

Turon, H el ene

Working Paper

Home Production of Childcare and Labour Supply Decisions in a Collective Household Model

IZA Discussion Papers, No. 12148

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Turon, H el ene (2019) : Home Production of Childcare and Labour Supply Decisions in a Collective Household Model, IZA Discussion Papers, No. 12148, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/196646>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor d urfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie d urfen die Dokumente nicht f ur  offentliche oder kommerzielle Zwecke vervielf altigen,  offentlich ausstellen,  offentlich zug anglich machen, vertreiben oder anderweitig nutzen.

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

Terms of use:

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

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

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

DISCUSSION PAPER SERIES

IZA DP No. 12148

**Home Production of Childcare and
Labour Supply Decisions in a Collective
Household Model**

Hélène Turon

FEBRUARY 2019

DISCUSSION PAPER SERIES

IZA DP No. 12148

Home Production of Childcare and Labour Supply Decisions in a Collective Household Model

Hélène Turon

University of Bristol and IZA

FEBRUARY 2019

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

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

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

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Home Production of Childcare and Labour Supply Decisions in a Collective Household Model*

This paper proposes a dynamic structural model of labour market and childcare choices for couples within a collective model of decision making. We formalise explicitly the need for childcare as a function of the age structure of the children population in the household then examine the determinants of the decision to supply labour. The fraction of home-produced childcare to household childcare needs is considered to be a public good within the household, for which preferences are heterogeneous across households. An important feature of our framework, which introduces one of the dynamic dimensions of the decision, is that we take into account the implications of today's labour supply decision on future wage growth and future bargaining power. The decision to leave (partially or not) the labour market is often taken within a couple but, in the event of divorce, the impact of this decision may not be borne by both parties equally, which may render the initial decision inefficient. Using data from the BHPS, we then present a structural estimation of our model to quantify these various components of the choice of home childcare vs. labour supply. We are able to quantify each household's sensitivity to potential childcare policies and find that a large part of the dispersion in these responses comes from households' valuation of home-produced childcare.

JEL Classification: J12, J13, J22, J31, J38

Keywords: household, labour supply, collective model, childcare, commitment

Corresponding author:

Helene Turon
Department of Economics
University of Bristol
Priory Road Complex
Bristol BS8 1TU
United Kingdom

E-mail: helene.turon-lacARRIERU@bristol.ac.uk

* I would like to thank Pierre-André Chiappori, Grégory Jolivet, John Knowles, Albert Marcet, Ramon Marimon, Chris Muris and seminar and conference participants in Journées Louis-André Gérard Varet, Royal Economic Society conference, Families and the Macroeconomy conference, Ecole CNRS TEPP Aussois, Kent Workshop in Labour Economics, VU Amsterdam, UAB Barcelona, European University Institute, Université Aix-Marseille, Trinity College Dublin, ESRI Dublin and colleagues in Bristol for very useful comments and discussions.

1 Introduction

Subsidies to childcare and reforms of the amount and eligibility to these subsidies are pervasive in developed countries.¹ Both the rationale for the existence of these subsidies and their effectiveness are recurring topics of policy debate.² One ground to subsidise childcare is to encourage maternal labour supply. The price of childcare has been used as the main policy tool to influence this choice.³ However, the size of the response of female labour supply to existing and potential childcare subsidies are not well understood. This paper proposes a structural estimation of the costs and benefits of using market-provided childcare and offers an explanation for the heterogeneity in responses to childcare reforms across households and an estimate of the aggregate population response to counter-factual policies. In our (British) data, a striking observation is that a *monetary* cost-benefit analysis of the choice of labour supply vs. home childcare does not explain the data well. We argue that heterogeneous preferences regarding home-produced childcare play an important role in this choice.

The main ingredients of our framework are the following. We formalise explicitly the need for childcare as a function of the age structure of the children population in the household. The fraction of home-produced childcare to household childcare needs is considered to be a public good in the household, for which preferences are heterogeneous across households. We model the labour supply decision as a discrete choice between non-participation, part-time work and full-time work and take into account the cost of working less than full-time on future labour market value. Household choices are made collectively and maximise a weighted sum of the lifetime utilities of the two spouses.⁴ Bargaining weights are adjusted when one spouse raises a credible threat to leave the marriage. These dynamics of the bargaining weights are modelled within a limited commitment model. They are endogenous in our framework because labour market choices affect future wages, which affect divorce values in an asymmetric manner. This in turn affects choices of the forward-looking household, which, depending on the current weights, may value positively or negatively this future change in bargaining weights. Indeed, the decision to leave (partially or not) the labour market is often taken within a couple but, in the event of divorce, settlement rules are such that the impact of this decision is not borne by both parties equally. Divorce regulations thus also play a role in the labour supply choice.

Using data from the British Household Panel Survey, we present a structural estimation of our model to quantify these various components of the choice of labour supply vs. home childcare. For example, comparing maternal full-time work combined with no home childcare to non-participation by the mother combined with full-time home childcare, we estimate the difference in current household budget through loss of current labour income and market childcare costs, in current household utility through its taste for home-produced childcare, in future household budget through the loss of potential wage growth and

¹Examples include the Childcare Voucher Scheme in the UK (introduced in 1989, phased out in 2018), the Equal Opportunities Childcare Programme in Ireland (2000-2006), the parental education allowance in France (reformed in 1994 and 2004), the reforms aimed at enhancing the educational quality of childcare in Sweden in 2008, and the 2008 German reform (Kinderfrderungsgesetz) establishing a right to a day care spot for every child from the age of one.

²See for example Domeij and Klein (2012) on the optimality of subsidising childcare.

³See for example Gustafsson and Stafford (1992), Baker et al. (2008), Blau and Robins (1988) and Ribar (1995).

⁴We only study two-headed households.

in future bargaining position within the couple. We find that, whilst much of the emphasis of the public debate is on the first component, i.e. the current net income loss of lower labour market participation, the next two components relating to the taste for home-produced childcare and the expected impact on future labour earnings are each of similar magnitude. On the other hand, the impact of non-participation on future bargaining weights within marriage is relatively modest and does not play a large role in the current household decision. By contrast, in the event of divorce, consequences of lower market participation can be sizeable for the spouse who has provided the home-made childcare since the loss of labour market value incurred is not taken into account in the divorce settlement calculation.

Using our results and the structure of our framework, we are then able to compute differences in lifetime utility for all households in our sample between alternative labour supply choices. This allows us to gauge the fractions of households in our sample who are at the margin, i.e. likely to alter their labour supply choice in the event of an decrease in the price of childcare, of households who will supply joint labour even if the price of childcare was much increased, and of households who will carry out domestic childcare even if the price of childcare was much decreased.⁵ This allows us to quantify the fraction of households who are likely to respond to a change in the price of childcare and provides us with a better understanding of mothers' labour market elasticity to childcare subsidies.

The related literature is plentiful. Ever since the seminal paper by Chiappori (1992), collective models of household decisions have been used to understand household consumption patterns and to estimate the sharing rule, e.g. Browning et al. (2006), sometimes in the presence of children, as in Blundell et al. (2005), sometimes in the presence of a public good, as in Donni (2009). A recent stream of papers uses the revealed preference approach to identify the model parameters with minimal use of parametric assumptions, led by the work of Cherchye et al. (2010), Cherchye et al. (2011) and Cherchye et al. (2012). Much of this literature is set in the static context, but a recent branch has introduced a dynamic dimension to household decisions. This raises the question of the type of commitment prevailing in households. When divorce legislation is such that divorce can be triggered by one of the parties alone (unilateral divorce as opposed to divorce by mutual consent, see Mazzocco (2007)) an adequate description of bargaining within couples is the limited commitment model as described in Marcet and Marimon (2017), Ligon et al. (2002), Chiappori and Mazzocco (2017) and Lise and Yamada (2019). Within this model, bargaining weights vary and this in turn affects household choices, see Voena (2015), Bronson (2014) and Stevenson (2007). Post-divorce outcomes also influence these choices directly, by affecting the incentives to specialisation within marriage (Foerster (2018)) and married women's decisions to work part-time (García-Morán and Kuehn (2018)). This paper contributes to this literature by incorporating the need for childcare into the household decision problem and considering the choice of home childcare versus labour supply within the dynamic perspective of its impact on future labour earnings and future household bargaining. A novel feature is that the non-work time of the spouse providing childcare is valued by both spouses but cannot be bargained over through current or future transfers between spouses.

⁵Specific figures can be found in section 5.4.

Our results contribute to various policy-related debates. Papers connected to the questions we examine are Adda et al. (2017) on the career costs of children, Guner et al. (2018) on household labour supply and childcare subsidies, Bick (2016) on childcare enrolment rate and mother’s labour force participation, Blundell et al. (2016) on the impact of the tax and benefits system on mothers’ labour supply, and Chiappori et al. (2002) on the impact of marriage market and divorce legislation on household labour supply. We provide a structural understanding of the labour supply elasticity of a population to childcare subsidies and of the inefficiencies in household labour supply choices in the presence of young children arising from the design of divorce regulations. The specificity of our approach lies in the central role played by households’ taste for home-produced childcare and its interactions with our two policy tools of interest, namely the price of childcare and divorce regulations.

The rest of the paper is organised as follows. The next section outlines our model. Section 3 presents descriptive statistics of our dataset and estimates of the auxiliary regressions used later on. In section 4 we detail our estimation procedure before presenting the results and our policy analysis in section 5. Finally, section 6 concludes.

2 The model

We set out a partial equilibrium model of labour supply, childcare and consumption choices of individuals in two-partner households and of their decision to divorce or remain in the marriage.⁶ Our focus is on the decision to carry out childcare at home or to buy childcare services to free up time for labour market work and/or leisure, given a population of children in the household and the (potential) wages of the two partners. Individuals value consumption, leisure and home-produced childcare in a manner described in section 2.2. In order to keep the model reasonably tractable and parsimonious, we keep the fertility hazard as exogenous. Couple formation, similarly, is kept out of the picture as we will follow couples already in existence at the beginning of the sample. We model couple decision making within a limited commitment model described in section 2.5, which also predicts endogenous divorce.

The environment is one of dynamic choice set in discrete time, where the (common) discount factor is denoted β . The dynamic dimension of choices is twofold and hinges on the fact that participation in the labour market and the intensity of this participation have an impact (in expectation) on future earnings (see section 2.3) and future bargaining power (see section 2.5). We rule out saving and borrowing behaviour as well as the concept of home childcare being an investment in the ‘quality’ of the child in future periods. The unit time period is five years. This is mainly to ease the computational burden but also because we view the choice between domestic and market-bought childcare as being made rather infrequently and kept to for some years. We follow working-age couples and assume for simplicity that both partners have the same age, a . The age structure of the children population is denoted κ and will be defined precisely in section 2.1. We will only model time use within the standard working time in a week, i.e. 40 hours, as this is the time where

⁶We make no distinction between married and cohabiting couples.

the alternative between home-produced and market-bought childcare seems the most relevant. We ignore the possibility of working staggered shifts to combine work and home childcare, that is, if both parents work full-time all childcare needs (within the 40 working hours) must be bought on the market.⁷

In this framework, a household member g has three possible uses of his/her time, namely work, leisure, L_g and childcare dom_g . We assume that labour market participation lm_g can only take three forms: non-participation, $lm_g = 0$, part-time work $lm_g = 1$ or full-time work, $lm_g = 2$. The individual time constraints are thus the following:

$$40 = 20 \cdot lm_g + L_g + dom_g, \quad g = m, f. \quad (1)$$

Given that we have no savings in this framework, the only state variables are the age of two partners, a , the age composition of the children population, κ , and the wages (or potential wages, which we keep track of for individuals out of the labour force) of the male and the female partners, w_m and w_f and the Pareto weights of the two partners. The control variables are both partners' labour market choices, lm_g , hours of domestic childcare, dom_g , hours of leisure, L_g and private consumption, C_g . Labour market choices are discrete ; the other choices are continuous.

The household faces three types of constraints: a joint budget constraint, individual time constraints, and a joint constraint in childcare need imposed by the age structure of the children population in the household. Two additional constraints are that both spouses must be better off within the couple than in the divorced state. If neither of them is, divorce ensues. If at least one of them is better off within the marriage, an adjustment of the Pareto weight may allow the marriage to continue, as described in section 2.5. These dynamics of the Pareto weights is an interesting outcome of our framework, both for its own sake and for its role in the labour supply decision.

