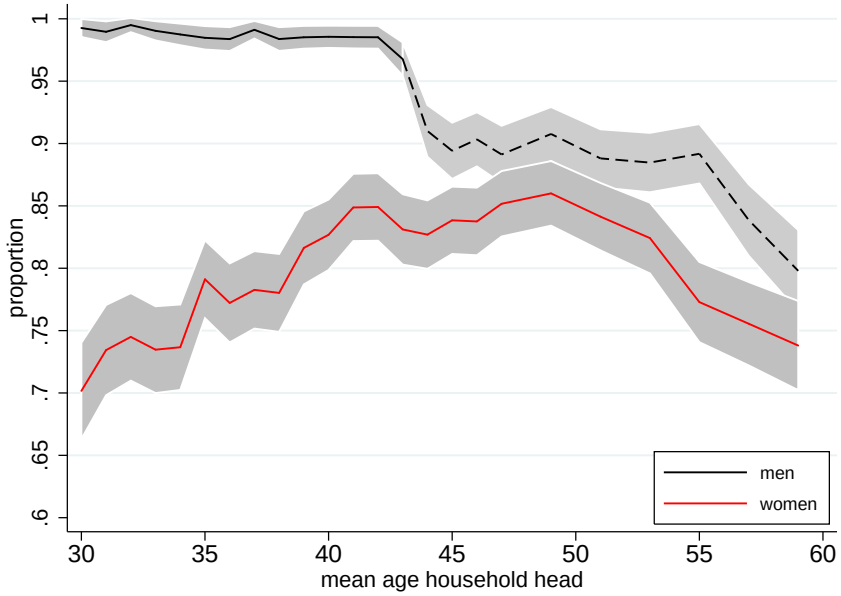
Figure 1: Average annual hours worked in the market



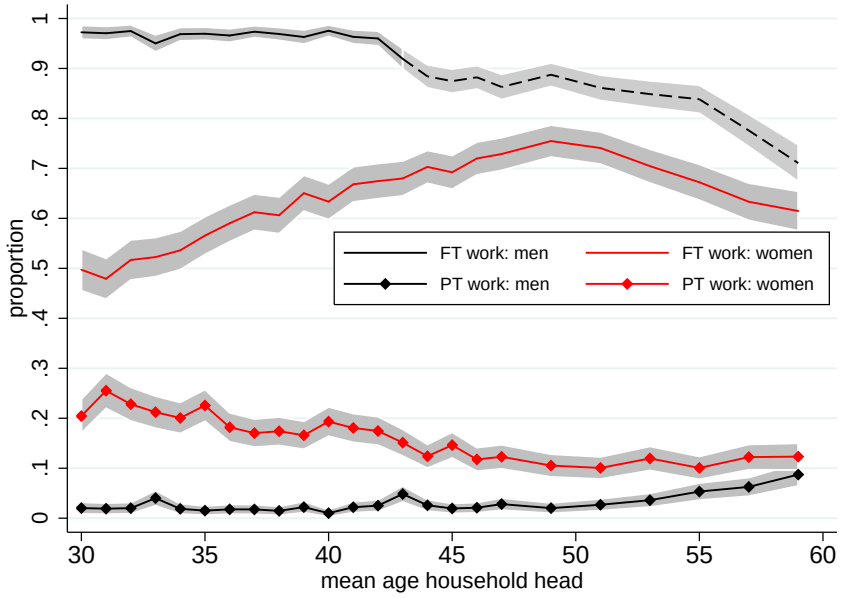
Notes: This figure plots average annual hours of market work for workers and non-workers. A 95% confidence interval appears in gray shade. *Source:* PSID 1980-2009.

To understand these trends better, figure 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. 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 it starts declining in the last few years.

Figure 2: Employment trends: market participation, full-time and part-time 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:* PSID 1980-2009.

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 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. 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 2b delves deeper into the employment trends and plots the proportions of people working full- or part-time in the labor market (among all 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 is a combination of a strong increase in the extensive margin of labor supply (figure 2a) and a smaller increase in the intensive margin (figure 2b).

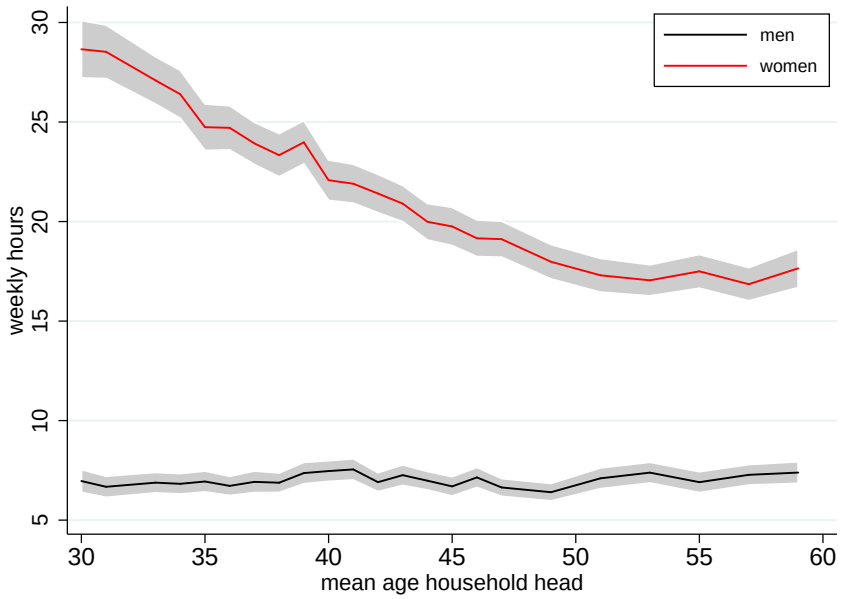
Turning to household work (time devoted to household chores), figure 3 plots *weekly* hours of household work for married men and women including those who report 0 such hours. Household work refers to any work in and around the household, such as cooking or cleaning, but excludes time spent with children. 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 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, the big drop in women's household work in figure 3 is almost exclusively due to a decrease in its intensive rather than extensive margin.

Finally, 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 5 redraws the initial graphs for market and household hours splitting the sample by 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

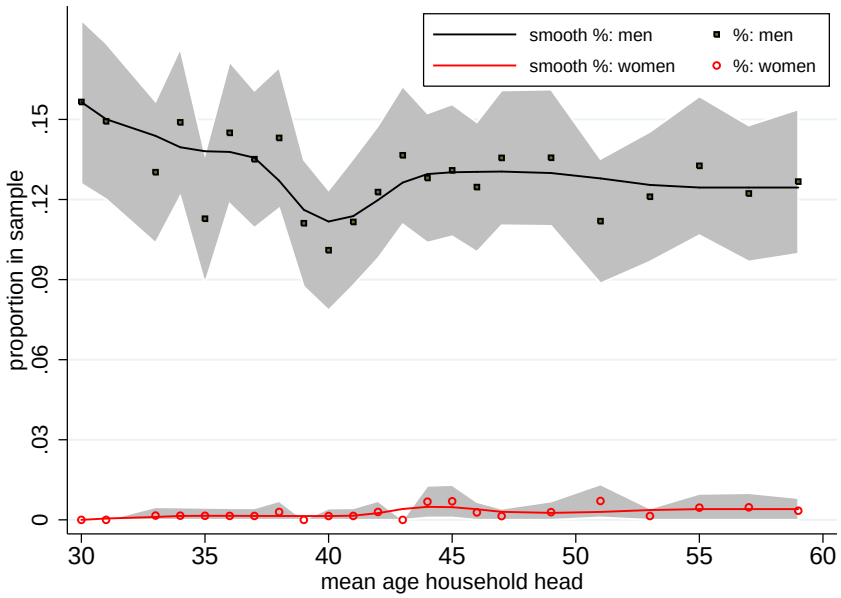
to men (albeit at different magnitudes). These facts are true for both work in the market and work in the household.

Figure 3: Average weekly hours worked in the household



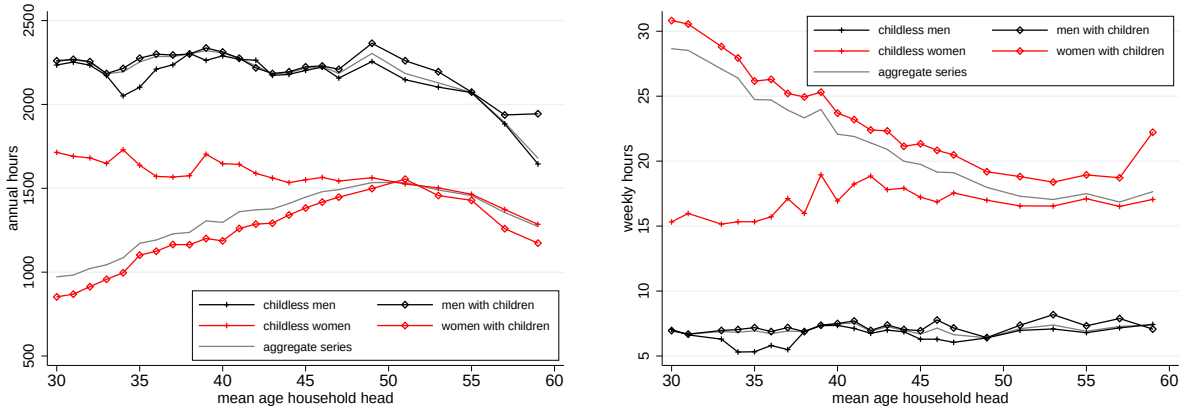
Notes: This figure plots average weekly hours of household work for household workers and non-workers. A 95% confidence interval appears in gray shade. *Source:* PSID 1980-2009.

Figure 4: Non-participation in home production



Notes: This figure plots the proportion of spouses who report zero weekly hours of household work. A 95% confidence interval around the original (non-smoothed) proportions appears in gray shade. *Source:* PSID 1980-2009.

Figure 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:* PSID 1980-2009.

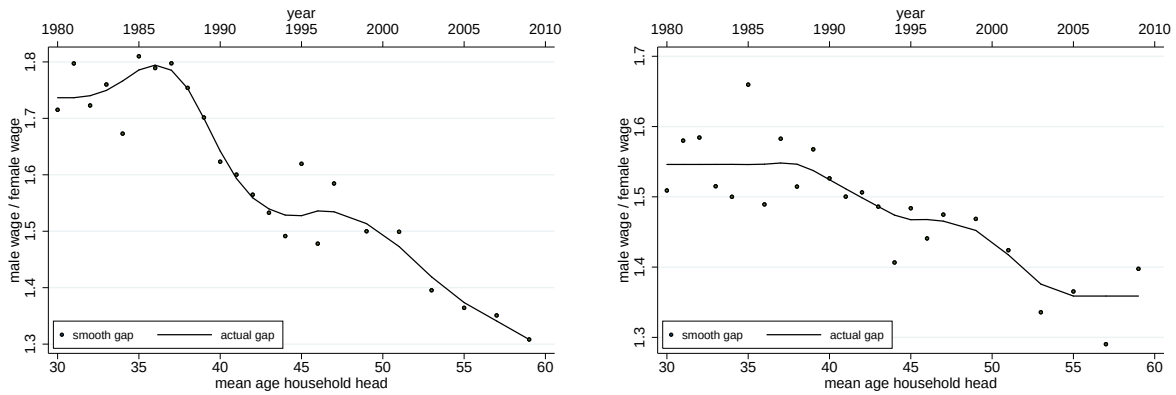
2.3 The Gender Wage Gap

As the paper focuses currently on one cohort, the gender wag gap statistics I report below refer to this cohort only. I study the evolution of the wage gap over calendar years 1980-2009; these correspond to the lifespan of the aforementioned cohort and, as such, calendar time coincides with mean age of the household head. I calculate the unconditional gender wage gap in two alternative ways: (a) as the ratio of median male to median 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 6 plots both measures 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 this period: 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 cohort[13] 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.