The dynamics of κ , the age structure of the children population in the household, is driven by the (deterministic) ageing of existing children and the stochastic exogenous fertility process. The dynamics of wages are stochastic, driven by independent shocks and by (endogenous) labour market choices.

We will now describe each component of the model in more detail and present the formal equations in section 2.6.

2.1 Childcare need

The children population of the household is described by its age composition and size. Because we are interested in household decisions in terms of domestic production versus market provision of childcare services, we use age categories that determine childcare needs. A feature of this paper is that we include within the household constraints a need for childcare which is explicitly formulated in terms of the age composition of the children population.

⁷This is counterfactual for some subset of the population who work staggered working days to accommodate childcare needs. It however comes at the cost of time spent together out of work. Hamermesh (2002) estimates that couples put some value on this togetherness and synchronise their schedules. The analysis of this tradeoff is beyond the scope of this paper and we ignore it altogether with the assumption that all labour market work has to be carried out within the same 40 hours. We also ignore informal childcare carried out by e.g. grandparents.

There are 4 age categories, characterised by the amount of childcare that they require: under 5 years old, 5-11, 11-16 and over 16, which correspond, in the UK, to the ages of children who are pre-school, in primary school, in secondary school and above the legal school-leaving age (within our sample). As our focus is on time use decisions within normal working hours, we will define childcare needs only with respect to the subset of these needs that fall within working hours.

Without much loss of generality, we model families as having up to 3 children as very few (1.3%) families in our sample have more than three children.⁸ The vector $\kappa = (\kappa_k)_{k=1..3}$ represents the age structure (at next birthday) of the children in the household where $\kappa_k = 0$ if there are less than k children (i.e. the k -th child is unborn). The dynamics of κ are driven by the ageing of existing children and the birth of news ones.

Just as adults, children age by 5 years every period. Fertility events are exogenous and depend on the age structure of the existing children in the household, κ , and on the age of the parents, a . So for example, a family with two children aged 4.5 and 6.2 is represented by $\kappa_t = (7, 5, 0)$. In the following year, we will have either $\kappa_{t+1} = (8, 6, 1)$ in the event of a new birth or $\kappa_{t+1} = (8, 6, 0)$ otherwise.

These dynamics are represented by the matrix A^K of transition probabilities between different values of κ . This matrix only depends on a the age of the two partners in the household.

We denote the number of hours of childcare needed for the aforementioned four age categories as $\{\gamma_j\}_{j=1..4}$. Intuitively, $\gamma_{j+1} \leq \gamma_j$ for $j = 1..3$ since older children require less childcare than younger ones. These γ parameters will be calibrated in the estimation section below to reflect evidence on childcare use for children of different age categories. Note that these hours needed are all nested in one another in that the γ_3 hours weekly childcare that a child aged 11-16 needs are a subset of the γ_2 hours of weekly childcare that a primary school child needs in a timetabling sense. We denote CC_k the number of hours of childcare needed by the child indexed k in the household. The total time needed for childcare if this childcare is carried out by any (or both) adult(s) from the household is: $\max\{CC_k\}_{k=1..3}$. This expression reflects the fact that, when devoting one hour of time to childcare, the adult may look after one to three children.⁹

Adults in the household may decide to carry out some or all of the childcare needed by spending time dom_g ($g = m, f$) on this activity. We rule out¹⁰ the possibility that childcare times performed by parents within the 40 weekly ‘working’ hours overlap so that the total time devoted to the home production of childcare within the household is $dom_m + dom_f$. By “working hours” we mean hours that could be used for labour market work. Since childcare needs must be covered either by home production or market-bought childcare services, the household spending on childcare services is:

$$p_{CC} \cdot \sum_{k=1}^3 \max\{CC_k - (dom_m + dom_f), 0\}, \quad (2)$$

where p_{CC} is the unit (hour/child) price of childcare services.¹¹ The expression (2) reflects the fact that, contrasting with home-produced childcare, one hour of childcare services needs to be bought for each child

⁸In the rare instances where households have more than three children, the ages of the youngest three children are used.

⁹We abstract from any potential difference in childcare quality depending on the number of children cared for by an adult in the home –or by market-provided childcare.

¹⁰It is in fact a dominant strategy in our framework.

¹¹One could argue that this price should vary with parents’ education levels to capture the fact that better-off households

needing childcare. This difference between the time required for home-produced childcare and the number of hours of childcare services needed to be bought for a given family structure means that the relative ‘price’ of one hour devoted to childcare or to work depends on the age structure of the children population in the household κ .

2.2 Taste for home-produced childcare

One specificity of our approach is that we consider that home-produced childcare D enters individuals’ preferences as well as consumption C and leisure L . Each spouse’s instantaneous utility is denoted $\mathcal{U}_g(C_g, L_g, D)$ for $g = m, f$. Individuals do not derive utility from carrying out childcare *per se*, but from the fact that a higher proportion of the household childcare need is being carried out within the home, by whatever parent as opposed to a market provider (nursery, childminder). This proportion D is formally defined as follows:

$$D = \frac{dom_m + \lambda \cdot dom_f}{\max_k \{CC_k\}}, \quad (3)$$

D equals 1 for households without children and households who opt to carry out all their childcare domestically. This fraction D of total childcare time need carried out within the home is a **public good** within the household.

We allow households to value the domestic times of the two parents differently via the factor λ . Possible reasons why the domestic times of the parents may not be valued equally are social custom on ‘traditional’ role of the mother as care-giver, biological comparative advantage in breastfeeding or inherited preference for parent giving care in the early months. We will discuss the role of λ in more depth in section 5.5.

The marginal rate of substitution between D and consumption is one of the driving forces of the choice of childcare mode. A contribution of this paper is our attempt to quantify this ‘taste’ for home-produced childcare and its heterogeneity across households. Indeed we claim that this is key to an understanding of the response of the population of household to changes in the price of childcare services. Our estimate of the distribution of the sensitivity of households to the unit price of childcare is presented in section 5.4.

The relative ‘price’ of D depends on both static and dynamic considerations. In the current period, the amount of household consumption lost or gained with a marginal increase in D depends on the number of children needing childcare at the margin, the overall household childcare need, the unit price of childcare and the wage of the parent increasing his/her domestic time. In future periods, other components of the trade-off include the loss of labour market value of the partner working less as well as the impact of this on Pareto weights (see section 5.4). We assume away any dynamic consequence of children being looked after by a parent versus a market provider in terms of children’s well-being or outcomes in the long run.

Components of this relative price will be discussed at length below, but we retain from this that it will depend on κ , (w_m, w_f) and the Pareto weights of the two partners.

are likely to use childcare services at higher prices on average than less well-off households. We have ruled this out, though it could easily be included by considering that this price is a fraction of the household potential income. Similarly, we have ignored potential discounts on the price of childcare services when more than one child in the family is using the facility.

2.3 Wage processes

In order to preserve the parsimony of the model and to keep the computational burden manageable, our modelisation of the wage process is very stylised. We model wage dynamics as a first-order Markov process. Labour market choice can only take three values: non-participation ($lm = 0$), part-time ($lm = 1$) or full-time ($lm = 2$), corresponding to 0, 20 and 40 weekly hours of work respectively. We characterise all wages by their quintiles in the overall wage distribution and assume that transition probabilities between quintiles p and p' over the next period depend only the current quintile p and labour market choice, lm . We denote these probabilities $a_{lm}^W(p, p')$. Each individual, whether earning an actual wage or not, will be assigned a 'market value' or potential wage, which will carry on evolving over time according to the above process even when the individual is out of the labour force.

As mentioned above, labour supply choices are assumed to be restricted to three states: full-time employment, part-time employment and non-participation. Part-time and full-time employment may differ among three dimensions: the amount of time they take, the hourly wage and the rate of wage growth. Indeed, it is a well-documented fact that part-time jobs tend to pay less and offer less scope for promotion than their full-time equivalent, even when accounting for selection (see Connolly and Gregory (2009), Harkness (1996) and Manning and Petrongolo (2008)). An important feature of these differences is that they are not uniform across the wage distribution: the expected part-time penalty in terms of wage growth is not the same at low and high quintiles of the distribution. Evidence of this in our data will be shown in section 3.2.

Our assumptions are the following: part-time work does not carry an instantaneous penalty, i.e. there is no hourly wage drop in the current period when taking a part-time as opposed to a full-time job. Over time, however, part-time work (and non-participation) deteriorate the labour market value of an individual relative to full-time work, in expectation. New labour market entrants are given a market value distribution that depends solely on education. From then on, wages experience growth spurts or drops which occur at Poisson rates $a_{lm}^g(p, p')$. These wage shocks are assumed to be independent between spouses and independent of the household fertility shock.

Admittedly this is a crude representation of income dynamics, which have been documented to be better represented by richer specifications (see for example Meghir and Pistaferri (2004) and Guvenen et al. (2015)). We think nevertheless that our simple process allows us to capture the mechanisms of interest in this paper, one of them being that any parent's decision to take time off the labour market to produce home childcare does have consequences on his or her wage progression. Besides, this expected cost depends on the initial wage of the individual.

2.4 Divorce and retirement

The main objective of this model is to analyse labour market and childcare choices of individuals living as a couple. However, because we also consider the possibility of divorce and the impact that this possibility has on household choices, we need to model preferences and choices of divorced individuals. We follow

trajectories of couples whose initial status is a marriage or partnership. Thus we do not model household formation.

2.4.1 Divorce

There is no straightforward way to treat the allocation of time and money resources upon divorce. Our modelling aim in this respect is twofold. First, a concern for parsimony leads us to adopt a very stylised picture of post-divorce outcomes. Second, we wish to capture, even if in a simplified way, how childcare need and individual potential wages are passed on to each individual post divorce. In the related literature, Mazzocco (2007) and Bronson (2014) assume the divorced state to be equal to the state of being single while Voena (2015) allows for the remarriage probability of divorcees to differ from that of singles and for a division of assets depending on divorce laws. Since our framework focuses on childcare we need to specify divorce outcomes with respect to custody and child support. We use the government regulations (UK Government (2017)) to set our assumptions on divorce values. As noted by Voena (2015), marriage contracts are rare and difficult to enforce so it seems reasonable to assume that spouses expect to follow government regulations in the event of divorce.¹²

We assume that the female partner receives custody of the children population and the attached childcare needs. She also receives child support y from the male partner which depends on the age structure of the children population and on his labour market value at the time of divorce:

$$y = y(\kappa, w_m^D) \quad (4)$$

Note that the amount of child support does not depend on the female partner's labour market value or past labour market choices.

Her subsequent choice as to whether to buy childcare services or to produce them domestically is now an individual choice, but she inherits the household taste for home childcare. The male partner inherits this taste too, his utility will therefore be affected by her labour market choices for as long as there is a childcare need for their children. He however does not participate in that choice. Since the child support transfer is set according to the male partner's wage rate at the time of divorce, the dynamics of the male labour market value cease to matter thereafter in the female decision problem.

We also assume that remarriage and fertility hazards post-divorce are zero.¹³

2.4.2 Retirement

Retirement occurs at age R for the household (both spouses have the same age) and the remaining life expectancy is denoted T . Pension income is based on labour market values in the last period before retirement,

¹²The question of sharing the household's assets is irrelevant in our model since we rule out savings.

¹³In a robustness check, we allow remarriage to take place, in expectation over two periods, i.e. 10 years and make the following assumptions regarding this new union: the next spouse is in the same wage quintile as the ex-spouse (to preserve the degree of assortative matching), there is no further fertility and the new spouse does not influence the choice of childcare mode for the children of the first marriage, and there is no further divorce. To keep things tractable, we assume that the next spouse's labour supply is the same as the ex-spouse's. The results obtained, available upon request, are very similar to the ones presented here, presumably because the event of remarriage is distant in the future and heavily discounted by households.

with a replacement rate of ρ . The (constant) household budget between age R and age $R + T$ is thus:

$$z^R = 40 \cdot \rho \cdot (w_{m,R-1} + w_{f,R-1}) \quad (5)$$

We assume post-retirement divorce and fertility hazard to be zero, so that the only variations in the state variables in retirement relate to the ageing of children under 16 at the time of retirement.

Similarly, for retired divorcees, the only source of variation in the two ex-spouses environment is the ageing of the remaining children.¹⁴ The fraction of the male pension income allocated to the female ex-spouse is included in the child support specification (see section 4.1 for our specific assumptions).

2.5 Family decision making

We set our framework within the collective model literature and consider that the household decision making consists in maximising a weighted sum of the two spouses' asset values, subject to the constraint that both spouses are willing to remain in the marriage. The weight μ is the relative bargaining weight of the male partner in the household decision. Given values of the state variables, s (wages, age, ages of the children), the household chooses x (labour supply, domestic times, consumption shares) in order to maximise:

$$\mu \cdot V_m(s, x) + (1 - \mu) \cdot V_f(s, x)$$

where μ and $1 - \mu$ are the relative bargaining weights of the two spouses, V_i , $i = m, f$ are the asset values of both spouses (more formally defined in section 2.6).

The resulting values of this optimisation for both partners are denoted $V_i^M(\mu, s)$ and depend on the Pareto weights. In addition to this, both partners value their marriage through a 'love' term ϵ , which captures the quality of the match. This term is common to both partners and is subject to shocks. Each partner is content to stay in the marriage if it yields a valuation $V_i^M(s, \mu) + \epsilon$ that is greater than their outside option. We will consider that outside options refer to divorce and are denoted $V_i^D(s)$.

We assume that divorce can be triggered unilaterally so this bargaining weight can be adjusted in the event of one of the partners threatening to dissolve the marriage.¹⁵ With unilateral divorce and forward-looking agents we are in the presence of forward-looking constraints relating to the continuing participation of both spouses in the marriage. Such optimisation problems can be represented as recursive contracts, as developed by Marcat and Marimon (2017), whereby bargaining weights adjust whenever a participation constraint binds and the Lagrange multiplier attached to it is non-zero. This yields a limited commitment model as in Ligon et al. (2002) and Chiappori and Mazzocco (2017) where the partnership may survive after an adverse shock leading to one of the parties desire to leave it through a renegotiation of the bargaining weights.

Note that $V_m^M(s, \mu)$ is increasing in μ and we can define $\mu_m^*(s, \epsilon)$ such that $V_m^M(s, \mu_m^*(s, \epsilon)) + \epsilon = V_m^D(s)$ as the bargaining weight that makes the male partner indifferent between divorce and staying in the marriage

¹⁴And remarriage in our robustness check mentioned above.

¹⁵As shown in Mazzocco (2007) and Voena (2015), when divorce is not an option or can only be legally initiated when both spouses agree to it, bargaining weights are fixed over time. This is the full commitment model.

with this bargaining power. Of course, there may be no value of $\mu \in (0, 1)$ such that this condition is satisfied, in which case the male partner prefers divorce even if he has all the bargaining power in the marriage, i.e. $\mu = 1$. Similarly, we define $\mu_f^*(s, \epsilon)$ such that $V_f^M(s, \mu_f^*(s, \epsilon)) + \epsilon = V_f^D(s)$, the bargaining weight that makes the female partner indifferent between marriage and divorce (if it exists). Since $V_f^M(s, \mu)$ is decreasing in μ , a renegotiation triggered by the female partner will yield an decrease in μ . The workings of this limited commitment model are illustrated in Figure E-1 and explained in more detail in Appendix E.

Formally, the dynamics of the Pareto weights can be summarised as follows:

$$\mu_t = \mu_{t-1} \text{ if } \begin{cases} V_m^M(s_t, \mu_{t-1}) + \epsilon_t & \geq V_m^D(s_t) \\ V_f^M(s_t, \mu_{t-1}) + \epsilon_t & \geq V_f^D(s_t) \end{cases}$$

$$\mu_t = \mu_m^*(s_t, \epsilon_t) \text{ if } \begin{cases} V_m^M(s_t, \mu_{t-1}) + \epsilon_t & < V_m^D(s_t) \\ V_f^M(s_t, \mu_m^*(s_t, \epsilon_t)) + \epsilon_t & \geq V_f^D(s_t) \end{cases}$$

$$\mu_t = \mu_f^*(s_t, \epsilon_t) \text{ if } \begin{cases} V_m^M(s_t, \mu_f^*(s_t, \epsilon_t)) + \epsilon_t & \geq V_m^D(s_t) \\ V_f^M(s_t, \mu_{t-1}) + \epsilon_t & < V_f^D(s_t) \end{cases}$$

in other cases, divorce happens, with probability $\delta(s_t)$. Note that this probability does not depend on μ_{t-1} .

Both partners anticipate the above dynamics, which are endogenous in that current labour market choices impact on the law of motion of some state variables (wages) which in turn impact on future values of the bargaining power. This feedback effect plays a role in household decisions that we will discuss further in section 2.7.

2.6 Summary and formalisation

We now turn to the formalisation of the model, including the various ingredients described above. We will derive in turn the constraints, the laws of motion of state variables and the objective functions relating to the optimisation problems of the married household, of divorced individuals and of retired households, either married or divorced.

Constraints For the married household pre-retirement, the constraints are:

$$\begin{aligned} 40 &= 20 \cdot lm_g + L_g + dom_g \text{ for } g = m, f, \\ 20(w_m \cdot lm_m + w_f \cdot lm_f) &= C_m + C_f + p_{CC} \cdot \sum_{k=1}^3 \max\{CC_k - (dom_m + dom_f), 0\} \\ 0 \leq dom_g &\leq \max\{CC_k\} \text{ for } g = m, f, \\ dom_m + dom_f &\leq \max\{CC_k\}. \end{aligned} \tag{6}$$

In retirement, they become:

$$\begin{aligned}
40 &= L_g + dom_g \text{ for } g = m, f, \\
z^R &= C_m + C_f + p_{CC} \cdot \sum_{k=1}^3 \max\{CC_k - (dom_m + dom_f), 0\} \\
0 \leq dom_g &\leq \max\{CC_k\} \text{ for } g = m, f, \\
dom_m + dom_f &\leq \max\{CC_k\}.
\end{aligned} \tag{7}$$

where z^R is the household pension as defined in section 2.4.

For the male divorcee, the relevant constraints do not include childcare any more since he does not have custody of the children:

$$\begin{aligned}
40 &= 20 \cdot lm_m + L_m \\
20w_m \cdot lm_m - y(\kappa, w_m^D) &= C_m.
\end{aligned} \tag{8}$$

For the female divorcee, the amount of time devoted to childcare is still relevant:

$$\begin{aligned}
40 &= 20 \cdot lm_f + L_f + dom_f \\
20w_f \cdot lm_f + y(\kappa, w_m^D) &= C_f + p_{CC} \cdot \sum_{k=1}^3 \max\{CC_k - dom_f, 0\} \\
dom_f &\leq \max\{CC_k\}.
\end{aligned} \tag{9}$$

In retirement, the male divorcee consumes his time endowment as leisure and his pension z_m^R minus the amount of child support. The female divorced pensioner shares her time between leisure and childcare if there is any remaining need and consumes her pension z_f^R plus the child support.

Laws of motion Denote s the vector of state variables:

$$s = \{a, \kappa, w_m, w_f\}, \tag{10}$$

The ageing process of the household is deterministic (we rule out death hazards). The dynamics of κ are exogenous and described by the transition matrix A^K , as explained in section 2.1. In both divorce and retirement the fertility rate is zero so that the dynamics of κ boil down to the ageing of the population of existing children and are captured by a transition matrix A_0^K . The dynamics of each spouse's wage quintile is endogenous as they depends on labour market choices. They are described by the transition matrices $A^W(lm)$ as mentioned in section 2.3.

As in Marcet and Marimon (2017), we consider that the bargaining power μ is an additional state variable, the dynamics of which are captured by the mechanism described in section 2.5. In the empirical section we will discretise the set of values of μ with a set M so that the law of motion of the male Pareto weight can be captured by a transition matrix too. Given a distribution of the 'love' shock ϵ and the set M , we can calculate the transition matrix A^μ between these values as a function of the model parameters. These

dynamics also depend on all other state variables and on the labour market choice: $A^\mu(s, lm)$, where we denote lm the vector of labour market choices: (lm_m, lm_f) .

Since we have assumed that the shocks to the various processes are independent, we can write the law of motion of the vector of state variables (s, μ) as the following products:

$$\begin{aligned}
\pi((s', \mu')|s, lm, \mu) &= \mathbf{a}^{\mathbf{K}}(\kappa, \kappa') \cdot \mathbf{a}^{\mathbf{W}}(w_m, w'_m|lm_m) \cdot \mathbf{a}^{\mathbf{W}}(w_f, w'_f|lm_f) \cdot \mathbf{a}^\mu(\mu, \mu'|s, lm) \\
\pi_m^D(s'|s, lm) &= \mathbf{a}_0^{\mathbf{K}}(\kappa, \kappa') \cdot \mathbf{a}^{\mathbf{W}}(w_m, w'_m|lm_m) \\
\pi_f^D(s'|s, lm) &= \mathbf{a}_0^{\mathbf{K}}(\kappa, \kappa') \cdot \mathbf{a}^{\mathbf{W}}(w_f, w'_f|lm_f) \\
\pi^R(s'|s) &= \mathbf{a}_0^{\mathbf{K}}
\end{aligned} \tag{11}$$

where π denotes the transition probability within a married household pre-retirement, π_m^D (respectively π_f^D) the transition probabilities for divorced males (respectively females) before retirement and π^R the transition probabilities for all retired households and divorced individuals.

Optimisation As seen above the flow utilities for each partner are $\mathcal{U}_g(C_g, L_g, D)$. In the first instance, let us note that, given labour market choices, the decision regarding C_g, L_g and dom_g does not have any future consequences. As a result, we can consider that this decision is made in the static framework where the household maximises the weighted sum of instantaneous utilities:

$$\max_{(C_m, C_f, L_m, L_f, dom_m, dom_f)} \mu \cdot \mathcal{U}_m(C_m, L_m, D) + (1 - \mu) \cdot \mathcal{U}_f(C_f, L_f, D) \tag{12}$$

given the choice lm , the value of the state variables s and the constraints seen above. The optimal level of household instantaneous utility thus achieved is denoted $\tilde{\mathcal{U}}_h(lm, s, \mu)$.

In retirement both partners cease participation in the labour market and the only remaining dynamics relate to the ageing of children yet to reach the age of 16. The household asset value at retirement is thus:

$$V^R(s_R, \mu_R) = \tilde{\mathcal{U}}_h((0, 0), s_R, \mu_R) + \sum_{\tau=R+1}^{R+T} \beta^{\tau-R} E \left[\tilde{\mathcal{U}}_h((0, 0), s_\tau, \mu_R) | s_{\tau-1} \right] \tag{13}$$

Now turning to (discrete) labour market choices, recall from sections 2.3 and 2.5 that these choices will have an impact on future periods through the hazard rates of wage progression and bargaining power dynamics which depend on labour market status. The dynamic problem faced by household h is to choose $lm = (lm_m, lm_f)$ to maximise expected lifetime household utility:

$$\begin{aligned}
\tilde{\mathcal{U}}_h(lm_t, s_t, \mu_t) &+ \sum_{\tau=t+1}^{R-1} \beta^{\tau-t} E \left[\tilde{\mathcal{U}}_h(lm_\tau, s_\tau, \mu_\tau) | lm_{\tau-1}, s_{\tau-1}, \mu_{\tau-1} \right] \\
&+ \beta^{R-t} E \left[V_h^R(s_R, \mu_R) | lm_{R-1}, s_{R-1}, \mu_{R-1} \right]
\end{aligned} \tag{14}$$

This yields an optimal labour market choice, \tilde{lm} , and the household asset values defined as:

$$V^M(s, \mu) = \mu V_m^M(s, \mu) + (1 - \mu) V_f^M(s, \mu) \tag{15}$$

and satisfy the following Bellman equations since we use the Pareto weights as an additional state variable and its law of motion described in section 2.5:¹⁶

$$V^M(s_t, \mu_t) = \tilde{U}_h(\widetilde{lm}_t, s_t, \mu_t) + \beta \cdot E \left[V^M(s_{t+1}, \mu_{t+1}) | \widetilde{lm}_t, s_t, \mu_t \right] \quad (16)$$

All expectations are taken conditional on state variables and labour market choices in the previous period since all processes are first-order Markov. Specifically:

$$E[V_h(s_\tau, \mu_\tau) | lm_{\tau-1}, s_{\tau-1}, \mu_{\tau-1}] = \sum_{(s_\tau, \mu_\tau) \in SxM} \pi((s_\tau, \mu_\tau) | lm_{\tau-1}, s_{\tau-1}, \mu_{\tau-1}) \cdot (\delta(s_\tau) [\mu_{\tau-1} V_m^D(s_\tau) + (1 - \mu_{\tau-1}) V_f^D(s_\tau)] + (1 - \delta(s_\tau)) V_h(s_\tau, \mu_\tau)) \quad (17)$$

where the asset values in divorce, V_m^D and V_f^D do not depend on the bargaining power at the time of divorce since the post-divorce allocation is governed by regulations as set out in section 2.4 and can be expressed as follows:

$$V_g^D(s_t) = \tilde{U}_g(\widetilde{lm}_{g,t}, s_t) + \beta \cdot E \left[V_g^D(s_{t+1}) | \widetilde{lm}_{g,t}, s_t \right] \quad (18)$$

for $g = m, f$ and where $\tilde{U}_g(lm_g, s)$ is the result of ex-spouse g maximising their instantaneous utility given (lm_g, s) subject to the constraints (8) and (9):

$$\begin{aligned} \tilde{U}_f(lm_f, s) &= \max_{dom_f} \mathcal{U}_f(C_f, L_f, D) \\ \tilde{U}_m(lm_m, s) &= \mathcal{U}_m(C_m, L_m, D) \end{aligned} \quad (19)$$

and \widetilde{lm}_g are their optimal individual labour supply. Note that, given his labour supply, the male divorcee does not have any remaining choice since C_m and L_m are both dictated by his budget and time constraints and D is determined by the female divorcee's choice of domestic time dom_f .