[13]The 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 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:* PSID 1980-2009.

The narrowing of the gender wage gap is robust to a number of richer specifications. Figure 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 that graph I define the gender wage gap as

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

where $\widetilde{w}_{jit}$ 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 observable characteristics and correcting women's wages for self-selection in the labor market.[14] Figure 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 that graph I define the gender wage gap as

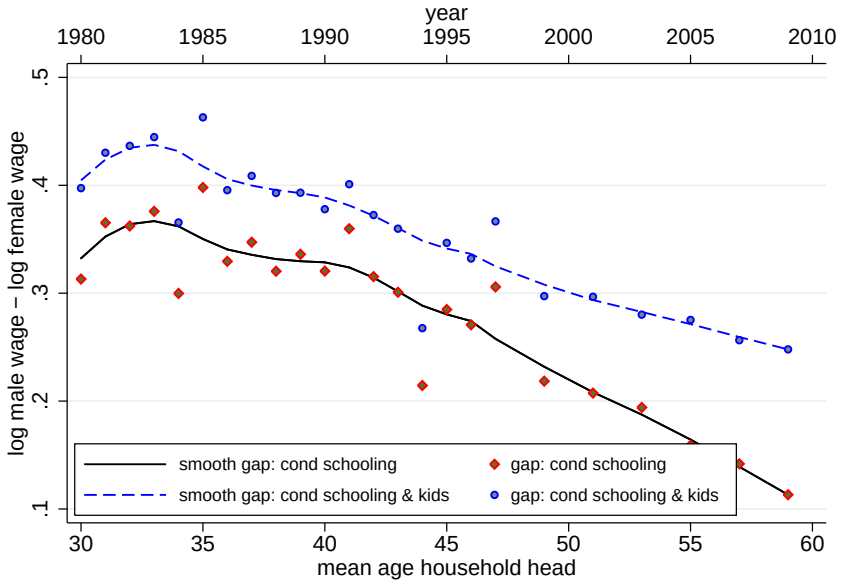
$$GWG_t = \text{median}(\widetilde{\Delta w}_{it})$$

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

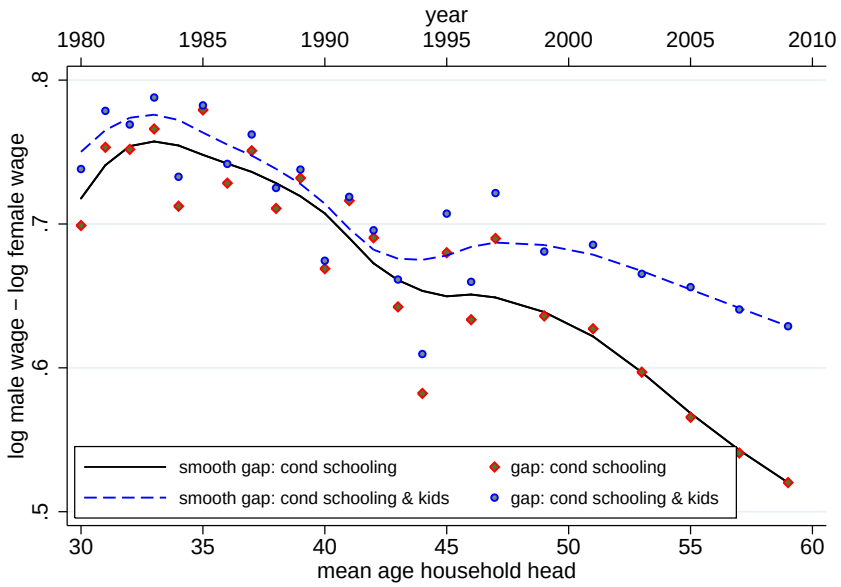
Across all figures the picture that emerges points to an improvement of the economic status of women relative to that of men (at least as reflected upon their wages). Such improvement is robust to a number of factors that potentially affect the gender wage gap, such as women's education, labor market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

[14] I 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 2a).

Figure 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 median($\widetilde{w}_{1it}$) − median($\widetilde{w}_{2it}$) where $\widetilde{w}_{jit}$ is 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 median($\widetilde{\Delta w}_{it}$) where $\widetilde{\Delta w}_{it}$ 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:* PSID 1980-2009.

In a series of papers, Blau and Kahn (1997, 2006) investigate the reasons 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 paper aims to investigate the extent to which an *exogenous* narrowing of the gender wage gap can help explain family time allocation.

A caveat is due here. As the paper 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 paper 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* or *across cohorts*.

3 A Life-Cycle Collective Model without Commitment

This section develops the life-cycle collective model of family time allocation, public consumption and savings, and 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 3.1 I summarize the key features of the model during the working period. The details are given in section 3.2, where I lay out the model's building blocks including its recursive formulation, and in section 3.3, where I detail the model's parametric specification. Section 3.4 describes the retirement period.

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

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

[15] However, see Dauphin et al. (2011) and Dunbar et al. (2013) for environments were children can act as decision makers.

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 B.

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

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

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.

I do not model marriage decisions; instead the focus of the paper is on the partners' choices after they have formed a household (i.e. conditional on marriage). However, the model accounts for initial conditions that arise from assortative patterns in the marriage market (see the wage process in section 3.3).

Ex post, 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 in the model are of the same age. However, 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 period. The participation constraints 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 paper, 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 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

[16] Consistent 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 consider non-cooperative cohabitation as an alternative outside option (see, for example, Lechene and Preston, 2011).

[17] By 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.

will later represent them.

The value of each spouse's time in the labor market is captured by the hourly wage they can earn. The model abstracts from human capital accumulation or 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-offs among 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 factor during the working period of life. To capture the impact of children on behavior I model an exogenous stochastic 'fertility' process that reproduces the life-cycle dynamics observed in the data. Individuals make choices conditional on their family composition rather than choosing 'fertility' explicitly (something that would complicate the model considerably).[19]

3.2 Model

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

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

subject to the following constraints

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

$$A_t + \sum_{j=1}^{2} w_{jt} h_{jt} = K_t + CC_t\big(h_{2t}, N_t\big) + \frac{A_{t+1}}{1+r} \qquad A_{t+1} \geq \underline{A}_{t+1} \tag{3}$$

$$U_1\big(Q_t, l_{1t}; \mathbf{z}_{1t}\big) \geq \bar{U}_1(w_{1t}; \mathbf{d}_{1t}; \mathbf{z}_{1t}) \tag{4}$$

$$U_2\big(Q_t, l_{2t}; \mathbf{z}_{2t}\big) \geq \bar{U}_2(w_{2t}; \mathbf{d}_{2t}; \mathbf{z}_{2t}) \tag{5}$$

$$Q_t = f\big(K_t, \tau_{1t}, \tau_{2t}; \mathbf{Z}_t\big) \tag{6}$$

$$l_{jt} + h_{jt} + \tau_{jt} = \mathcal{T} \qquad j = \{1, 2\}. \tag{7}$$

Constraints (3)-(7) must be satisfied in every period t. Expression (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 (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 (3) is the sequential budget constraint linking available resources to expenditure and savings in each period of working life, (4)-(5) are the

[18] The 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. (2016) or Chiappori et al. (2015) for a treatment of schooling choices in the context of a dynamic unitary or collective model respectively.

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

participation constraints in the household, (6) is the household production function, and (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, $CC_t\big(h_{2t}, N_t\big)$ 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 variable that affects 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)-(7). Consider representing the problem by its Lagrangian formulation. Let ν_2 be the Lagrange multiplier on (2); also let $\widetilde{\nu}_{1t}$ be the Lagrange multiplier on participation constraint (4) and $\widetilde{\nu}_{2t}$ on (5). Then the above problem is equivalent to

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

or, written more compactly, to

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

subject to constraints (3), (6) and (7) only (Chiappori and Mazzocco, 2017). In such case $\mu_{jt} = \nu_j + \frac{\widetilde{\nu}_{jt}}{\beta_t}$ is individual j's intra-family bargaining power at time t ($\nu_1 = 1$) or, equivalently, the weight his/her preferences carry in the household decision process at that time. Moreover,

[20] Chiappori 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).

if one imposes 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 (2) and, as a consequence, contribute to determining the initial weight ν_j. $\widetilde{\nu}_{jt}$ is the multiplier on j's participation constraint in period t. Whatever affects the outside option $\bar{U}_j$ at t will affect $\widetilde{\nu}_{jt}$ 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 ex ante to never exploit their outside options, which is equivalent to removing the participation constraints, $\widetilde{\nu}_{jt} = 0$ and $\mu_{jt} = \nu_j$ in each period ('full commitment' benchmark). If such commitment is not possible, j's bargaining power shifts when any of the time-dependent factors affecting the outside options, thus $\widetilde{\nu}_{jt}$ too, shifts ('no commitment' benchmark). As an example, an increase in j's wage may improve her outside option and result in $\widetilde{\nu}_{jt} > 0$; this will raise her bargaining power by $\frac{\widetilde{\nu}_{jt}}{\beta_t}$ and decrease her partner's power by the same amount.

Note that the model in this paper 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 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 paper 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 prevent the spouses from reaching the first-

[21] This 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 $t \neq s$.

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 without violating the prevailing participation constraints once information at time t is revealed; for details see section 6.2.2 in Browning et al. (2014) or Chiappori and Mazzocco (2017).

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

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

Program (1') can be written recursively as

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

subject to constraints (3) and (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 because I do not solve for the value of divorce numerically. After I initialize intra-family bargaining power at the start of the life-cycle, bargaining power changes subsequently with the gender wage gap so as to match married people's time allocation profiles; no further information on divorcees is used or needed by the model. The initialization draws on information from divorced individuals in the PSID in a reduced-form way (this point is discussed in section 5.3). This shortcut eases the computational needs of the model (saving the burden of solving for the divorced man's and woman's life-cycle problems) but comes at the cost of restricting the estimation sample to stable households only (as shown in section 2.1). I discuss the implications of this restriction in section 6.

3.3 Parametrization

In each period of their life in the household, which I take to be one year, the spouses maximize expected lifetime utility (1') taking as given their individual characteristics and economic circumstances. These are described by their common age (t), their common assets (A), the presence of kids in the household and the age of the youngest among them (N), their respective idiosyncratic productivity in the labor market (v_1, v_2) and the female's utility cost of work (θ_2).

Time allocations I assume time can take on discrete values across three activities: market work, work in the household, 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 2 summarizes the discrete values market and household work can take. Market work can take on three values ('no work', 'part time', 'full time') while work in the household four values ('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 2 serves as a discrete approximation to the empirical distribution of time observed in the data.

Table 2: Time into market and household work

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

For computational reasons, not all choices of market and household work are available to both men *and* women. I restrict men's household work to 'low' or 'low middle' and women's household work to 'high middle' or 'maximum' (consistent with men's and women's observed household hours after 1980). In addition, I restrict men's market work to 'full time' only as there are very few men in the PSID sample that do not work full-time.[22] 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 the non-separable function

$$U_j\big(Q_t, l_{jt}; \mathbf{z}_{jt}\big) = \frac{1}{1-\gamma}\big(Q_t/s(N_t)\big)^{1-\gamma} \times \exp\big(g_j(l_{jt}; \mathbf{z}_{jt})\big) \tag{8}$$

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.[23] Function $g_j(\cdot)$ reflects how

[22] Part-time work for men does increase slightly towards the end of the life-cycle (figure 2b). However, I shut down men's market work choice for reasons pertaining to the feasibility of the computations herein.

[23] The 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.

the marginal utility of consumption changes with leisure (and thus with market and household work) and depends on N_t through $\mathbf{z}_{jt}$. I specify

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

The sum $\sum_n$ is over the different values $\{n\}$ of the age of the youngest child (if any): for all possible values $\{n\}$, $\{\mathbf{1}[N_t = n]\}_n$ constitutes a set of mutually exclusive dummies, each of which becomes active whenever $N_t = n$. Parameters κ_j^n depend on the amount of leisure individual j enjoys and thus on the amount of market and household work he/she supplies.[24] 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.[25] 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 is one only relevant preference shifter affecting U_j, namely $\mathbf{z}_{jt} = N_t$.