2.7 Time inconsistency and inefficiency

The household collective decision-making unit is a hybrid of the preferences of the two partners. It thus has changing preferences over time since Pareto weights may vary, even though the preferences of both partners are constant. Today's household is characterised by the composite preferences determined by μ_t and makes choices in the anticipation that the future household will be guided by different preferences μ_{t+1} while today's choices have an impact on these preferences since labour market choices affect future labour market values, which affect divorce outcomes in asymmetric ways and yield variations in bargaining power (in expectation).

Time inconsistency resides in the fact that $V_m^M(s_{t+1}, \mu_{t+1})$ and $V_f^M(s_{t+1}, \mu_{t+1})$ are the results of the optimisation problem solved by the household with preferences μ_{t+1} whereas $\mu_t \cdot V_m^M(s_{t+1}, \mu_{t+1}) + [1 - \mu_t] \cdot V_f^M(s_{t+1}, \mu_{t+1})$ is the valuation of this result by the household with preferences μ_t . Note that this is not the object that will be maximised by the decision-making entity characterised by μ_{t+1} . Sally (2000) analyses decision making in a context of a forward-looking agent and changing preferences. The agent (Ulysses) makes

¹⁶As in Marcet and Marimon (2017), we are able to formalise household choices in this way after adding the co-state variable μ_t , the dynamics of which embody the forward-looking constraints that both partners are willing to stay in the marriage.

choices today in the understanding of the impact of these choices on future preferences, and maximising his welfare in terms of today's preferences. Manipulating tomorrow's individual (and preferences) to make choices that are consistent with the preferences of today's decision maker are one of the determinants of today's choices.

Within a marriage, these changing preferences are in the form of variations in the relative bargaining power. Basu (2006) shows the interplay between household decisions and the balance of power within the household in a dynamic setting, where agents are forward looking and anticipate the dynamics of the nature of the 'hybrid', which leads to the possibility of Pareto inefficiencies. Lundberg and Pollak (2003) showed that inefficient outcomes occur in marriages because spouses cannot commit not to exploit a bargaining advantage arising in the future as a result of their partner's investment in the marriage today.¹⁷

In our framework, the source of inefficiency lies in the fact that any spouse cutting on labour market participation to supply the public good by producing domestic childcare bears an expected loss in the event of divorce which cannot be compensated for, since contracts between spouses within a marriage are difficult to enforce. Indeed, divorce regulations do not usually allow for compensation for loss of labour market value. We can thus be in a situation where Pareto improvements could be achieved if the working spouse could borrow individually to transfer an asset to the spouse providing home-produced childcare.

3 Descriptive statistics

3.1 Data

The data we use comes from the British Household Panel Survey, which covers years from 1991 to 2008. We restrict our attention to observations relating to households with two adults aged 20 to 70 and follow these couples (married or cohabiting) from their entry into the survey if they are married in their first year of observation or their first marriage if they enter the survey as 'never married' until their exit from the survey or their divorce. We will use the 'marriage' denomination for both married and cohabiting couples. This selection leaves us with 4,467 couples whom we follow for 1 to 18 years, 2,660 of these staying in the sample for at least 6 years. The mean number of years in the sample is 8.22 years and we have 36,717 household-year observations.

The survey provides detailed information on the number and ages of the children present in the household. 44% of observations relate to households with no children present (these may have had children who do not live in the same residence anymore). Few households have three children or more (6.9% observations). The summary statistics of our sample are shown in Table 1.

¹⁷Similar issues abound in the literature. Mazzocco (2007) rejects the full efficiency model with US data and finds evidence of variations over time of the relative bargaining power of spouses. Duflo and Udry (2004) present evidence that expenditure patterns in households in Ivory Coast are not consistent with a Pareto efficient allocation of household resources. Aura (2005) discusses the impact of different divorce laws on consumption and saving of married couples that cannot commit. He focuses on the fact that future behaviour is constrained by the outcome of future renegotiation process and that today's choices affect this process. Stevenson (2007) also finds that divorce laws affect the extent of specialisation within marriage.

across households		
	mean	sd
Household age	38.5	13.6
Age gap	2.5	5.8
N children (max)		
	%	
0	44.1	
1	17.8	
2	25.4	
3 and more	12.8	
<i>N</i>	4,467	
across observations		
	%	
No children present	44.0	
Children under 5 present	18.9	
Children aged 5-11 present	24.0	
Children aged 12-18 present	13.1	
<i>N</i>	36,717	

Table 1: Descriptive Statistics

The age structure of the children population in the household drives the need for childcare as explained in section 2.1. In the model simulations, we will restrict our attention to cases where there are at most two children for two reasons. First, relatively few households (12%) have more than two children in our data. Second, the market price of two nursery places is roughly the same as that of a nanny so that capping the childcare need in terms of time and cost of market services to the need of the two youngest children in the households is a fair representation of actual time/money costs.

We will use the index K to represent the age structure of a children population κ as:

$$K(\kappa) = 100 \sum_{i=1}^3 \mathbb{1}_{\{11 < \kappa_i \leq 16\}} + 10 \sum_{i=1}^3 \mathbb{1}_{\{5 < \kappa_i \leq 11\}} + \sum_{i=1}^3 \mathbb{1}_{\{\kappa_i \leq 5\}} \quad (20)$$

so that, for example, a household with children aged 4 and 13 will be represented by an index $K = 101$. K can only take ten values to represent households with up to two children. The variations in K that are useful in our identification yield variations in time and money need for childcare. We thus regroup categories that yield similar needs on these two dimensions¹⁸ and focus on the following 7 types of children population, the distribution of which is reported in Table 2.

¹⁸The categories 100 and 200, 10 and 110, 1 and 101 have been regrouped.

across observations	all ages	ages 21-34	ages 35-44	ages 45-54
K	%			
0	56.00	48.14	24.58	58.47
1	8.93	19.35	8.70	1.25
2	3.72	8.72	4.63	0.52
10	10.21	6.80	20.59	12.97
11	5.93	10.53	12.27	1.65
20	7.51	5.67	17.46	5.90
100	7.71	0.80	11.77	19.23
N	37,644	10,590	10,073	7,270

Table 2: Children population

The dynamics of K reflect the ageing of the children the household and fertility hazards, which we assume to be driven only by the current κ and the household age. We assume fertility hazards to be exogenous and thus rule out a potential feedback mechanism of current labour supply decisions on future fertility. If fertility hazards were allowed to vary with parents' wages then fertility motives would enter decisions on labour supply since these affect future wages. This would add another mechanism to our analysis and we have excluded it on the grounds that it is likely to be quantitatively minor relative the mechanisms that we focus on. Given the small income elasticities typically found in the literature (for example Schultz (2005)), the size of the fertility changes that we ignore is probably small.¹⁹ Birth rates peak at the age of 28, increases in the presence of children under 5 in the household and decreases sharply in the presence of older children. The first-order Markov process of the dynamics of K is characterised by the $(7, 7)$ transition matrices A_K^{age} reported in the Appendix in Table A-1.

The joint distribution of education levels is reported in Table 3, where 'high' refers to degree level education, 'medium' refers to the completion of A-levels and 'low' refers to a maximum qualification of strictly less than A-levels.

Educ. F	high	medium	low
Educ. M			
high	6.22	2.66	7.97
medium	2.88	3.74	12.88
low	5.06	8.25	50.34
$N = 3,953$			

Table 3: Education joint distribution

We note some modest degree of assortative matching by education with 60% of the sample on the first diagonal. In 16% of households, the female spouse has a higher level of education than the male spouse. In 24% of households the reverse is true.

¹⁹Besides, our (unreported) estimates of fertility hazards with additional controls include insignificant coefficients on the education variables.

3.2 Stylised facts

Now turning to the joint distribution of wages, we consider quintiles of hourly labour market values for all individual in our sample. For all individuals in employment this is simply their monthly labour earnings divided by their hours of work. For individuals out of the labour force we use the last hourly wage quintile observed for this individual if they have been observed in employment at some previous date in the sample or the predicted wage quintile at age 21-25 for individuals of the same education and gender. This will allow us to trace all individuals labour market values, which is essential in our framework where these play a role in the household decision. In the following, we will refer to both actual wage quintile and inputted labour market value as ‘wage’.

Hourly wages are calculated as the ratio of monthly labour income and the sum of usual hours and overtime hours (with a weight of 1.5 on overtime hours to reflect overtime pay). The mean wage levels in each quintile are (3.71, 6.47, 7.86, 10.19, 18.67). About 10% of wages fall below the official minimum wage. The rate of non-compliance by employers with minimum wage regulation is probably not this large and we expect the bulk of these low wages to come from the fact that, in the BHPS, hours are reported as hours ‘normally worked per week’ while monthly income is reported for the previous month.

	female wage quintile				
	4.89	4.52	1.60	1.83	1.40
male	4.53	5.09	1.93	1.83	1.07
wage	4.14	7.50	2.05	2.19	1.19
quintile	5.15	6.92	3.80	5.22	2.90
	5.18	6.76	3.97	6.72	7.63

Table 4: Joint distribution of wage quintiles

The joint wage distribution is shown in Table 4 in our sample of 36,717 observations of individuals living in a couple. Consistent with the assortative matching in education observed above, we see some positive association of female and male wage quintiles, with nearly 25% of the joint distribution on the first diagonal. Observations where the male spouse is on higher wage quintile than the female spouse represent 55% of the sample while the reverse is observed in 20% of the sample. The joint distribution of wages is much less equal between genders than the joint distribution of education in a large part because the accumulated labour market experience differs markedly between genders as we will see shortly.

Wage processes, part-time, full-time and non-participation As detailed in section 2.3, we model wage dynamics as a first-order Markov process for wage quintiles, where transition hazards depend on initial labour market choices. Table B-1 in the Appendix displays our estimates of 5-year transition matrices \mathbf{A}_{lm} for $lm = 0, 1, 2$ relating to initial non-participation, part-time and full-time employment. Our sample size and the fact that few men work part-time precludes a reliable estimation of transition matrices for both genders in each labour market state so we have opted to estimate these matrices for both genders together.

This is a limitation since returns to experience have been documented to vary across genders. However, the bulk of this difference stems from the fact that some of the female labour experience is in part-time employment, which we do account for with our transition matrices which are specific to each labour market choice.

Two facts emerge from these transition matrices that are key to the approach we take in this paper: first, there are substantial costs attached to part-time work and non-participation in terms of earnings dynamics²⁰, so it is useful to include these in household labour market decisions by using a dynamic framework. Second, these costs vary depending on one's initial position within the wage distribution. For example, it is more costly in terms of expected wage progression to work part-time for individuals initially on the third wage quintile than for those on the first (lowest) wage quintile. This variation in the cost-benefit analysis of opting for part-time work or non-participation with individuals' current wage quintile will help the identification of the model in our estimation below.

Monetary cost-benefit analysis Without running our structural estimation below and uncovering more information about households' preferences, we can already take a look at the balance of purely monetary pros and cons of supplying labour or domestic childcare. For example, if we consider part-time relative to full-time work for the female spouse, based on current loss of labour income, expected loss of labour income and savings in childcare services over all households in our sample, we obtain the distribution of the net costs of switching to part-time that is displayed in Figure 1. Note that for some households in our sample there are no children in need of childcare within working hours, i.e. they may be absent or over 16, so there are no savings in nursery fees associated with this switch.

²⁰Most of the literature on the part-time pay penalty focuses on static differences rather than these differences in dynamics. Its key findings are that much of the raw penalty of part-time work comes from differences in education and occupation, and that occupation downgrading often occurs in transitions from full-time to part-time work (see Connolly and Gregory (2009), Harkness (1996) and Manning and Petrongolo (2008)). We do not account for an instantaneous pay penalty from switching to part-time. This penalty arises over time in the form of a reduced wage growth. Besides, the only way we account for selection is to compute transition rates specific to each initial wage quintile, thus controlling for initial earning capacity.

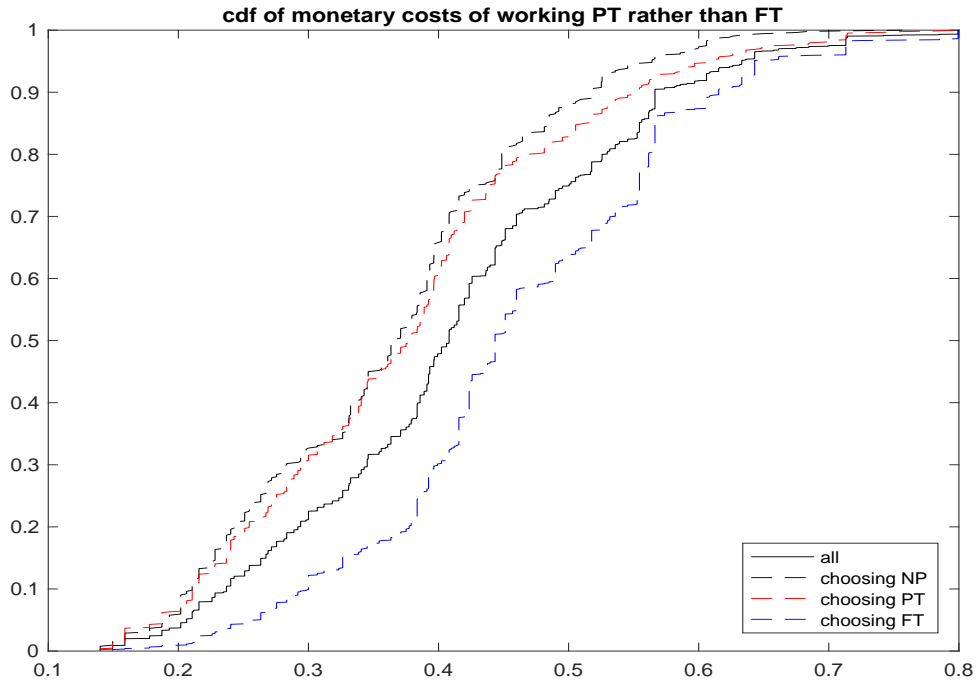


Figure 1: Distribution of monetary costs for households in the sample

Two striking features of this graph are the large variance of the distribution of costs and the fact that these costs are not systematically different between households choosing that the wife works part-time and those preferring that she works full-time. While these monetary costs are on average lower for households who choose part-time than for those who choose full-time (respectively 38% and 44% of potential²¹ household budget), there is a lot of overlap between the two distributions of costs. Moreover, these monetary costs are substantial: they exceed 1/3 of potential household budget in 65% of households who opt for less than full-time labour supply for the female spouse, and in 85% of households who opt for full-time labour supply of the wife. We conclude from this figure that a monetary comparison between labour supply choices does not explain the observed choices well and that a better understanding of households' heterogeneous preferences with respect to the mode of childcare is needed to enrich that picture.

3.3 Moments of interest

Labour market choices We will assess the fit of our model by comparing its predictions with the data in terms of labour market choices and labour market trajectories. Within our theoretical framework these will be chiefly affected by the age composition of the children population in the household and wages of the two spouses.

The distribution of household labour supply choices is reported in Table 5 by composition of the children composition in the household, embodied by the value of K and by values of each spouse's wage quintile. In

²¹Defined as the sum of the two spouses' full-time labour earnings.

our data, few men work less than full-time so we regroup all observations where this is the case into a single category.

	$lm_m < FT$	(FT, NP)	(FT, PT)	(FT, FT)
K				
0	0.08	0.09	0.14	0.69
1	0.11	0.27	0.33	0.29
2	0.16	0.41	0.29	0.13
10	0.09	0.18	0.41	0.33
11	0.13	0.35	0.38	0.15
20	0.10	0.21	0.47	0.23
100	0.08	0.11	0.37	0.45
w_m				
1	0.13	0.21	0.24	0.42
2	0.07	0.19	0.27	0.48
3	0.19	0.14	0.24	0.43
4	0.06	0.16	0.28	0.50
5	0.08	0.20	0.33	0.39
w_f				
1	0.10	0.16	0.39	0.35
2	0.14	0.31	0.24	0.31
3	0.07	0.15	0.25	0.53
4	0.07	0.11	0.21	0.60
5	0.08	0.08	0.32	0.51

Note: NP: non-participation, PT: part-time, FT: full-time.

Table 5: Household labour supply (lm_m, lm_f)

The dynamics of household labour supply choices are reported in the Appendix in Table C-1, aggregated over all ages 21-54 and by age category. Within our framework, these dynamics are driven by the dynamics of K (ageing of household’s children, fertility events), the dynamics of wages (stochastic wage changes given labour market choices) and shocks to marriage ‘quality’, ϵ .

Divorce hazard As seen in section 2.4, the possibility of divorce plays a crucial role in the dynamics of the bargaining power within the couple. It also affects labour supply decisions in that post-divorce outcomes are affected by labour market values, which are themselves influenced by labour market status. The five-year divorce hazard in the data averages at 9.4%.

4 Estimation procedure

We present here the estimation method and the assumptions we make. We will use estimated processes presented above of the dynamics of the age structure of the children population (section 3.1) and the dynamics of wage quintiles (section 2.3). These two processes are assumed to be exogenous and are both key in labour supply decisions of forward looking households.

Besides, we make a number of assumptions described below regarding preferences, divorce arrangements

and retirement (Section 4.1), childcare needs by child age and the initial distribution of power (Section 4.2). The method of estimation is described in Section 4.3 and identification is discussed in Section 4.4.

4.1 Assumptions

Preferences Individuals derive utility from consumption, leisure and from the public good described in section 2.2, i.e. the fraction of the childcare need carried out domestically, D as follows:

$$\mathcal{U}_g(C_g, L_g, D_h) = \exp[\beta_h(\alpha_g \log C_g + (1 - \alpha_g) \log L_g) + (1 - \beta_h) \log D_h] \quad (21)$$

Whilst α^g is individual-specific, the coefficient β_h is household-specific. We allow individuals and households to be heterogeneous with regard to their relative preferences for these.²² Spouses also derive utility from being in a marriage, ϵ_h , which is common to both and subject to shocks every period (5 years). These capture the utility from being in this couple (aka love). The mean and variance of these shocks, $(m_\epsilon, s_\epsilon^2)$, are two of the parameters that we calibrate (see below). The distribution of ϵ is assumed to be normal.

We match divorce hazards to calibrate the mean and variance of marriage quality shocks. These are set to 380 and 90 respectively.

Divorce As mentioned in section 2.4, the male ex-spouse transfers child support y to the female ex-spouse following their divorce. The size of this transfer varies with the ages of the children and are a fraction of the male spouse's full-time labour earnings at the time of divorce. Our calibration is the following:

$$y(w_m^D, \kappa) = \left[0.025 \sum_{i=1}^3 \mathbb{1}_{\{11 < \kappa_i \leq 16\}} + 0.05 \sum_{i=1}^3 \mathbb{1}_{\{5 < \kappa_i \leq 11\}} + 0.1 \sum_{i=1}^3 \mathbb{1}_{\{\kappa_i \leq 5\}} \right] 40 \cdot w_m^D \quad (22)$$

For retired divorcees, the child support transfer is assumed to be a fixed percentage, 15% of the male ex-spouse's earnings if there are still children under 16 from the dissolved marriage.

Retirement The retirement age is set at 65 for all, with a remaining life expectancy of 20 years. The retirement income replacement ratio is set to 0.50.

4.2 Other calibrated parameters

The childcare needs of each age category are calibrated as follows, bearing in mind that these are time needs within the 40-hour working week: 40 hours for children under 5, 15 hours for children aged 5 to 11 and 5 hours for children aged 11 to 16. The market price for a unit of childcare services (per hour-child) is set to £3.5.²³ Sensitivity of our model predictions to the price of childcare will be assessed in section 5.4.

The annual discount rate is set at 0.95. Since the amount of leisure and domestic childcare decided upon within our framework relate to the time within working hours, we allow for the fact that individuals and

²²The parameter illustrating the household's preference for home-produced childcare is similar to the parameter measuring the dis-utility of joint work in Guner et al. (2018).

²³According to Which? website, childcare costs average 4.6 per hour for a child under 2 in nursery in 2015. The lower figure we use is meant to account for the fact that some households use grandparents' childcare.

household enjoy ‘baseline’ levels of both of these outside these hours. Thus, the quantities L and D that enter the utility function are the sum of the leisure and domestic childcare chosen within the model and the baselines levels, L_0 and D_0 respectively. We set L_0 to 25 hours (weekly) and D_0 to 0.2.

Finally, the relative taste for the male domestic time dom_m to the female domestic time dom_f , i.e. the parameter λ in equation 3, is set to 0.7.²⁴

4.3 Estimation method

We are then able to solve the model by backward induction from retirement and predict household labour market choices for each (w_m, w_f, K, μ, age) . The parameters to estimate are the joint distribution of individual preferences for consumption and leisure and the households’ taste for domestic childcare $(\alpha_m, \alpha_f, \beta)$.

Here we restrict our estimation sample to households which we follow over 3 periods, i.e. 15 years. Since households are allowed to be heterogeneous in their preferences but with constant preferences over time, keeping households with full 15-year trajectories helps us with identification. The data is arranged in $N = 938$ household labour market histories over three 5-year time periods, as well as records of both spouses’ labour market values, household age and composition of children population: $\{lm_{i,t}, w_{m,i,t}, w_{f,i,t}, K_{i,t}, age_{i,t}\}_{i=1..N, t=1..3}$. Recall that $lm_{i,t}$ captures both partners’ labour market choices.

We model the heterogeneity in the parameters in the form of discrete distributions, with no restrictions on the correlations between the three dimensions of heterogeneity. Similarly, the distribution of Pareto weights is approximated by a set of discrete mass points over a finite number of values for μ .

For each parameter value p and household i , the predicted household labour market choices given their state variables $s_{i,t} = \{w_{m,i,t}, w_{f,i,t}, K_{i,t}, age_{i,t}\}$ and μ are denoted $\widehat{lm}_{i,t}(\mu, p)$ for each value of μ . The indicator $I_{i,t}(\mu, p)$ takes the value 1 if $\widehat{lm}_{i,t}(\mu, p) = lm_{i,t}$ and 0 otherwise. We assume for simplicity that the unemployment rate is the same for all individuals in the sample so that, whatever their desired labour supply, we observe individuals not participating with some (constant) probability.

The limited commitment model described in section 2.5 is informative on the dynamics of Pareto weights in response to shocks to wages, fertility and marriage quality but does not deliver any predictions regarding the Pareto weights at the time of marriage formation, apart from the fact that both parties have to derive a higher asset value from marriage than from being single. Since we do not model marriage formation and consequently not the value of being single pre-marriage, we take a shortcut by assuming an initial distribution of Pareto weights that is uniform over the discrete support of values that we use in the estimation procedure. An initial negotiation ensues so as to adjust μ to fulfil both spouses’ participation constraints. We have examined the sensitivity of our results to alternative choices for this initial distribution and found qualitatively similar outcomes.²⁵

²⁴This parameter helps to predict that households where the female spouse is on a higher wage quintile still choose for her to work less than full-time to fulfil childcare needs domestically, as often observed in the data. We have experimented with a higher value of $\lambda = 0.8$ and found qualitatively similar results and a lower fit.

²⁵These results are available upon request.