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

$$f\big(K_t, \tau_{1t}, \tau_{2t}; \mathbf{Z}_t\big) = K_t^{\phi}\big(\pi_1 \tau_{1t}^{\varphi} + \pi_2 \tau_{2t}^{\varphi}\big)^{\frac{1-\phi}{\varphi}} \tag{9}$$

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 (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 pension income in section 3.4). This is not a generous borrowing limit as lifetime earnings are hindered by the possibility that the female spouse abstains from market work. Wages w_{jt} and child care costs $CC_t\big(h_{2t}, N_t\big)$ 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

$$\begin{aligned} \ln w_{jt} &= \overline{W}_{jt} + v_{jt} + \xi_{jt} \\ v_{jt} &= v_{jt-1} + \zeta_{jt}. \end{aligned} \tag{10}$$

This process has been shown to fit the PSID data well (Blundell et al., 2016). The hourly wage

[24] For 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 normalize $\kappa_{j,0}^n = 0$, $\forall j, n$; 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.

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

is assumed exogenous and the individuals are viewed as price-takers in the labor market. $\overline{W}_{jt}$ is the mean of j's log wage at t which is common across people of the same gender j and, in principle, education.

The sum $v_{jt} + \xi_{jt}$ represents stochastic idiosyncratic productivity consisting 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

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

This process is estimated directly from the data and details are provided in section 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 for example); 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. (2016). 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 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

$$\mu_{1t} = \frac{\exp\left(\eta(m_t)\right)}{1 + \exp\left(\eta(m_t)\right)} \tag{11}$$

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

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

$\eta^{(n)}$ reflects how intra-family bargaining power changes with the gender wage gap *within* a particular family ($\frac{w_{1t}}{w_{2t}}$) and varies with family composition (N_t). In this specification $\mathbf{x}_j$, $\mathbf{d}_{1t}$ and $\mathbf{d}_{2t}$ are left empty; however, in a model where education of each spouse is explicitly modeled in the state space, education will likely enter $\mathbf{x}_j$.

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 $\kappa_j^{(n)}$, (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 depends on age and is directly estimated in the PSID.

The transition between fertility states depends on age 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

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

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 $CC_t\big(h_{2t}, N_t\big) = cc_t\big(h_{2t}, N_t\big) \times \text{Prob}_t(costs > 0 \mid N_t)$ 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 $cc_t\big(h_{2t}, N_t\big)$ 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 $cc_t\big(h_{2t}, N_t\big)$ can be summarized by

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

The hourly price of child care is $cchrate_t$ and varies with calendar time. Section 5.1 provides details on the estimation of $cchrate_t$ and the probability of positive child care costs.

3.4 Retirement

Retirement starts at time $T+1$ and ends at time T^R for both spouses. During this period the spouses 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.[26,27]

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 paper). 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 (8)).

Adopting a compact formulation equivalent to expression (1'), the household solves during retirement

$$\max_{\{Q_t, A_{t+1}\}_{t=T+1}^{t=T^R}} \sum_{t=T+1}^{T^R} \beta_t \Big[\mu_{1T} U_1\big(Q_t, l_{1t}; \mathbf{z}_{1t}\big) + \mu_{2T} U_2\big(Q_t, l_{2t}; \mathbf{z}_{2t}\big) \Big] \tag{1$^{\mathrm{R}}$}$$

subject to

$$A_t + \sum_{j=1}^{2} I_{jt} = Q_t + \frac{A_{t+1}}{1+r} \qquad A_{t+1} \geq \underline{A}_{t+1} \tag{12}$$

$$l_{jt} = \mathcal{T} \qquad j = \{1,2\} \tag{13}$$

$$A_{T^R} = 0. \tag{14}$$

Most of the notation has been introduced previously. Preferences U_j are given by (8) and the Pareto weight μ_{jT} by (11); the vector of observable taste shifters is now $\mathbf{z}_{jt} = (N_t = 1)$ as there is no fertility. The budget constraint is slightly different from (3) in that earnings are replaced by pension income I_{jt}, the public good Q directly enters the constraint (as there is no home production), and there are no child care costs. For each spouse, pension income is set to a deterministic 75% of their average full-time earnings at the start of their working life.

[26] Allowing 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 paper, other than the within-household gender wage gap, thus I fix the retirees' intra-family bargaining power to its last value prior to retirement.

[27] In 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.

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.[28] Q_t and A_{t+1} are the choice variables. Program (1^{R}) can be written recursively as

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

subject to the budget constraint (12), the time budget (13), and terminal condition (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.[29]

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

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

subject to the constraints of the working period.

4 Model Solution and Simulation

In this section I describe the steps I take to solve and simulate the model developed in section 3. 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 3.3. A time period is taken to be 1 year.[30]

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

[28] The 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 27.

[29] Fernández and Wong (2014) also model a 'cake-eating' and non-stochastic retirement period.

[30] Currently the paper 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.

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, 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 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 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' age t is the discrete state 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 2 of section 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 (11).[31] 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 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 3.3.

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

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 contained 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).[32] 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 family composition directly 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 (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 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.[33] 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 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. I start the simulations assuming households hold 0 initial assets.

The above solution and simulation routines are written in Julia.[34] They run in approximately 40 seconds in total on a 12-core Intel Xeon E5-2630 server at a 2.3GHz clock speed.

5 Identification and Estimation

In this section I describe the steps I take to estimate the structural model of section 3. I follow a two-step procedure. In the first step I estimate the external elements of the model, namely the

[32] I 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.

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

[34] Julia is a new high-performance programming language; documentation is available at http://julialang.org.

wage process for each spouse, the fertility process, and the child care costs. These are all taken as given and the model is conditional on them. In addition, I normalize intra-family bargaining power in the first years of the family life-cycle approximating a gender-specific value of divorce using reduced-form information on divorcees. The role and details of this normalization are 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.

Section 5.1 discusses identification and estimation of the external elements; section 5.2 discusses identification and estimation of the model parameters. This part requires information on the initial intra-family bargaining power, a matter that I discuss in section 5.3.

5.1 External Processes

Wages To construct wage grids, I require i) the mean wage over the life-cycle and ii) the wage variance net of variation in measurement error (transitory shock) for each gender. To integrate out future uncertainty, I require iii) 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 iv) the covariance matrix of permanent shocks over the life-cycle. To obtain simulated wage profiles, I also need v) 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 on the transition matrix for wages are also omitted (but available upon request).

Given the parametrization of the wage process in (10) the second moments of shocks can be readily identified from a combination of 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 shocks.

To obtain estimates of the second moments of shocks I run a minimum distance estimation matching the empirical covariance matrix of 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 C reports the estimation details and a full table of estimates. To reduce the effect of wage measurement error in the estimation of the model, 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 C.1 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 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.

Figure 8: Simulated against actual wages (means)

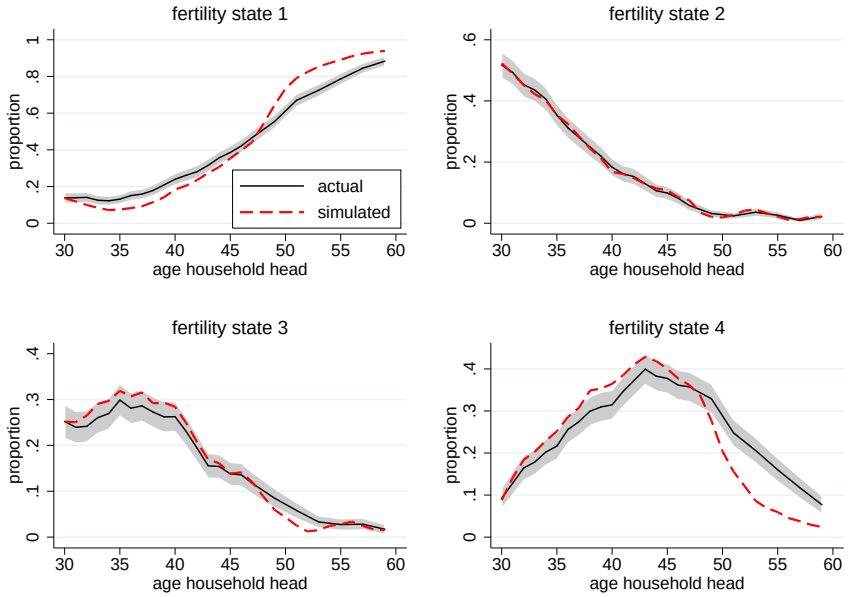


Notes: This figure plots average profiles of simulated and empirical wages over the life-cycle. Simulated wages are based on wage process (10). A 95% confidence interval around the empirical means appears in gray shade. *Source:* PSID 1980-2009 and own simulations.

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 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 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.

Figure 9: Proportion of families in various fertility states



Notes: This figure plots the proportion of families in each fertility state in the actual and simulated data over the life-cycle. A 95% confidence interval around the empirical means appears in gray shade. *Source:* PSID 1980-2009 and own simulations.

Child care costs The hourly rate of child care in function $CC_t(h_{2t}, N_t)$ is $cchrate_t$ 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.[35] This number is somewhat 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".

[35] Blau (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 $0.496 \approx 0.5$.

Given that the PSID data I use cover the years 1980-2009, I adopt the above percentage (50%) and I fix *cchrate* in year 1981, the mid-year between 1976 and 1986, at \$6.59 (expressed in 2010 dollars).[36] A question remains regarding *cchrate* 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) report 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 *cchrate* 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 paper focuses on). Whenever this rate is below the real federal minimum wage, I update *cchrate* to reflect this.[37] 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 over the 3 decades after 1980.

Finally, I calculate the probability of a family facing positive child care costs by estimating the proportion of families in a given fertility status who 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 C.2 in the appendix reports the estimated probabilities as well as the calibrated hourly price of child care over time.

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 an initialization of intra-family bargaining power in the first 10 years of the family life-cycle. The details of this initialization (normalization) are presented in section 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 on estimates available in the literature; these are summarized in table 3.

[36] $cchrate_{1981}$ = \$6.59 is 50% of the average female wage in the PSID between 1976 and 1986.

[37] I 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 3: Fixed parameters

Parameter		Value	Reference
r	interest rate	0.01	Attanasio et al. (2008)
β	discount factor	0.98	Attanasio et al. (2008)
γ	coefficient of relative risk aversion	1.5	Attanasio et al. (2008)
ϕ	output elasticity of public expenditure	0.8	Lise and Yamada (2014)
π	production share of male time	0.5	Lise and Yamada (2014)
φ	scale parameter of spousal time	0.5	Lise and Yamada (2014)

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. (2016). 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.[38]

There is a total of 26 remaining parameters: 4 pertaining to male preferences ($\kappa_1^{(n)}$), 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 allocation in the simulated dataset and I repeat using the actual data. Finally, I calculate a distance metric between the simulated and empirical moments and I repeat the process until the metric is minimized.

[38] Cherchye 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 paper 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.

Formally, the estimated parameters $\hat{\mathbf{\Theta}}$ solve the minimization problem

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

where $\widetilde{\mathbf{M}}_n$ is a $k \times 1$ vector of moments over n observations from the real data, $\mathbf{M}_s(\mathbf{\Theta})$ 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 empirical moments.[39] For the optimization I use `NLopt` (Johnson, 2014) and a number of algorithms therein;[40] 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
  - '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 these same years identify their preferences over full-time and part-time market work ($k_2^{(n)}$, θ_2). Similarly, men's and women's rates of household work in the same years identify their respective preferences over home production time ($k_1^{(n)}$, $k_2^{(n)}$, θ_2). These preferences may differ by family composition and variation across fertility states identifies such differences.