We also predict dynamics of Pareto weights given parameter values $p = (\alpha_m, \alpha_f, \beta)$ –indeed, since p affects the two partners’ values in and outside of marriage, p will affect the dynamics of μ , embodied in the transition probabilities $\pi(\mu_{t+1}|\mu_t, s_{i,t+1})$, within our limited commitment model. Given an initial (assumed uniform) distribution of Pareto weights, partners engage in a first renegotiation upon forming their match yielding a distribution of weights $f(\mu|p, s_{i,1})$ where both partners are at least weakly better off in the marriage than in divorce at all points where the density is positive. We assume that the distribution of power in the couples in their first year in our sample is still the distribution obtained this way, even though most couples are not in the initial year of their marriage.

The likelihood of observing a household labour market trajectory given their history of state variables and values of parameters p takes the following expression:

$$\ell_i(p) = \sum_{\mu_1} \left(f(\mu_1|p, s_{i,1}) I_{i,1}(\mu_1, p) \left(\sum_{\mu_2} \pi(\mu_2|\mu_1, s_{i,2}) I_{i,2}(\mu_2, p) \left(\sum_{\mu_3} \pi(\mu_3|\mu_2, s_{i,3}) I_{i,3}(\mu_3, p) \right) \right) \right). \quad (23)$$

In the spirit of Train (2008) we specify a discrete support of values $\{\alpha_g^p\}_{g=m,f;p=1..P}$ and $\{\beta^p\}_{p=1..P}$ for these coefficients and then estimate the attached weights. Relative to estimating both mass points and weights, this makes the estimation much quicker and allows for more support points.

Considering that each set of parameter values $\{\alpha_m, \alpha_f, \beta\}$ represents a latent class, we use the following technique²⁶ to estimate the P^3 shares of each class:

- Start with an initial set of (prior) shares ω_p^0 for $p = 1 \dots P^3$.
- At each iteration j , obtain a set of households posteriors $\omega_p^j(i)$:

$$\omega_p^j(i) = \frac{\omega_p^{j-1} \cdot \ell_i(p)}{\sum_{p'} \omega_{p'}^{j-1} \cdot \ell_i(p')} \quad (24)$$

- Update the class shares as:

$$\omega_p^j = \frac{1}{N} \sum_i \omega_p^j(i) \quad (25)$$

- Recalculate posteriors with these shares until convergence to obtain $\hat{\omega}_p$.

These shares $\hat{\omega}_p$ are our estimates of the discrete joint distribution of $\{\alpha_m, \alpha_f, \beta\}$ over its P^3 points of support.

4.4 Identification

We do not have a formal proof of identification but can point to several sources of variations that help identifying the parameters of interest. First, as we saw in section 3.2, the cost of working less than full-time, both in the current period and in expectation depends on one’s initial wage quintile and age. Second, depending on the number of children needing childcare, one hour of a parent’s domestic time may save from

²⁶This is similar to an EM algorithm without the maximisation step since we keep the points of support fixed.

1 to 3 hours of market childcare services. The trade-off between this loss of income, the possible increase in leisure time –since labour supply choices are discrete, in some cases, working part-time to fulfil the childcare need entails increased leisure time too– and the increased amount of the public good (the fraction of needed childcare carried out domestically) depends on the household’s preferences on consumption, leisure and home-produced childcare.

Third, working less than full time causes two additional losses for the *individual* concerned: the expected future loss in labour market value will weaken his/her bargaining power, in expectation, as well as his/her asset value post divorce. From an individual point of view, variations in bargaining power impact mostly on their share of the household consumption, so an individual will take a stronger interest in their bargaining power if their taste for consumption is high.

If an individual has a high Pareto weight at the time of the decision, they are likely to influence it by leaning against the option of them working less than full time (all other things equal). Indeed, the above two costs are not internalised by the spouse not expecting to bear them, which makes each spouse more likely to push for the other one to drop out of full-time. This, however, may be offset by the fact that the household values the domestic time of the mother more than that of the father (see section 2.2 and the parameter λ).

We have experimented with different numbers of latent classes by varying the size of the set of values that each parameter can take from 3 to 20. Using either AIC or BIC criterion points to the need for ever larger number of classes, but we observed that the computational cost of doing so was large relative to marginal increases in fit. A reassuring feature of the results with respect to identification is that the joint distribution of model parameters obtained is qualitatively very similar across all specifications. In the following we will present results for the specification where each parameter can take 10 values, which offers a good tradeoff between fit and running time, which allows us to run robustness checks in a reasonable running time.²⁷

5 Results and policy analysis

5.1 Fit

The fit of the model is presented in Tables D-1 and D-2 for the labour supply choices and labour supply transitions. The model offers a good fit and captures the main stylised facts in all these dimensions. Figure D-1 shows the distribution of household joint labour market trajectories in the data and as predicted by the model. Here too we see that the model does a very good job at fitting this distribution.

5.2 Coefficient estimates

The distribution of coefficient estimates are displayed in Table 6 and illustrated in Figures 2 and 3. We note a substantial degree of heterogeneity both in terms of the consumption-leisure trade-off and in terms of households’ taste for home-produced childcare. Some households (with a high β) do not derive much utility

²⁷Results for alternative specifications are available upon request.

from the fact that their children are looked after by a parent rather than a market provider of childcare, so take their labour supply versus childcare decision mostly based on its expected financial consequences and its impact on future bargaining power. Some other households, on the other hand, place a large value on domestic childcare and this taste will play a prominent role in their decision to supply labour in the presence of young children. Section 5.3 presents an illustrative example of these contrasting tastes.

Distribution of	α_m	α_f	β
Mean	0.636	0.548	0.534
Variance	0.078	0.084	0.079
Correlations	(β, α_m)	(β, α_f)	(α_m, α_f)
	0.053	0.159	-0.233

Table 6: Distributions of coefficient estimates

Figure 2: Marginal distributions of coefficients

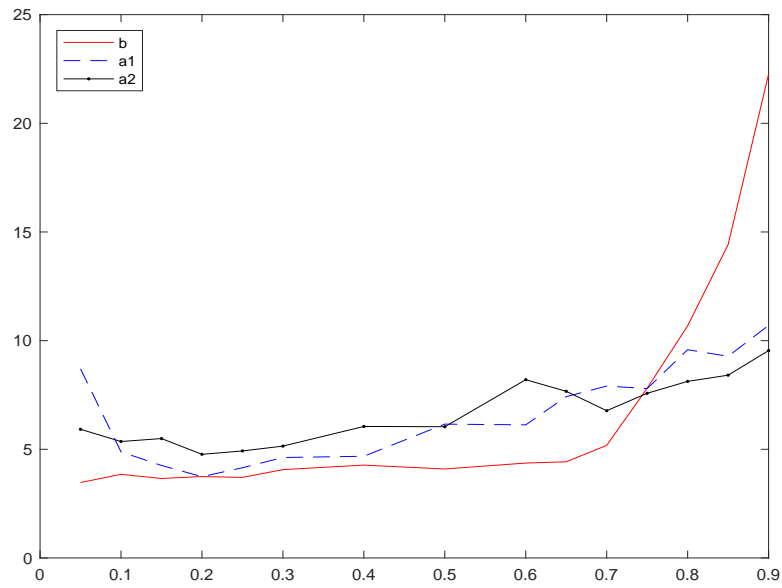
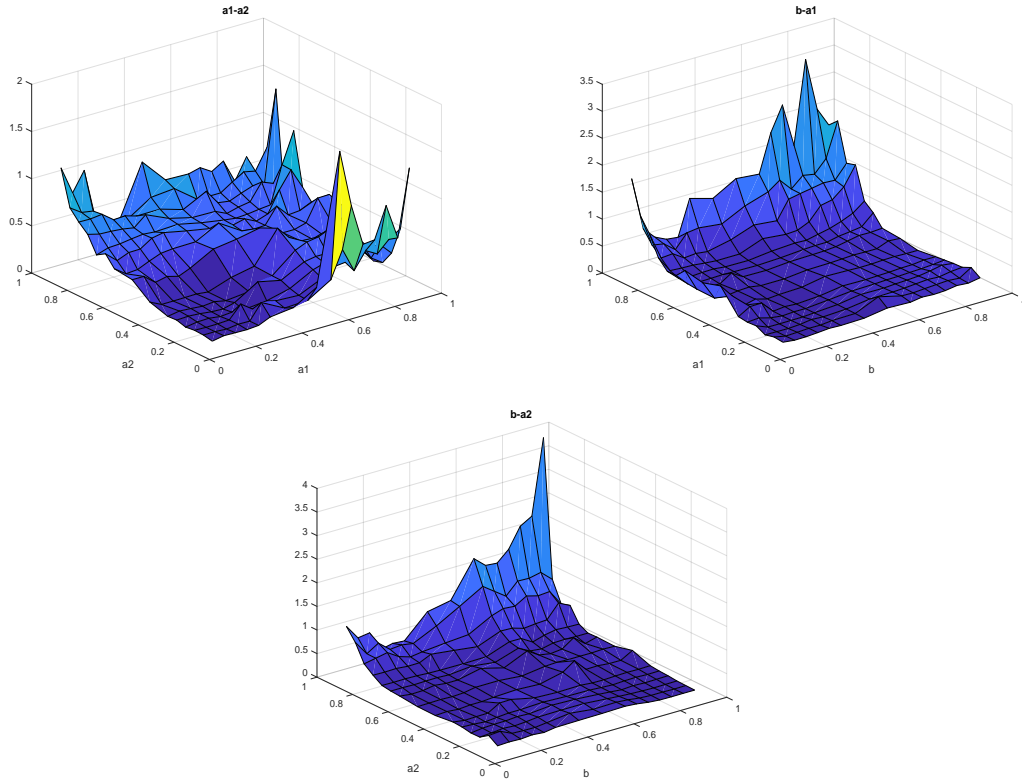


Figure 3: Bivariate distribution of coefficients



Estimates of the mean and variance of the marriage quality shocks give rise to dynamics of the Pareto weights within our limited commitment model which are illustrated by the following transition matrix between the values on the discrete support of μ – a grid of 10 evenly spaced values over the interval $(0.05, 0.95)$ – where the current μ 's are represented by rows and the μ 's in the next period are represented by columns. This matrix is for illustration purposes only and shows transitions that are calculated for utility parameter values of $\beta = 0.7$, $\alpha_m = 0.5$ and $\alpha_f = 0.4$ and for a household where $(w_m, w_f, K, age) = (3, 3, 1, 3)$ – as an example. In this matrix, transitions occur solely as a result of shocks to marriage quality, keeping wages and children constant.

$$\begin{pmatrix} 0.28 & 0.18 & 0.26 & 0.16 & 0.11 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.46 & 0.26 & 0.16 & 0.11 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.72 & 0.16 & 0.11 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.89 & 0.11 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 1.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 1.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.10 & 0.90 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.10 & 0.14 & 0.76 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.10 & 0.14 & 0.15 & 0.61 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.10 & 0.14 & 0.15 & 0.19 & 0.41 \end{pmatrix}$$

Table 7: Transition matrix between values of the Pareto weights

The shape of this transition matrix illustrates the workings of the limited commitment model: for values

of the Pareto weights near 0.5, both partners have near equal surplus in marriage relative to the state of divorce, there is thus no possible (discrete) adjustment of the Pareto weight that would increase one party’s surplus above zero (should it become negative due to an adverse shock to ϵ) without pushing the other party’s surplus below zero. Consequently, Pareto weights are fully persistent around 0.5 within continuing marriages. On the other hand, when starting from a very low value of the Pareto weight, e.g. on the first row of the above matrix, the female partner’s surplus is much higher than that of the male partner, there is thus scope to adjust the Pareto weight upwards if the match quality shock pushes the male partner’s surplus below zero but not the female partner’s (recall that the match quality is common to both spouses). There is however no case in which the male partner would be able to raise the Pareto weight above 0.5, since at that point the two surplus are near equal and we are back in the previous situation. These mechanisms account for the two triangles on non-zero values in the above matrix.