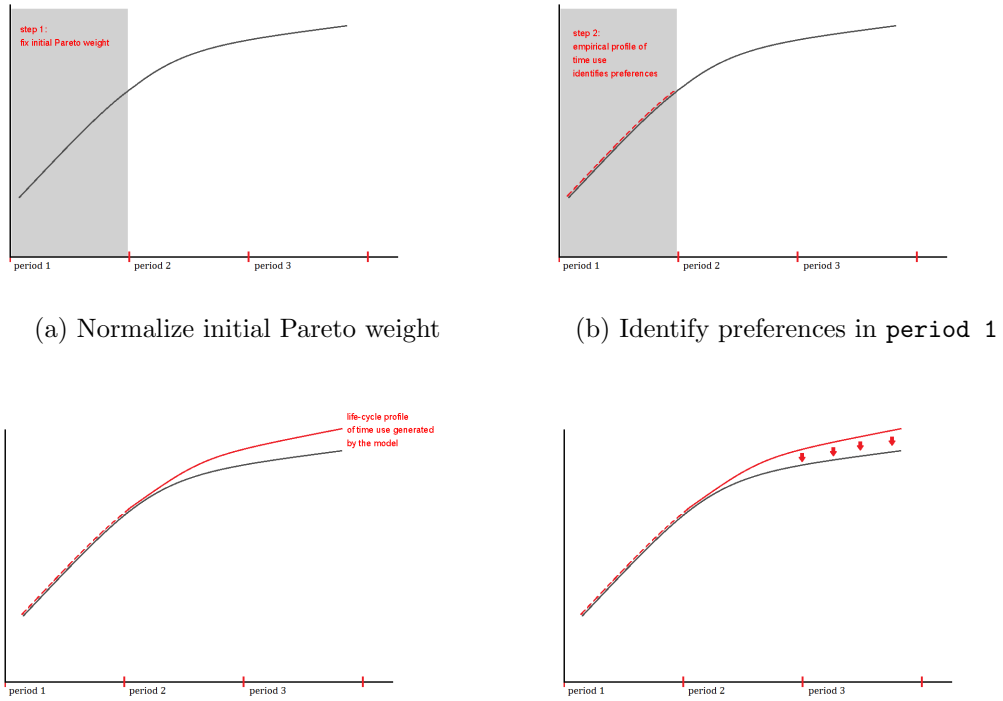
Keeping preferences and the initial normalized Pareto weight fixed, I use the model to generate life-cycle profiles of family time allocations. Then, I allow intra-family bargaining power to shift parametrically in response to changes in the gender wage gap in life-cycle years 11-30, so as to minimize the wedge, if any, between the model-generated and the empirical profiles of

[39] I use the inverse *diagonal* of the covariance matrix of the empirical moments in the light of evidence of small-sample biases with the optimal weighting matrix; see Altonji and Segal (1996).

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

time allocations. This identifies $\eta^{(n)}$.[41] 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.

Figure 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.

These points can be made clearer using the graphical illustration in figure 10. Suppose the life-cycle consists of three periods only, `period 1`, `period 2`, and `period 3`. Suppose also that the solid black line in figure 10 depicts the empirical average profiles of female full-time market work over the life-cycle. I normalize the Pareto weight in `period 1` (shaded gray in graph 10a). Conditional on it, I match the empirical profile in `period 1` with the aim to identify women's preferences for full-time market work (red dashed line in graph 10b). Holding preferences and the initial normalization fixed, I use the model to generate the full profile of female full-time

[41] Identification 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.

market work over the life-cycle (red dashed and solid line in graphs 10b and 10c). Subsequently, I shift the Pareto weight in `period 2` and `period 3` *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 10d).

How do I normalize initial intra-family bargaining power? I discuss this extensively in the following section. As a brief illustration here, I proxy decision power of married spouses by comparing *at the start of their life-cycle* their hypothetical lifetime earnings should they get divorce. To do so I project each spouse's lifetime earnings in the hypothetical scenario of divorce given observable characteristics and information on divorcees from the PSID.[42] This normalization uses information on divorcees because divorce is taken as the relevant outside option. Alternatively I could use an arbitrary normalization as in Fernández and Wong (2014).

5.3 Initialization of the Pareto Weight

The projection of lifetime earnings in the hypothetical scenario of divorce requires information on earnings of divorced persons. I obtain this information from the minor PSID sample of singles whose discussion I postponed previously in data section 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 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.[43] The resulting dataset consists of 4561 divorced male-year and 7614 divorced female-year observations. Some key descriptive statistics are presented in table 4 and appendix A discusses this sample at a greater length. A few differences are apparent between married (table 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

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

where Y^d_{jt+n} is j's earnings as divorcee at time $t+n$ when divorce was enacted at time t. The main difficulty is to estimate expected future earnings.

[42] Other 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.

[43] This sample includes those separated alongside those formally divorced.

Table 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 12^{th} grade. Annual earnings are expressed in 2010 dollars. *Source:* PSID 1980-2009.

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

$$Y_{kt} = \boldsymbol{\chi}_k' \mathbf{a}_k + \boldsymbol{\psi}_{kt}' \mathbf{b}_k + \epsilon_{kt}, \qquad k = \{1, 2\} \tag{15}$$

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 (15) along with information known at t about the concerned person j. I construct $\widehat{Y}_{jt+s} = \boldsymbol{\chi}_j' \widehat{\mathbf{a}}_k + \boldsymbol{\psi}_{jt+s}' \widehat{\mathbf{b}}_k$, $k = j$ and $j = \{1, 2\}$, and I use it in place of $E_t(Y^d_{jt+s})$. To generate the sequence of expected future earnings I assume that individuals work until age 65 and $\rho = \beta^{-1} - 1$.

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 as

$$\mu_{1t} = \frac{\text{Human Wealth Divorced}_{1t}}{\text{Human Wealth Divorced}_{1t} + \text{Human Wealth Divorced}_{2t}}$$

where subscripts 1 and 2 point to the male and female spouses respectively. 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.[44]

[44]The way intra-family bargaining power is constructed favors mechanically the youngest 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 by an equivalent annual annuity; specifically I divide 'Human Wealth Divorced$_{jt}$' 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.

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 (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 (15) appear in table C.3. 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 couple's life-cycle given the restriction of my estimation sample to one cohort only). Table 5 reports descriptive statistics for the derived bargaining power of married males in years 1980-1989; figure 11 is a graphical representation of its cross-sectional distribution.

Table 5: Initialized intra-family bargaining power of men in 1980-1989

	mean	median	std.dev.	minimum	maximum	N
Men's bargaining power	0.6358	0.6208	0.0673	0.4022	0.9106	5970

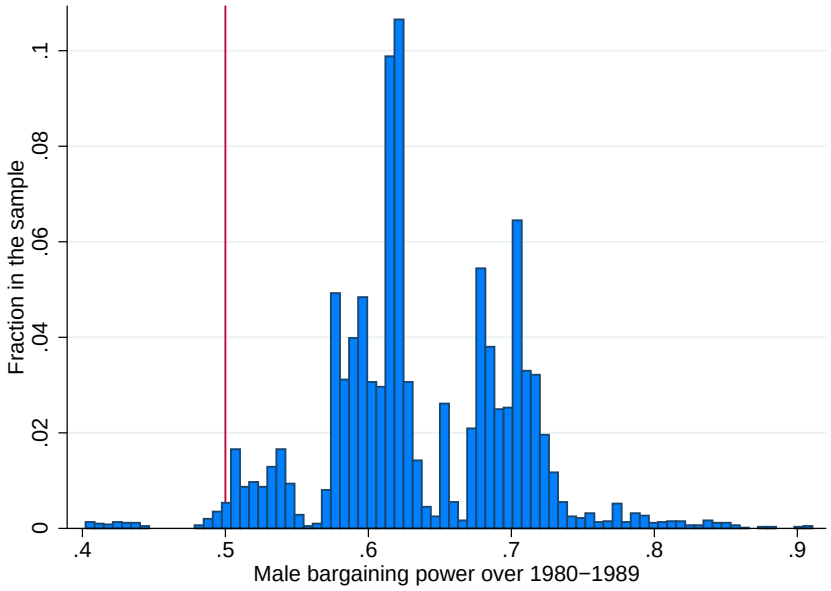
Notes: This table reports descriptive statistics for the derived intra-family bargaining power of married men in the major PSID sample in years 1980-1989. 'N' refers to the number of married household-year observations.

I estimate an average $\overline{\mu}_1 = 63.58\%$ in years 1980-1989; the median is slightly lower at 62.08%.[45] The maximum value is approximately 91% but there are also households where 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 a structural model that postulates that $\mathbf{x}_j = \{educ_j\}$ affects the initial allocation of bargaining power (but here I abstract from this feature). Appendix table C.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 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 sign of coefficients in the Pareto weight function or the gradient of each spouse's preference parameters with respect to leisure.

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

Figure 11: Histogram of initialized intra-family bargaining power of men in 1980-1989



Notes: This figure plots the distribution of 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 weights between spouses).

6 Results

This section presents the estimates of the structural parameters and displays the model fit. Table 6 reports the estimates for the parameters of preferences, namely the components of the $g_j(\cdot)$ functions in (8). Panel A reports preferences for market work and panel B reports preferences for household work.

Reading through table 6 note that *within* a type of time use, for example within women's full-time market 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, induces greater disutility as utility in (8) is negative due to γ set equal 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. (2016) 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 (panels A(I) and A(II), row 1).

Table 6: Estimates of preferences

	A. Women's market work			
	(I) full-time work		(II) part-time work	
	value	st.error	value	st.error
(1) No children	0.174	(0.0023)	0.143	(0.0024)
(2) Youngest child: up to 5	0.262	(0.0025)	0.185	(0.0024)
(3) Youngest child: 5-10	0.227	(0.0029)	0.183	(0.0015)
(4) Youngest child: 10+	0.222	(0.0024)	0.158	(0.0015)
Unobserved types:				
(5) Type I: utility cost of work	0.131	(0.0023)	0.020	(0.0022)
(6) Type II: utility cost of work	-0.321	(0.0099)	-0.023	(0.0023)
	B. Household work			
	(I) men 'low middle'		(II) women 'maximum'	
	value	st.error	value	st.error
(1) No children	0.0600	(0.0004)	0.109	(0.0004)
(2) Youngest child: up to 5	0.0613	(0.0000)	0.106	(0.0004)
(3) Youngest child: 5-10	0.0608	(0.0001)	0.097	(0.0004)
(4) Youngest child: 10+	0.0620	(0.0007)	0.104	(0.0003)
Unobserved types:				
(5) Type I: utility cost of work			0.012	(0.0008)
(6) Type II: utility cost of work			-0.004	(0.0003)

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

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 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 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 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. However, the amount of time each gender devotes to household work is very different. 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 unobserved 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 7 reports the coefficients $\eta^{(n)}$ on the gender wage gap from the Pareto weight function (11). Note again that, as these are parameters corresponding to different family compositions (rows 1-4), they are mutually exclusive in (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 (11), lowers men's bargaining power in the household, and increases women's by an equal amount.

Table 7: Estimates of gender wage gap coefficients 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 (11). Asymptotic standard errors are reported in parentheses.

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 done in section 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.[46] 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

[46]Using the estimates in table 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 (11) and I assume wages in the household change intertemporally according to the gender wage gap in figure 6. The

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.[47] Second, targeting '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 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 8 reports selected moments of time allocation in the data alongside their counterparts from the model simulations. These moments are weighted means of targeted moments but are not otherwise targeted directly. The full set of targeted empirical moments, alongside their model counterparts, appears in tables D.1-D.2 in the appendix.

Table 8: Proportions of married people in different time allocations

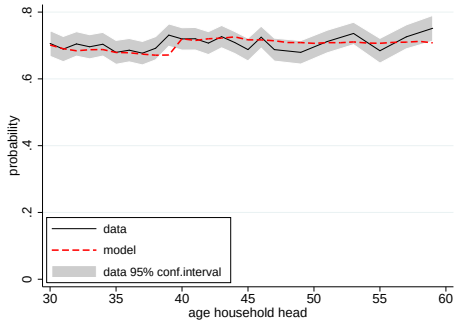
	Data	Model
A. Men		
'low middle' household work	0.705	0.703
'low' household work	0.295	0.297
B. Women		
FT market work and		
'maximum' household work	0.295	0.318
'high middle' household work	0.352	0.370
PT market work and		
'maximum' household work	0.121	0.098
'high middle' household work	0.039	0.037
No market work and		
'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' housework refer to different amounts of time put into home production; see table 2 for details. *Source:* PSID 1980-2009 and model simulations.

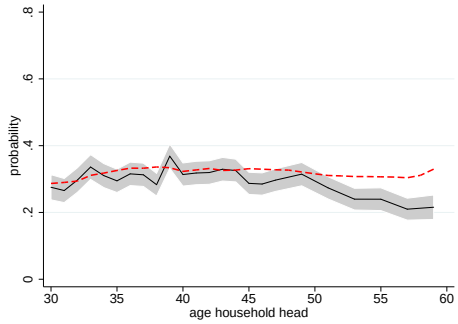
magnitude of this change, unlike its direction, is not of practical interest as it is subject to the initial normalized Pareto weight and the cardinalization of preferences.

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

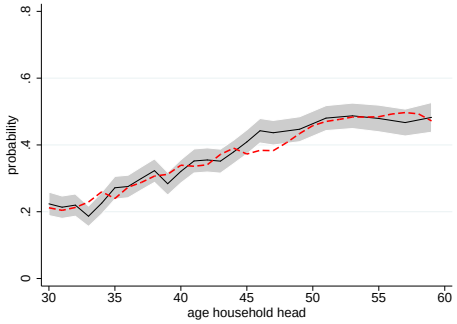
Figure 12: Life-cycle model fit



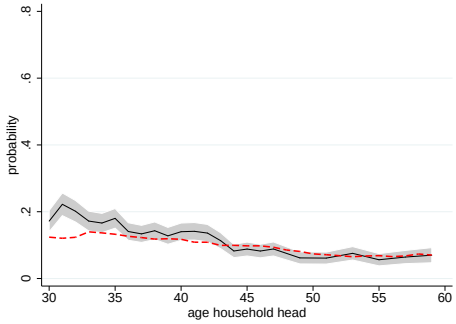
(a) Men: 'low middle' household work



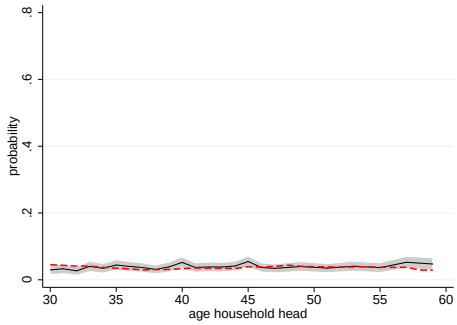
(b) Women: full-time market work & 'maximum' household work



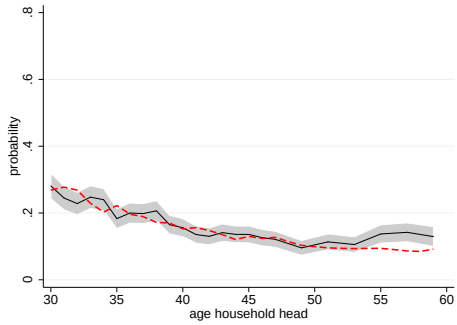
(c) Women: full-time market work & 'high middle' household work



(d) Women: part-time market work & 'maximum' household work



(e) Women: part-time market work & 'high middle' household work



(f) Women: no market work & 'maximum' household work

Notes: This figure plots the life-cycle model fit averaged over different family compositions. A 95% confidence interval around the empirical 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 2 for details. *Source:* PSID 1980-2009 and model simulations.

Figure 12 illustrates the life-cycle model fit across different time allocations of men and women (averaged over various family compositions). The model performs reasonably well. 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 (figure 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 12f). There is currently no mechanism in the model inducing early retirement, such as compulsory receipt of social security benefits crowding out labor earnings (French, 2005); 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 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 D.1 in appendix D. The model succeeds in replicating most of those moments too.

7 Implications of Model for Behavior

This section discusses the implications the narrowing gender wage gap has for married people's time allocation. Specifically, I illustrate three effects (income, substitution, bargaining) changes in the gender wage gap induce on behavior.

Changes in wages or the gender wage gap affect family behavior through a number of channels:

1. A higher female wage likely renders 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.[48]

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 boost 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 on all outcome variables.

In the data, the average within-family gender wage gap, which I feed into the model through budget constraint (3) and intra-family bargaining power (11), narrows down on average by approximately 10% over the family lifetime (graph 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).

[48] An inspection of the production function (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 substitute one another.

My aim is to separate and quantify the two aggregate types of effects. I proceed as follows. Using the preference estimates from section 6 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 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.[49] 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 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 9 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.

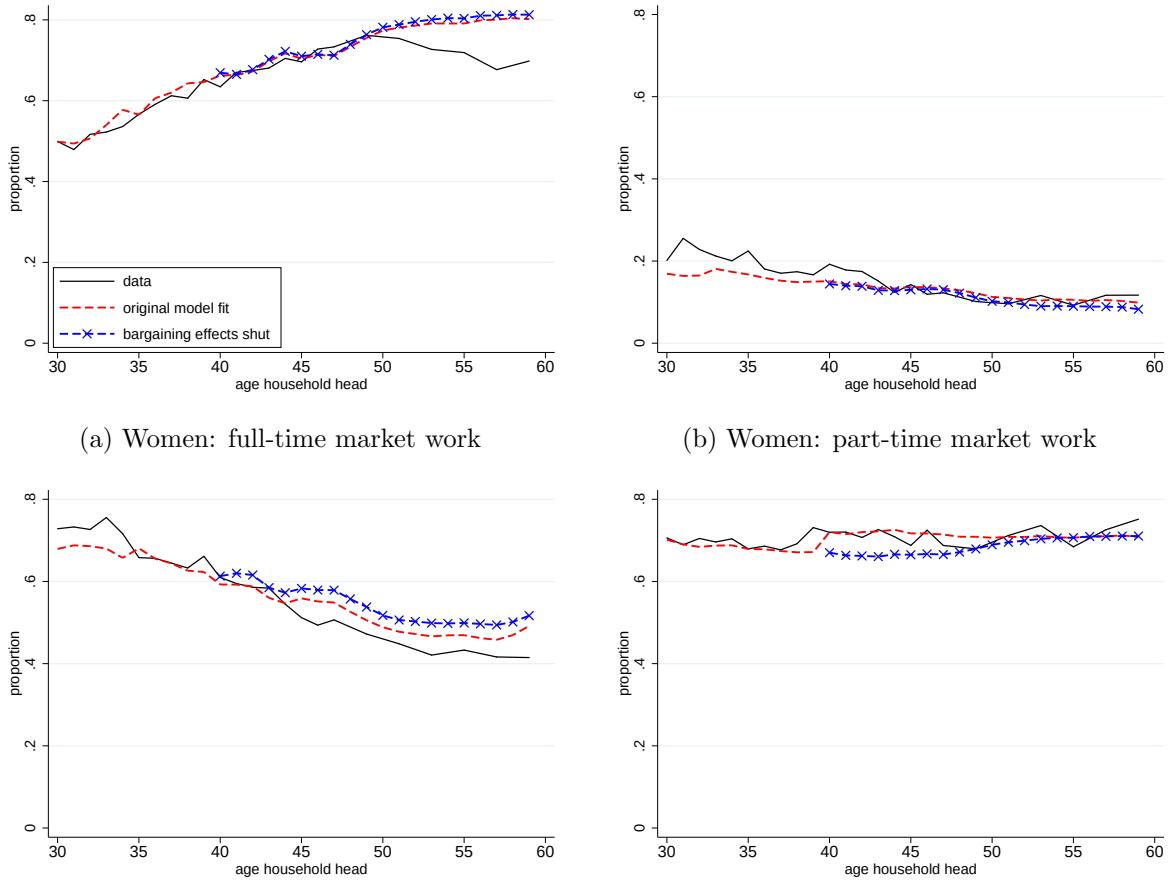
In the first application, women's bargaining power is not allowed to improve alongside the narrowing gender wage gap. From table 9, this induces *up to* 3.06% more women into working 'maximum' hours in the household and *up to* 4.98% fewer men into supplying 'low middle'. When compared to the baseline 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%.[50] Up to 1.03% more women 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 supply more time to both the market and the household, thus enjoy less leisure. The opposite holds for men.

[49] Specifically 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.

[50] The 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 13: Bargaining effects of a narrowing gender wage gap



(a) Women: full-time market work

(b) Women: part-time market work

(c) Women: 'maximum' household work

(d) Men: 'low middle' household work

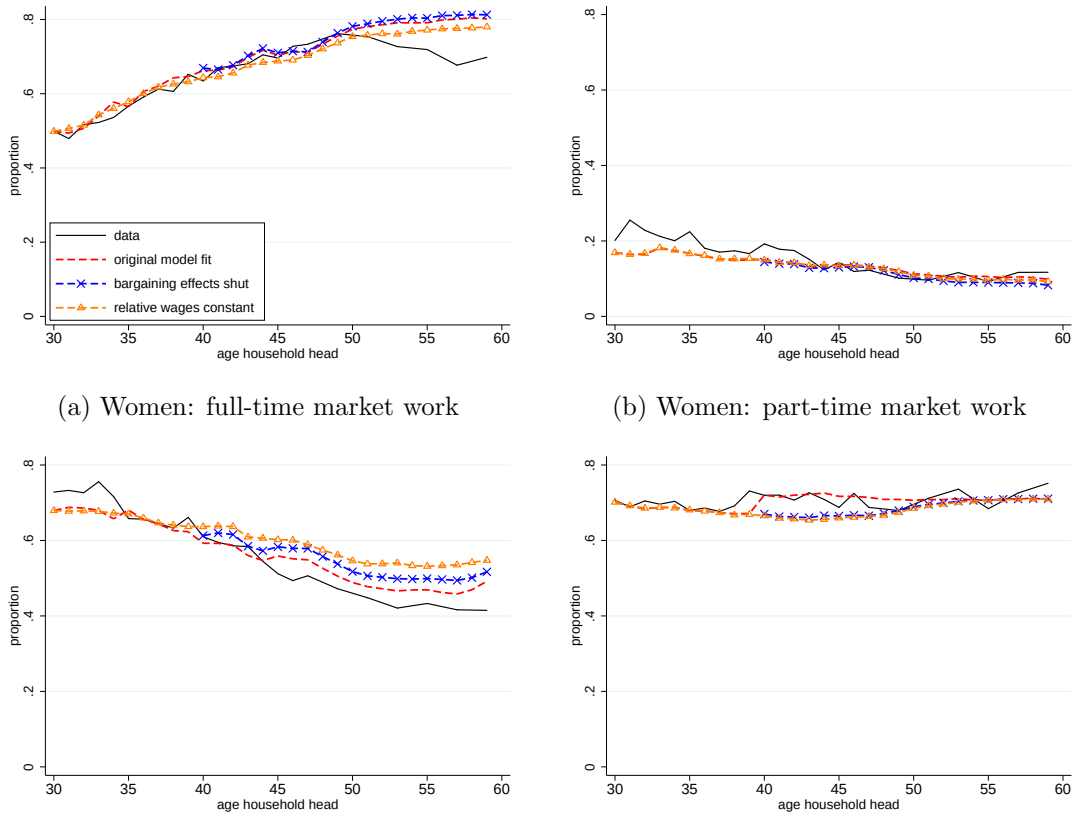
Notes: This figure plots life-cycle profiles of family time allocations when intra-family bargaining power does not respond to the gender wage gap (bargaining effects shut; blue dashed line through the `X`'s). The solid line depicts the original data and the red dashed line depicts the original model 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 2 for details. *Source:* PSID 1980-2009 and model simulations.