The table below shows our predicted average Pareto weights for different values of (w_m, w_f) , which demonstrates the influence of relative wages (or labour market values) on bargaining power and justifies the concern for future power in today’s choice of labour supply, given its consequences on expected wages. These figures are calculated for utility parameter values of $\beta = 0.7$, $\alpha_m = 0.5$ and $\alpha_f = 0.4$ and relate to households aged 1 without children -as an example. We note that we find much less amplitude in Pareto weights than Lise and Yamada (2019), where the range of μ goes from 0.2 to 0.75. Like them though we find that adjustments to Pareto weights in continuing marriages are fairly small (see Table 7 above).

	female wage quintile				
	0.508	0.505	0.482	0.484	0.463
male	0.512	0.504	0.506	0.487	0.461
wage	0.531	0.509	0.506	0.507	0.480
quintile	0.535	0.518	0.514	0.510	0.481
	0.553	0.539	0.537	0.514	0.507

Table 8: Average Pareto weights by (w_m, w_f)

As a robustness check we have run our estimation with a different value of the discount rate illustrating a scenario where households are more myopic than in the benchmark estimation. Indeed, households may choose to have one spouse dropping out of full-time work not because they value domestic childcare highly but because they are shortsighted with regards to the future impact of this decision on future wages. Our alternative calibration of the annual discount rate is 0.90. The results we obtain²⁸ are qualitatively similar but include a distribution of the taste for home-produced childcare that is more skewed towards 0 and a lower fit to our data.²⁹

²⁸ Available upon request.

²⁹ Another robustness check that we perform is to experiment with a higher values of λ , equal to 0.8. Again, results are qualitatively similar but the fit is less good.

5.3 Illustration

Our coefficient estimates allow us to quantify the various components of the cost-benefit analysis underlying households' labour supply decisions. In this section we provide an illustration of the relative sizes of these components for some example values of the coefficients and the state variables. We are considering a family aged 35 with two children under 5, with utility parameters $(\alpha_m, \alpha_f, \beta) = (0.6, 0.6, 0.75)$ and calculating the costs and benefits of carrying out all childcare domestically by switching from (FT, FT) to (FT, NP) in three cases of wage quintiles for w_m and w_f : $(1, 3)$, $(3, 3)$ and $(5, 1)$. In the first two cases, the net benefit of switching is negative (see last column in Table 5.3) so these households would choose to supply full-time labour from both spouses. In the last case, the net benefit is positive (absent considerations regarding the future of bargaining weights, which we will discuss below), so the household would opt for non-participation in the labour for the wife in this period.

The components of the trade-off between these two alternatives are the following. First, denoted C_1 in Table 5.3, the home production of childcare entails one of the partners, here the wife, taking time off work³⁰ and foregoing current labour market income. Second, there is the sum of discounted losses in household income through loss of wife's income growth, denoted C_2 . Fourthly, we evaluate the current value of domestic childcare as *equivalent* household income gain, denoted B_1 . This willingness to pay is directly related to the coefficient β . In terms of monetary budget, this is to be compared with the instantaneous savings in terms of market childcare services, denoted B_2 .

(w_m, w_f) quintiles	C_1	C_2	B_1	B_2	Net benefit
$(1, 3)$	-0.64	-0.50	0.63	0.28	-0.34
$(3, 3)$	-0.49	-0.38	0.63	0.21	-0.03
$(5, 1)$	-0.19	-0.18	0.63	0.22	0.41

Table 9: All in percentages of (pre-childcare costs) household budget when both spouses work FT.

The most striking feature of this table is that, in spite of C_1 and B_2 attracting most of the attention in the public debate about mothers' labour supply, the other two components are each of similar magnitude to C_1 and B_2 . The current and expected monetary costs C_1 and C_2 of the mother's non-participation vary mainly with her labour market value (and with her spouse's too since these costs are expressed as percentages of the potential household budget). Variations in B_2 only depend on (w_m, w_f) because B_2 represents constant childcare costs across households with the same age children as a fraction of potential household budget. By contrast, B_1 , the valuation by the household of having their children looked after by the mother rather than by a market provider, is heterogeneous across households with different β 's. The variance of β thus plays a key role in explaining the variance of household choices given a monetary comparison between labour market choices (this is the puzzle shown in Figure 1). Indeed, since B_2 can be large and varies across households, it can account for households making choices that seem costly from a monetary point of view.

³⁰We assume that households do not decide to pay for childcare in normal working hours in order to increase the amount of leisure of an adult in the household, so that the relevant alternative to childcare is work. This seems to us a reasonable assumption in most households and will be checked with our estimated coefficients.

The last component of this trade-off is that an individual’s labour market value may (depending on the love shocks in future periods) impact on his/her bargaining power in household decision making. If that is the case, taking time off work in the current period means foregoing bargaining power in future household decisions, in expectation. Given that the cost of non-participation in terms of future income growth depends on current market value, the cost/benefit of non-participation in terms of Pareto weight in the next period depends on the current pair (w_m, w_f) . Our estimated average variations in Pareto weights with relative wages are shown in Table 8. With regards to our example, starting from a pair $(3, 3)$, the expected loss of Pareto weight for the female spouse is 0.018. This translates into a loss of consumption share of 3.5% of household consumption. As in Basu (2006), these anticipated changes in future negotiating power affect current labour market choices.

5.4 Households sensitivity to childcare policy

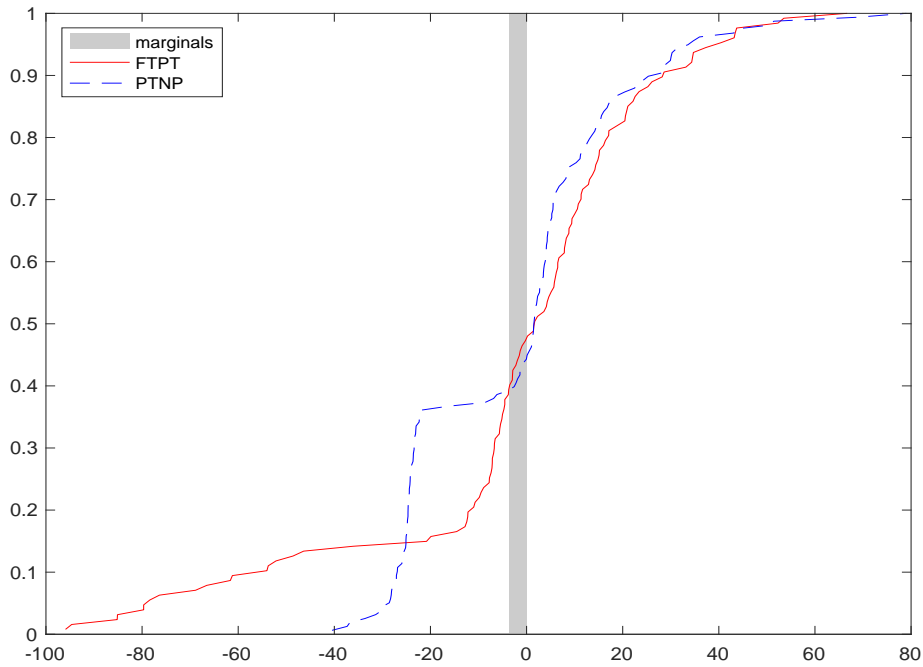
The example above illustrates how the cost-benefit analysis of labour supply versus home childcare depends on the household’s preference parameters, current Pareto weights, spouses’ wages and children’s ages. The only households who are likely to be responsive to a change in childcare prices are those for whom the net cost/benefit (in terms of lifetime utility) of changing their labour supply is close to zero. All others will value their current choice much more than the alternatives and will not react to a change in the childcare policy. From the policy-maker’s point of view it is interesting to quantify the mass of households who are marginal in the population in order to predict the effectiveness of a potential policy change and this is something we can do with this structural approach.

Our estimation allows us to compute the predicted household asset values of alternative labour supply choices. We can thus assess how far households are from changing their labour supply behaviour, i.e from being indifferent between two labour supply choices. We will call this their ‘distances to break-even’, which are formally defined as follows:

$$\begin{aligned}\Delta_{FT} &= V_h(FT, FT) - V_h(FT, PT) \\ \Delta_{PT} &= V_h(FT, PT) - V_h(FT, NP).\end{aligned}\tag{26}$$

These measure the gap in asset values between the labour supply choices (FT, FT) and (FT, PT) and between (FT, PT) and (FT, NP) . Figure 2 shows the c.d.f of the distances to break-even among our sample of households aged 30 to 35 and with at least one child under 5. The line labelled ‘FTPT’ (respectively ‘PTNP’) refers to the distribution of this distance among households where the mother either works full-time (resp. does not participate) or works part-time. The childcare need for these households is 40 units of parental time or at least (depending on other children) 40 units of childcare services. As is clear from our framework, these gaps will depend on the households taste parameters $(\alpha_m, \alpha_f, \beta)$. For each individual in our subsample, we use the posterior of this individual belonging to each parameter ‘class’ as well as the corresponding distribution of Pareto weights for this household in each class.

Figure 4: C.d.f.s of estimated distances to break-even



solid: Full-time vs Part-time; dashed: Part-time vs Non-participation

If the policy maker is considering a decrease in childcare prices equivalent to the width of the grey area,³¹ the marginal households are those for whom our calculated quantities are in the range of the shaded area. The households for whom these quantities are positive already supply labour before the policy change and are infra-marginal. The households for whom these quantities are below the shaded area will not be enticed to supply more labour by this policy change.

For the subsamples under consideration, Figure 4 tells us that 52% of households who are working either full-time or part-time are infra-marginal, i.e. would benefit from the policy but not alter their labour supply behaviour, while 8% of households are marginal and would switch from part-time to full-time with this new subsidy and 16% of these households would not switch to full-time even if the subsidy was 5 times as large. Among the subsample of households where the mother works part-time or not at all (dashed line), 56% are inframarginal in the sense that they are already participating, 5% would start participating with the proposed subsidy, but 36% of these households would still not participate (for the mother's labour supply) even if the subsidy was 5 times as large. This exercise allows us to quantify the heterogeneity of potential responses by households to policies of varying sizes and to identify the fractions of the target population who are very inelastic to childcare prices.

³¹For this exercise, we have set this width equivalent to a gain in utility provided by an addition to the household budget equal to the price of a full-time nursery place, for a household with average tastes for consumption and leisure.

5.5 Policy questions

We discuss here various policy implications of our results. Incentives to specialisation depend on several features of our model. First, taste for mother time (λ): custom, biology, time persistence (acquired taste). Second, taste for domestic childcare ($1 - \beta$) combined with expected loss of income due to part-time work or non-participation. Third, the initial distribution of Pareto weights. Divorce regulations such that partner incurring loss of labour market value from working less than full time is not compensated for this loss post-divorce. This implies that this partner not only has a lower post-divorce continuation value but also loses bargaining power within the marriage. This means that, all other things equal, the spouse with a high bargaining power initially is less likely to push the household labour supply decision where he/she works less than full-time.

A policy that would compensate any partner who has lost potential labour market earnings by providing the household's domestic childcare would thus have an impact on both post-divorce values and the balance of power in continuing marriages. The impact of such a policy on specialisation depends on the distribution of bargaining power at the time of the domestic childcare versus labour supply decision. For example, in a situation where the father has most of the bargaining power and the household decision is for the mother to specialise in childcare while the father specialises in labour supply, the above policy would lessen the incentive to specialisation as the male partner (and thus the household, in which decisions he weighs heavily) now internalises the expected cost of domestic childcare even post-divorce (in this case, the Pareto weights would exhibit more persistence). On the other hand, in a situation where the female partner has most of the bargaining power, the above policy lowers the expected cost of specialisation and increases the incentive for this household to specialise. In other words, this policy would shift some of the expected costs of working less than full-time from the parent providing childcare to the other parent. This would thus entice a household where the childcare provider has most of the power to specialise more and a household where the parent staying in the labour force has most of the power to specialise less. Both the policy itself and these responses would then lead to a more balanced division of power within households. The direct effect results from the fact that both partner's outside options would be affected more equally by the decision to produce childcare domestically. The indirect effect comes from the fact that households where the mother has more power have more incentives to adopt a choice (specialisation) that will reduce her power whereas households where the father has more power have less incentives to specialise, so less incentive to decide that the mother produces domestic childcare and loses bargaining power.

The crucial role of divorce laws on the dynamics of bargaining and on investments within marriages have been pointed out by Stevenson (2007). She shows that changes in divorce laws affect both the likelihood of divorce and the incentive to invest in the marriage. In our framework, since we ignore any long-term effect of domestic childcare on children outcomes that may enter the parents' utility, the time taken off labour market supply to look after children is a reverse investment in that benefits are reaped off in the current period while costs are spread over future periods.

One could take the view that if households choose to specialise because they value domestic childcare highly then there is no need for the policy maker to interfere with these decisions. However, the policy maker may find the long-term effects of specialisation undesirable for two reasons. First, since the childcare provider is often the woman and since working less than full-time damages future earnings, specialisation leads to greater future gender gaps in labour market value. Second, specialisation may produce higher numbers of children living in poor households. Indeed, when divorce occurs in the presence of children, custody is mostly given to the mother, who, in specialised households, has lost potential labour market value. The post-divorce household in which the children live is thus poorer when the pre-divorce household was specialised. Since gender gaps and child poverty are both items of the political agenda, the specialisation of households is indeed of interest to the policy maker.