Table 9: Foregone 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 'low middle' 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 initialization of the Pareto weight in section 5.3). The definitions of 'maximum' or 'low middle' refer to different amounts of time put into home production; see table 2. *Source:* Model simulations.

Figure 14: Income/substitution effects of a narrowing gender wage gap



(a) Women: full-time market work

(b) Women: part-time market work

(c) Women: 'maximum' household work

(d) Men: 'low middle' household work

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 2 for details. *Source:* PSID 1980-2009 and model simulations.

Table 10: Foregone 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 'low middle' 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. When the gender wage gap narrows down, it induces income/substitution and bargaining effects. Keeping the gender gap constant, as this application suggests, removes both effects. The table reports exclusively the changes due to the foregone income and substitution effects. The figures presented inside the square brackets are the model's rates factoring in (adding) the bargaining effects of table 9. The definitions of 'maximum' or 'low middle' refer to different amounts of time put into home production; see table 2. *Source:* Model simulations.

The results from the second application suggest that keeping the average gender wage gap constant at its beginning-of-life level has, through removing the income and substitution effects, important implications for women's labor supply and household work. Figure 14 presents life-cycle profiles of family time allocations (orange lines through the triangles) when women's wages grow as per men's observed wage growth and, as a result, the gender wage gap remains unchanged. Table 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 foregone income/substitution effects should the gender wage gap remain fixed; model rates (in %) factoring in the foregone bargaining effect appear in square brackets on the side.[51]

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.

Inspecting figure 14c, 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 (orange lines through the triangles vs. red dashed line). Half of its effect is due to the higher monetary reward of women's market work, thus to women switching to some form of market work (orange lines through the triangles vs. blue dashed line through the `X`'s). The other half is due to women becoming relatively stronger in the household decision process, and therefore able to enjoy more leisure (blue dashed line through the `X`'s vs. red dashed line).

The narrowing of the gender gap is also important for women's full-time market work, contributing to its overall rise over the life-cycle. Two opposite forces operate here: the higher monetary reward pushes women's market work up (dominating force; blue line through the `X`'s vs. orange lines through the triangles in figure 14a) while women's improved bargaining power pushes work down replacing it with leisure (red dashed line vs. blue line through the `X`'s).

Finally, as the income women bring in the household rises, the spouses are in a better financial position to replace certain household chores with services purchased on 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.

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 14 illustrates. This suggests that life-cycle considerations, fertility dynamics and child care costs are still jointly quite important even after accounting for the gender wage gap.

[51] The figures presented in the square brackets are the model's original baseline rates adding the foregone bargaining effect of table 9. 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 10 reports exclusively the former as the latter are already reported in table 9.

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 around the world pledge to implement.[52] With its explicit focus on the gender wage gap, the model allows me to investigate equal pay, abstracting, however, from potential important issues pertaining to education or experience.

I investigate the implications of equal pay through three counterfactual experiments. The experiments have similar 'flavor' as across all three of them I shift the mean of female wages towards the mean of male wages. The timing that this shift occurs within married people's lifetime differs across experiments.[53]

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 $\overline{W}_{2t} = \overline{W}_{1t}$ at all times (where $\overline{W}_{1t}$ 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 the observed average levels $\overline{W}_{1t=0}$ and $\overline{W}_{2t=0}$ 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 on average the same wages (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 the empirical gender-specific variation in wages. Table 11 summarizes the features of the three counterfactual experiments.

Table 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

[52] 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". His full speech is available at www.whitehouse.gov/blog/2014/04/12/weekly-address-ensuring-equal-pay-equal-work. 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.

[53] In a more detailed model with multiple education levels, the counterfactual experiments would involve equalizing the *mapping* from education to wages between genders.

The results across all experiments are illustrated numerically in table 12 and visually in figure 15. Figure 15 plots life-cycle profiles of time allocations; across all graphs 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.

Gender wage equality has important implications for men's and women's time allocations. Equal pay induces women's entry in the labor market even during the child bearing years. It increases the rates of full-time market work up to 18.36 percentage points (*ppt.*) in certain years and decreases part-time work, albeit by less. It lowers women's likelihood of supplying 'maximum' hours in the household up to 14.55*ppt.* and increases men's merely by 2.18*ppt.* in certain years.

The most sizable effects on married people's time allocations are seen in counterfactual #1 where men's and women's average wages are equal 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,[54] equal pay implies that up to 75.36% of women now 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*ppt.* more likely to work full-time. Equal pay makes women aged 30-39 up to 2.28*ppt.* *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 are generally 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 implies.

The proportion of women aged 30-39 supplying 'maximum' hours to home production drops 14.55*ppt.* (from a baseline of 66.2% to 51.7%), whereas the proportion of men supplying 'low middle' increases 2.18*ppt.* points (from a baseline of 68.2% to 70.4%). 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 at ages 30-39 is equivalent to up to 7 hours of home work less per week (see figure 15e).[55]

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*ppt.* (8.64*ppt.*) 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*ppt.* (3.63*ppt.*) 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.

[54] The 'model baseline rates' refer to the proportions observed in the originally simulated data after fitting the model to the PSID.

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

Table 12: Counterfactual family time allocation

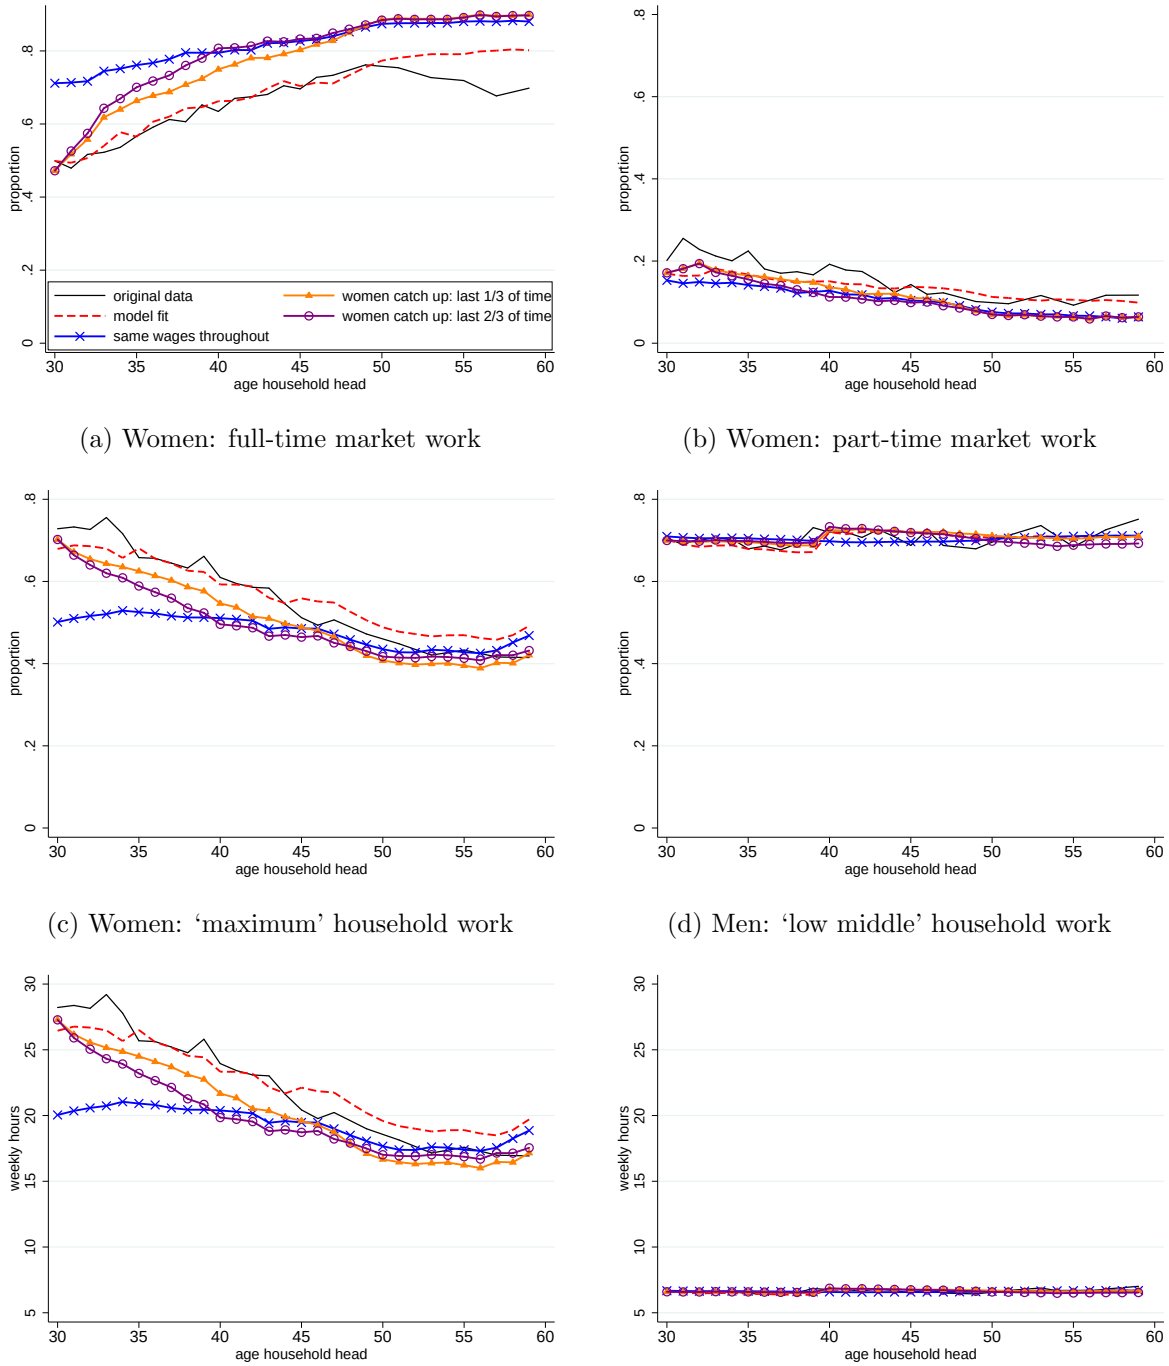
	(1) women full time work		(2) women part time work		(3) women 'max' home work		(4) men 'low middle' home work	
Experiment 1: equal average wages throughout life-cycle								
ages 30-39	+18.36	[57.0%]	-2.28	[16.3%]	-14.55	[66.2%]	+2.18	[68.2%]
ages 40-49	+12.29	[70.3%]	-3.00	[13.6%]	-7.27	[55.7%]	-1.99	[71.7%]
ages 50-59	+8.64	[79.2%]	-3.71	[10.5%]	-3.63	[47.2%]	+0.03	[70.9%]
Experiment 2: women catch up with men in year 21 *of* 30								
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%]
Experiment 3: women catch up with men in year 11 *of* 30								
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 2. *Source:* Model simulations.

Counterfactual #1 generates a higher and flatter female life-cycle profile of full-time market work compared to the data. If men and women earn on average the same wages, the model predicts that more than 75% of women work full-time right from the beginning of their working life; that rises 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 few only differences between all three experiments in the last 10 years of the life-cycle (when for all three average wages are equal between spouses) but there are some noticeable differences in the earlier years.

Figure 15: Counterfactual life-cycle profiles of time allocation



(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 allocation 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.

Counterfactual #2 (orange lines through the triangles) 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*ppt*. increase in full-time market work rates compared to the baseline model (table 12, experiment #2, column 1). Experiment #3 (purple lines through the hollow circles) 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*ppt*. increase in women's rates of full-time work compared to the original model (tab.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*ppt*. in women's rates of part-time work compared to the original model (tab.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.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 experiment #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.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 and a more equal allocation of household work between men and women (even though women still work more in the household than men). The time mothers spend with children is not modeled in this paper and it remains an open question how equal pay impacts on this important dimension (for example Del Boca et al., 2014).

9 Discussion and Conclusions

Over the last three decades the wage gap between men and women in the United States narrowed down substantially. This paper 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 paper 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. PSID data 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 (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.

The model reproduces the empirical life-cycle profiles of family time allocations 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, thus 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. 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 paper investigates how the gender gap affects family time allocations *over a particular life-cycle*. Extending the paper 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 4 decades. Wages are taken as exogenous and the model abstracts from human capital; this raises 'reverse causality' concerns especially: a narrower gender wage gap may be driving family time allocations but women's gradual entry into the labor market may also 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. It will also require to cardinalize preference differences between married and divorced persons. However, an additional advantage is that solving for divorcees 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 a wider range of 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 or services beyond the price of child care, as in Greenwood et al. (2005), is likely to be important for the patterns of household work. Finally, use of consumption data can help identify a number of consumption parameters and possibly characterize consumption allocations between spouses.

References

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., H. Low, and V. Sanchez-Marcos (2008). Explaining changes in female labor supply in a life-cycle model. *American Economic Review 98*(4), 1517–1552.

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, and C. Meghir (2005). Collective labor supply with children. *Journal of Political Economy 113*(6), 1277–1306.

Blundell, R., M. Costa Dias, C. Meghir, and J. Shaw (2016). Female labor supply, human capital, and welfare reform. *Econometrica 84*(5), 1705–1753.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2015). Children, time allocation and consumption insurance. IFS Working Papers W15/13, Institute for Fiscal Studies.

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

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.

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., 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.

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. *American Economic Review 102*(7), 3377–3405.

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., 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., 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 (2017). Static and intertemporal household decisions. *Journal of Economic Literature 55*(3), 985–1045.

Chiappori, P.-A. and C. Meghir (2014). Intrahousehold inequality. NBER Working Papers 20191, National Bureau of Economic Research, Inc.

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. *The Review of Economic Studies 81*(1), 137–185.

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. *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.

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. *The Review of Economic Studies 72*(2), 395–427.

Gemici, A. (2011). Family migration and labor market outcomes. Unpublished manuscript.

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.

Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. *Econometrica 47*(1), 153–61.

Johnson, S. G. (2014). The NLopt nonlinear-optimization package [Software]. http://ab-initio.mit.edu/nlopt.

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.

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.

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.

Lechene, V. and I. Preston (2011). Noncooperative household demand. *Journal of Economic Theory 146*(2), 504–527.

Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *The Review of Economic Studies 69*(1), 209–44.

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.

Manser, M. and M. Brown (1980). Marriage and household decision-making: A bargaining analysis. *International Economic Review 21*(1), 31–44.

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.

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.

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. *The Review of Economic Studies 63*(2), 199–235.

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. Washington, DC: Center for the Child Care Workforce.

Appendices

A Data: Sample Selection and Variables

The paper 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 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 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 as 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 age-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.[56]

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 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 and the available measure excludes some previously included minor components of earnings such as the labor part of business income.[57] 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 2) using a 3-point

[56]The proportion of households that are 'stable', among all households satisfying the other selection criteria laid out in this paragraph, is approximately 81%.

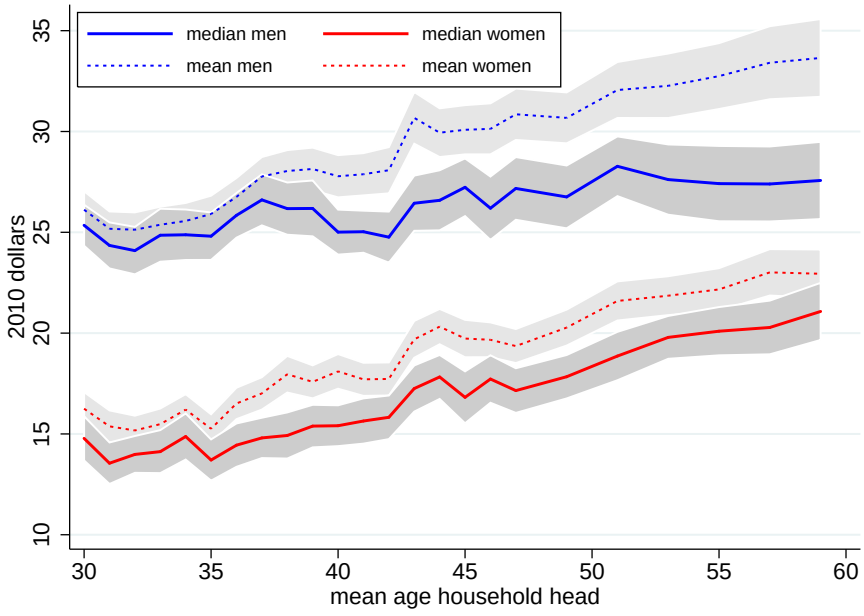
[57]Despite this, the PSID officially treats men's earnings series as consistent over time. For female earnings two different series are provided (one prior to 1993 and one after). I combine the two into a single female series.

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), 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 model hours.

Age of the youngest child is classified in four groups to reflect the way stochastic fertility is modeled in section 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 ($N_t = 4$).

Figure A.1: Evolution of wages against 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:* PSID 1980-2009.

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, farmers, 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.

Gender wage gap and female participation Suppose P^*_{2it} is women's latent participation choice. One observes $P_{2it} = 1$ if $P^*_{2it} > 0$ and $P_{2it} = 0$ if $P^*_{2it} \leq 0$, where P_{2it} is a binary variable for whether the female spouse participates in the labor market. Consistent with the structural model women's latent choice depends on (i) their husband's wage w_{1it}, (ii) their own offered wage w^*_{2it}, (iii) the presence and age of their youngest child N_{it}, and (iv) their assets A_{it}. As I do not observe assets consistently, I specify

$$P^*_{2it} = \gamma_0 + \gamma_1 w_{1it} + \gamma_2 w^*_{2it} + \sum_n \delta_n \mathbf{1}[N_{it} = n] + u_{it}. \tag{A.1}$$

A problem in (A.1) is that w^*_{2it} is a latent variable itself. Consistent with the model, I write w^*_{2it} as a function of the husband's wage (to capture possible assortative patterns between spouses) and age, namely

$$w^*_{2it} = \beta_0 + \beta_1 w_{1it} + \sum_t \beta_t \mathbf{1}[age_{it} = t] + \epsilon_{it}. \tag{A.2}$$

In practice one only observes $w_{2it} = w^*_{2it}$ if $P^*_{2it} > 0$. Replacing w^*_{2it} with (A.2) in (A.1), the latent participation equation becomes

$$\begin{aligned} P^*_{2it} &= (\gamma_0 + \gamma_2\beta_0) + (\gamma_1 + \gamma_2\beta_1) w_{1it} + \gamma_2 \sum_t \beta_t \mathbf{1}[age_{it} = t] + \sum_n \delta_n \mathbf{1}[N_{it} = n] + (\gamma_2 \epsilon_{it} + u_{it}) \\ P^*_{2it} &= \boldsymbol{\alpha}'_{it} \mathbf{x}_{it} + v_{it}, \end{aligned} \tag{A.3}$$

where $\mathbf{x}_{it} = (1, w_{1it}, age_{it}, N_{it})'$ and $v_{it} = \gamma_2 \epsilon_{it} + u_{it}$. From this it becomes clear that v_{it} is correlated with ϵ_{it} and OLS on (A.2) is inconsistent. However, following Heckman (1979), the *observed* wage equation can be written

$$w_{2it} = \beta_0 + \beta_1 w_{1it} + \sum_t \beta_t \mathbf{1}[age_{it} = t] + \sigma_{v\epsilon} \lambda(\widehat{\boldsymbol{\alpha}}'_{it} \mathbf{x}_{it}) + \varepsilon_{it}$$

where $\sigma_{v\epsilon}$ is the correlation between v_{it} and ϵ_{it}, λ is the inverse Mills ratio, and their product accounts for women's participation selection. Consistent estimates of β_0, β_1 and β_t can be obtained by the typical two-step approach. This then allows me to predict offered wages for participating and non-participating women. Age of the youngest child N_{it} provides an exclusion restriction in (A.3) so identification does not come through non-linearities only. This procedure can accommodate additional covariates such as education or *number* of children in the household.

B Model: Private Consumption

This appendix extends the model in section 3 to allow spouses to consume public *as well as* private consumption goods. In this case, individual j has preferences $\widetilde{U}_j$ given by

$$\widetilde{U}_j\big(Q, q_j, l_j; \mathbf{z}_j\big)$$

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)-(7) after replacing individual preferences with $\widetilde{U}_j$ and the budget constraint (3) with

$$A_t + \sum_{j=1}^{2} w_{jt} h_{jt} = K_t + p_t \sum_{j=1}^{2} q_{jt} + CC_t\big(h_{2t}, N_t\big) + \frac{A_{t+1}}{1+r}.$$

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

$$\widetilde{U}_j\big(Q_t, q_{jt}, l_{jt}; \mathbf{z}_{jt}\big) = \frac{1}{1-\gamma}\left(\alpha_j\big(Q_t/s(N_t)\big)^{1-\gamma} + (1-\alpha_j) q_{jt}^{1-\gamma}\right) \times \exp\left(g_j(l_{jt}; \mathbf{z}_{jt})\right)$$

which is a straightforward extension of (8). The leisure sub-utility $g_j(\cdot)$ remains unchanged. Here α_j serves as the utility weight of public consumption which may further 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.

C Estimation: Exogenous Elements

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 $\widetilde{\mathbf{W}} = (\widetilde{\mathbf{W}}'_{1980}, \widetilde{\mathbf{W}}'_{1981}, \ldots, \widetilde{\mathbf{W}}'_{2009})'$ where

$$\widetilde{\mathbf{W}}_t = \begin{pmatrix} \mathbb{E}[(\Delta \ln w_{1t})^2] \\ \mathbb{E}[\Delta \ln w_{1t} \Delta \ln w_{1t+1}] \\ \mathbb{E}[(\Delta \ln w_{2t})^2] \\ \mathbb{E}[\Delta \ln w_{2t} \Delta \ln w_{2t+1}] \\ \mathbb{E}[\Delta \ln w_{1t} \Delta \ln w_{2t}] \\ \mathbb{E}[\Delta \ln w_{1t} \Delta \ln w_{2t+1}] \\ \mathbb{E}[\Delta \ln w_{2t} \Delta \ln w_{1t+1}] \end{pmatrix}, \quad t \in [1980, 2009]$$

I ignore any auto-covariances of order higher than 1. In the PSID these are mostly insignificantly different from 0 thus motivating the specific parametrization (10).

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 given by

$$\hat{\boldsymbol{\vartheta}}_{\text{MD}} = \arg\min_{\boldsymbol{\vartheta}} (\widetilde{\mathbf{W}} - \mathbf{W}(\boldsymbol{\vartheta}))' I (\widetilde{\mathbf{W}} - \mathbf{W}(\boldsymbol{\vartheta}))$$

where I is the identity matrix. The estimates of $\boldsymbol{\vartheta}$ appear in table C.1 alongside their standard errors; to calculate those I adopt the block bootstrap with 500 replications.

The point-estimates in table C.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 well as the refined ones, appear graphically in figure C.1.

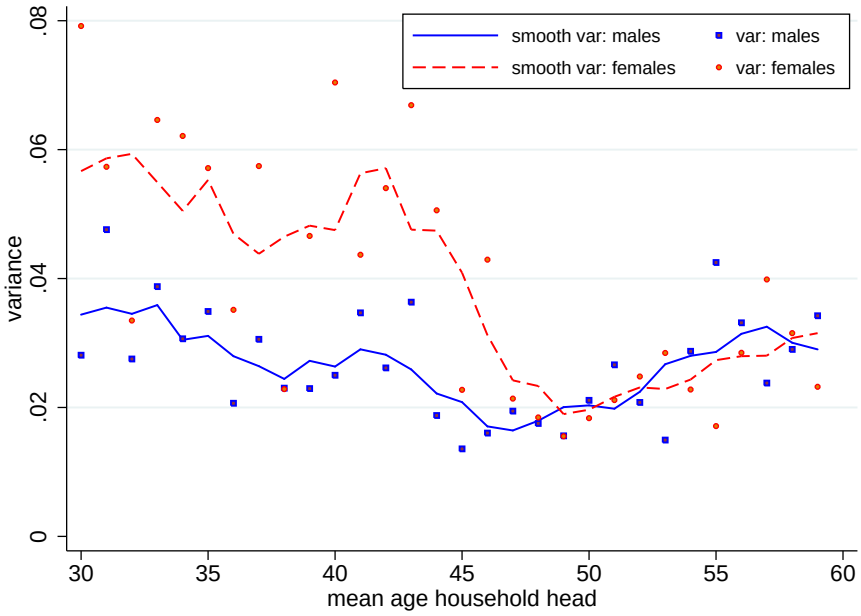
Child care costs I calibrate *cchrate* at a constant \$6.59 (expressed in 2010 dollars) throughout the 1980-2009 period; see section 5.1 in the main text for details. Whenever this rate is below the real federal minimum wage, I update *cchrate* to reflect this. Essentially the hourly wage of child care 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 estimating the proportion 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 C.2 reports the hourly rate of child care over time (accounting for the adjustment for

the federal minimum wage) as well as the estimated probabilities of positive child care costs in 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 C.2).

Initialization of the Pareto weight To project lifetime earnings if spouses get divorce, I first pool earnings of divorced persons 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 (15) in the main text and the results appear in table C.3 here. 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 (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 5.3 in the main text for details). Table C.4 reports how the derived intra-family bargaining power of the male spouse correlates with a number of characteristics of each individual.

Figure C.1: 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:* PSID 1980-2009 and own calculations.

Table C.1: Second moments of wage shocks

Year	Mean age head	I. Permanent shocks			II. Transitory shocks		
		Men	Women	Covari-ance	Men	Women	Covari-ance
1980	30	0.0281	0.0792	0.0081	0.0214	0.0546	0.0039
		(0.0080)	(0.0226)	(0.0080)	(0.0058)	(0.0164)	(0.0052)
1981	31	0.0476	0.0573	-0.0009	0.0253	0.0710	0.0060
		(0.0107)	(0.0200)	(0.0079)	(0.0055)	(0.0211)	(0.0051)
1982	32	0.0275	0.0335	0.0097	0.0226	0.0697	0.0037
		(0.0079)	(0.0211)	(0.0073)	(0.0046)	(0.0228)	(0.0063)
1983	33	0.0388	0.0646	0.0093	0.0176	0.0602	-0.0019
		(0.0071)	(0.0194)	(0.0075)	(0.0045)	(0.0155)	(0.0047)
1984	34	0.0307	0.0621	0.0062	0.0254	0.0380	-0.0010
		(0.0070)	(0.0146)	(0.0064)	(0.0052)	(0.0127)	(0.0049)
1985	35	0.0349	0.0571	0.0121	0.0239	0.0725	-0.0045
		(0.0062)	(0.0169)	(0.0081)	(0.0051)	(0.0161)	(0.0049)
1986	36	0.0207	0.0351	0.0102	0.0305	0.0737	-0.0032
		(0.0061)	(0.0146)	(0.0057)	(0.0059)	(0.0155)	(0.0051)
1987	37	0.0306	0.0575	0.0172	0.0253	0.0789	-0.0017
		(0.0065)	(0.0249)	(0.0059)	(0.0048)	(0.0251)	(0.0038)
1988	38	0.0230	0.0229	0.0108	0.0290	0.0638	-0.0053
		(0.0062)	(0.0153)	(0.0060)	(0.0057)	(0.0123)	(0.0045)
1989	39	0.0229	0.0466	0.0099	0.0337	0.0199	0.0027
		(0.0060)	(0.0122)	(0.0050)	(0.0073)	(0.0071)	(0.0044)
1990	40	0.0250	0.0704	0.0043	0.0178	0.0408	0.0055
		(0.0064)	(0.0177)	(0.0058)	(0.0049)	(0.0089)	(0.0041)
1991	41	0.0347	0.0437	0.0084	0.0259	0.0466	-0.0010
		(0.0059)	(0.0116)	(0.0064)	(0.0060)	(0.0092)	(0.0039)
1992	42	0.0261	0.0540	0.0165	0.0469	0.0343	-0.0118
		(0.0097)	(0.0147)	(0.0070)	(0.0101)	(0.0097)	(0.0053)
1993	43	0.0363	0.0669	0.0243	0.0673	0.0566	-0.0032
		(0.0123)	(0.0143)	(0.0072)	(0.0116)	(0.0123)	(0.0059)
1994	44	0.0188	0.0506	0.0173	0.0595	0.0885	0.0035
		(0.0085)	(0.0175)	(0.0069)	(0.0125)	(0.0165)	(0.0055)
1995	45	0.0136	0.0227	0.0071	0.0337	0.0514	0.0011
		(0.0077)	(0.0135)	(0.0063)	(0.0057)	(0.0099)	(0.0044)
1996	46	0.0160	0.0429	0.0128	0.0106	0.0589	-0.0017
		(0.0055)	(0.0128)	(0.0071)	(0.0032)	(0.0139)	(0.0058)
1997	47	0.0194	0.0214	-0.0004	0.0413	0.0628	0.0083
		(0.0059)	(0.0069)	(0.0040)	(0.0101)	(0.0118)	(0.0051)
1999	49	0.0156	0.0155	0.0024	0.0436	0.0494	-0.0048
		(0.0046)	(0.0061)	(0.0028)	(0.0090)	(0.0092)	(0.0042)
2001	51	0.0266	0.0212	0.0057	0.0522	0.0359	0.0112
		(0.0083)	(0.0062)	(0.0033)	(0.0102)	(0.0065)	(0.0063)
2003	53	0.0149	0.0285	0.0034	0.0571	0.0415	-0.0049
		(0.0052)	(0.0087)	(0.0036)	(0.0126)	(0.0091)	(0.0059)
2005	55	0.0425	0.0171	0.0019	0.0621	0.0411	-0.0003
		(0.0103)	(0.0045)	(0.0042)	(0.0155)	(0.0067)	(0.0051)
2007	57	0.0238	0.0399	0.0032	0.0430	0.0390	0.0054
		(0.0075)	(0.0078)	(0.0038)	(0.0129)	(0.0105)	(0.0059)
2009	59	0.0342	0.0232	0.0008	0.0431	0.0304	-0.0004
		(0.0091)	(0.0069)	(0.0053)	(0.0098)	(0.0123)	(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:* PSID 1980-2009 and own calculations.

Table C.2: Child care: price and take-up 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 a family reports 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 (1988). *In even years after 1997, when the PSID did not collect data, I use the mid-point of probabilities in adjacent years.

Table C.3: Earnings regressions: male and female divorcees

	Dependent variable: Annual Earnings			
	I. Male divorcees		II. Female divorcees	
Regressors:	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. The number of observations reflects the number of male/female divorcees-year observations in 1980-1989.

Table C.4: Married men's initialized bargaining power: correlation with spousal attributes

	Dependent variable: men's bargaining power	
Regressors:	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 for 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.

D Estimation: Model Fit and Overidentification

The following two tables present the full set of targeted moments used in the structural estimation (see section 5.2 in main text). The tables report the values of the moments in the data as well as their counterparts from the model simulations.

In addition, figure D.1 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 estimation, are aligned along or around the 45-degree line (this line indicates a prefect match).

Table D.1: Targeted moments: proportions men

	fertility state 1		fertility state 2	
	Data	Model	Data	Model
Mean age head: 30-39				
'Low middle' household work	0.732	0.704	0.734	0.713
Mean age head: 40-49				
'Low middle' household work	0.680	0.681	0.784	0.749
Mean age head: 50-59				
'Low middle' household work	0.705	0.705	0.679	0.737
	fertility state 3		fertility state 4	
	Data	Model	Data	Model
Mean age head: 30-39				
'Low middle' household work	0.686	0.716	0.620	0.585
Mean age head: 40-49				
'Low middle' household work	0.742	0.694	0.693	0.752
Mean age head: 50-59				
'Low middle' household work	0.673	0.728	0.786	0.738

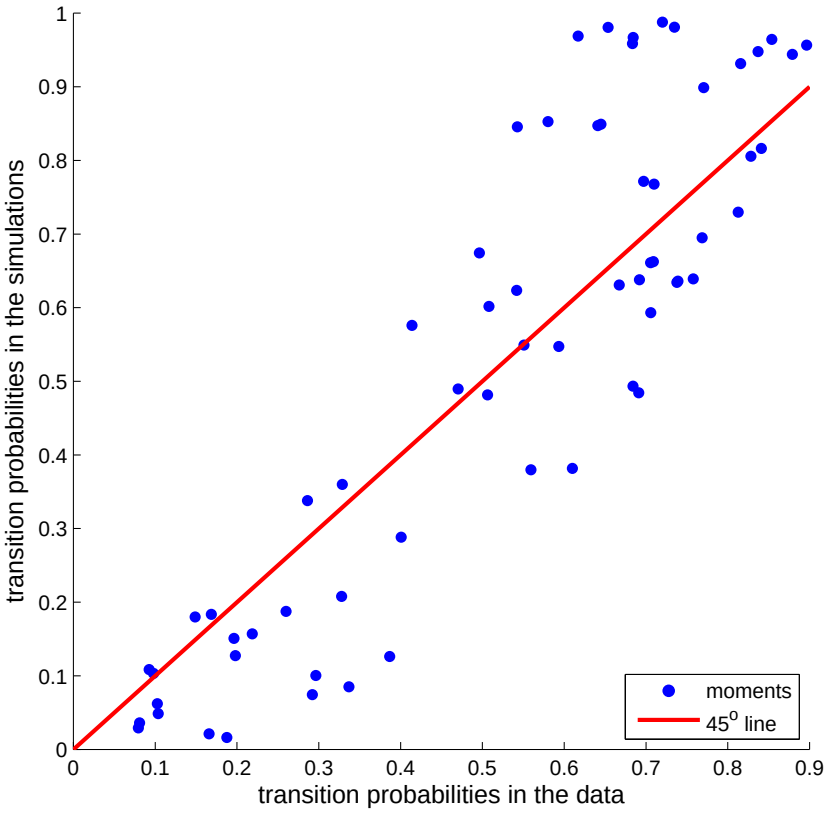
Notes: This table reports the values of men's targeted moments in the data and the model simulations. These moments are proportions of married men engaging in 'low-middle' household work by men's mean age and family composition. For the definition of 'low middle' household work refer to table 2 in the main text.

Table D.2: Targeted moments: proportions women

	fertility state 1		fertility state 2		fertility state 3		fertility state 4	
	Data	Model	Data	Model	Data	Model	Data	Model
Mean age head: 30-39								
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
Mean age head: 40-49								
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
Mean age head: 50-59								
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 2 in the main text.

Figure D.1: 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 Prob$\left[\text{spouse}_j \text{ time allocation}_t \mid \text{spouse}_k \text{ time allocation}_{t-s}\right]$ 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: spouse_1 time allocation_t = {'low middle' household work}; for women: spouse_2 time allocation_t = {FT market work, PT market work, 'maximum' household work}.