Our analysis falls short of a prediction within a general equilibrium model of a response to policy changes,³² but we can nevertheless conjecture that policies aiming to increase the quality of market childcare (thereby shifting the preference distribution for home-produced childcare), to decrease the price of these services, to sustain human capital and labour market skills of childcare home providers, or to include larger compensation for loss of market earnings into divorce settlements, would all have an impact of the decisions analysed in this paper.

Finally, a compelling reason for policy intervention is the presence of inefficiency. In this context, as we saw in section 2.7, the fact that today's decisions alter future bargaining power is a source of inefficiency, arising mainly from the fact that future outside options are affected in an asymmetric manner by today's decisions, i.e. the post-divorce outcome of the spouse providing childcare is more affected than that of the other spouse. A policy to subsidise childcare would have no direct impact on these dynamics and hence the source of inefficiency, while a policy that would account for past investments in household public good in post-divorce settlements would reduce the scope for such inefficiency to arise.

6 Conclusion

In this work we aim to reevaluate the trade-offs faced by households when making labour supply decisions in the presence of children. We argue that the picture of this trade-off usually discussed in the public debate is incomplete. We take into account both the instant (and expected) monetary costs and benefits of supplying labour or home-made childcare and the households' valuation of parental childcare, which we allow to be heterogeneous. Our estimates show that this valuation is large and exhibits a large variance across households.

This is, to the best of our knowledge, a different angle on labour supply choices in the presence of children from the existing literature. Our estimates of the taste for home-produced childcare are relevant for policy design since this taste is bound to affect the response of households to childcare subsidies. Indeed,

³²It is also limited by our assumptions of exogenous fertility. This preclude any feedback mechanism whereby firms may alter their job creation behaviour in response to changes in female labour supply, or whereby fertility rates or wage determination are affected by these changes.

households with a strong preference for home-produced childcare will be expensive to entice into labour market participation from a policy maker point of view. Our results allow us to compute a structural estimate of the fractions of our sample who are likely or not to alter their labour supply choices when presented with subsidies to the price of childcare of various sizes. We find that a substantial fraction of households with young children are unresponsive to childcare policies of standard magnitude. Our analysis suggests that policies aimed at changing the perception of the quality of market childcare, hence shifting the distribution of the taste for parental childcare, are potentially as effective as subsidising the price of childcare.

In addition, we analyse the impact of one spouse working less than full-time on the future balance of power within the couple. In expectation, this spouse will incur a cost in terms of his/her future bargaining power and consumption share. An interesting component of the model is that this cost cannot be fully internalised by the household because of the hazard of divorce and because divorce settlements do not compensate for past labour market choices and public good provision. The inability to compensate the partner providing home-produced childcare and thus the household public good in the current period or in future settlements gives rise to a market failure within the household and to inefficiency of labour supply decisions.

A better understanding of these mechanisms has the potential to help policy makers to tackle gender gaps and household income inequality and warrants continued growth of this area of the literature. Two important ingredients that are not included in our framework are endogenous savings and fertility. Incorporating these in the model while keeping it tractable is a possible avenue for further research. In terms of the role of policy, we have restricted our attention to childcare subsidies and divorce regulations but have ignored the tax system. In Bick and Fuchs-Schündeln (2017) international differences in wages and tax schedules have been shown to explain a large part of international differences in married couples' labour supply. Guner et al. (2011) also find a substantial impact of tax reforms (joint versus separate filing) on married women labour supply and Gayle and Shephard (2019) derive an optimal tax system for married couples and its effect on household labour supply. This literature has shown that monetary incentives inherent in the tax system are important determinants of the household labour supply decision. Understanding its interaction with the mechanisms outlined in this paper could be a fruitful line of future work.

References

- ADDA, J., C. DUSTMANN, AND K. STEVENS (2017): “The career costs of children,” *Journal of Political Economy*, 125, 293–337.
- AURA, S. (2005): “Does the balance of power within a family matter? The case of the Retirement Equity Act,” *Journal of Public Economics*, 89, 1699–1717.
- BAKER, M., J. GRUBER, AND K. MILLIGAN (2008): “Universal child care, maternal labor supply, and family well-being,” *Journal of Political Economy*, 116, 709–745.
- BICK, A. (2016): “The quantitative role of child care for female labor force participation and fertility,” *Journal of the European Economic Association*, 14, 639–668.
- BICK, A. AND N. FUCHS-SCHÜNDELN (2017): “Taxation and Labour Supply of Married Couples across Countries: A Macroeconomic Analysis,” *The Review of Economic Studies*, 85, 1543–1576.
- BLAU, D. M. AND P. K. ROBINS (1988): “Child-care costs and family labor supply,” *The Review of Economics and Statistics*, 374–381.
- BLUNDELL, R., P.-A. CHIAPPORI, AND C. MEGHIR (2005): “Collective Labor Supply with Children,” *Journal of Political Economy*, 113, 1277–1306.
- BLUNDELL, R., M. COSTA DIAS, C. MEGHIR, AND J. SHAW (2016): “Female labor supply, human capital, and welfare reform,” *Econometrica*, 84, 1705–1753.
- BRONSON, M. A. (2014): “Degrees are forever: Marriage, educational investment, and lifecycle labor decisions of men and women,” *Unpublished manuscript*.
- BROWNING, M., P.-A. CHIAPPORI, AND V. LECHENE (2006): “Collective and Unitary Models: A Clarification,” *Review of Economics of the Household*, 4, 5–14.
- CHERCHYE, L., B. D. ROCK, AND F. VERMEULEN (2010): “An Afriat Theorem for the collective model of household consumption,” *Journal of Economic Theory*, 145, 1142–1163.
- (2011): “The Revealed Preference Approach to Collective Consumption Behaviour: Testing and Sharing Rule Recovery,” *Review of Economic Studies*, 78, 176–198.
- (2012): “Married with Children: A Collective Labor Supply Model with Detailed Time Use and Intrahousehold Expenditure Information,” *American Economic Review*, 102, 3377–3405.
- CHIAPPORI, P.-A. (1992): “Collective Labor Supply and Welfare,” *Journal of Political Economy*, 100, pp. 437–467.
- CHIAPPORI, P.-A., B. FORTIN, AND G. LACROIX (2002): “Marriage market, divorce legislation, and household labor supply,” *Journal of political Economy*, 110, 37–72.

- CHIAPPORI, P.-A. AND M. MAZZOCCO (2017): “Static and intertemporal household decisions,” *Journal of Economic Literature*, 55, 985–1045.
- CONNOLLY, S. AND M. GREGORY (2009): “The part-time pay penalty: earnings trajectories of British Women,” *Oxford Economic Papers*, 61, i76–i97.
- DOMEIJ, D. AND P. KLEIN (2012): “Should day care be subsidized?” *Review of Economic Studies*, 80, 568–595.
- DONNI, O. (2009): “A Simple Approach to Investigate Intrahousehold Allocation of Private and Public Goods,” *The Review of Economics and Statistics*, 91, 617–628.
- DUFLO, E. AND C. UDRY (2004): “Intrahousehold Resource Allocation in Cote d’Ivoire: Social Norms, Separate Accounts and Consumption Choices,” NBER Working Papers 10498.
- FOERSTER, H. (2018): “The Impact of Post-Marital Maintenance on Dynamic Decisions and Welfare of Couples,” Unpublished manuscript.
- GARCÍA-MORÁN, E. AND Z. KUEHN (2018): “Till mess do us part: Married womens market hours, home production, and divorce,” Unpublished manuscript.
- GAYLE, G.-L. AND A. SHEPHARD (2019): “Optimal Taxation, Marriage, Home Production, and Family Labour Supply,” *Econometrica*, forthcoming.
- GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2011): “Taxation and household labour supply,” *The Review of economic studies*, 79, 1113–1149.
- (2018): “Child-Related Transfers, Household Labor Supply and Welfare,” Unpublished manuscript.
- GUSTAFSSON, S. AND F. STAFFORD (1992): “Child care subsidies and labor supply in Sweden,” *Journal of Human resources*, 204–230.
- GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2015): “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?” NBER Working Papers 20913.
- HAMERMESH, D. (2002): “Timing, togetherness and time windfalls,” *Journal of Population Economics*, 15, 601–623.
- HARKNESS, S. (1996): “The gender earnings gap: evidence from the UK,” *Fiscal Studies*, 17, 1–36.
- 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, 209–244.
- LISE, J. AND K. YAMADA (2019): “Household sharing and commitment: Evidence from panel data on individual expenditures and time use,” *The Review of Economic Studies*, forthcoming.

- LUNDBERG, S. AND R. A. POLLAK (2003): “Efficiency in marriage,” *Review of Economics of the Household*, 1, 153–167.
- MANNING, A. AND B. PETRONGOLO (2008): “The Part-Time Pay Penalty for Women in Britain*,” *The Economic Journal*, 118, F28–F51.
- MARCET, A. AND R. MARIMON (2017): “Recursive contracts,” *Unpublished Manuscript*.
- MAZZOCCO, M. (2007): “Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment,” *Review of Economic Studies*, 74, 857–895.
- MEGHIR, C. AND L. PISTAFERRI (2004): “Income variance dynamics and heterogeneity,” *Econometrica*, 72, 1–32.
- RIBAR, D. C. (1995): “A structural model of child care and the labor supply of married women,” *Journal of Labor Economics*, 13, 558–597.
- SCHULTZ, T. P. (2005): “Fertility and income,” *Yale University Economic Growth Center Discussion Paper*.
- STEVENSON, B. (2007): “The impact of divorce laws on marriage-specific capital,” *Journal of Labor Economics*, 25, 75–94.
- TRAIN, K. (2008): “EM Algorithms for nonparametric estimation of mixing distributions,” *Journal of Choice Modelling*, 1, 40–69.
- UK GOVERNMENT, . (2017): <https://www.gov.uk/looking-after-children-divorce>.
- VOENA, A. (2015): “Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?” *The American Economic Review*, 105, 2295–2332.

APPENDIX

A Dynamics of K

Five-year transition matrices between the K categories $\{0, 1, 2, 10, 11, 20, 100\}$

All ages						
0.82	0.09	0.04	0.02	0.02	0.00	0.00
0.04	0.01	0.05	0.35	0.39	0.16	0.01
0.03	0.00	0.03	0.02	0.22	0.69	0.00
0.17	0.03	0.01	0.22	0.03	0.05	0.49
0.03	0.00	0.01	0.29	0.09	0.56	0.02
0.04	0.01	0.01	0.41	0.01	0.07	0.46
0.91	0.01	0.00	0.01	0.00	0.00	0.07
Ages 21-34						
0.46	0.31	0.15	0.02	0.06	0.00	0
0.05	0.02	0.08	0.28	0.48	0.10	0.00
0.07	0.01	0.02	0.03	0.32	0.54	0.01
0.05	0.14	0.03	0.25	0.21	0.11	0.21
0.04	0.01	0.02	0.18	0.24	0.52	0.00
0.03	0.06	0.00	0.56	0.00	0.15	0.21
0.50	0.00	0.00	0.50	0.00	0.00	0.00
Ages 35-44						
0.55	0.21	0.09	0.05	0.08	0.01	0.01
0.02	0.01	0.04	0.36	0.39	0.18	0.00
0.02	0.00	0.03	0.01	0.21	0.72	0.00
0.08	0.04	0.02	0.29	0.04	0.10	0.44
0.03	0.01	0.01	0.28	0.08	0.58	0.02
0.02	0.01	0.01	0.47	0.01	0.08	0.40
0.74	0.01	0.01	0.07	0.00	0.00	0.17
Ages 45-54						
0.93	0.02	0.00	0.03	0.01	0.00	0.00
0.05	0.00	0.01	0.52	0.17	0.20	0.03
0.01	0.00	0.00	0.02	0.07	0.90	0
0.22	0.00	0.00	0.19	0.01	0.02	0.56
0.03	0.00	0.00	0.36	0.05	0.52	0.04
0.03	0.00	0.00	0.36	0.01	0.05	0.55
0.91	0.01	0.00	0.01	0.00	0.00	0.07

Table A-1: Transition matrices between K values

Note that some of these transitions are estimated on small samples: in the age range 21-34, the transitions from categories 20 and 100 (two children 5-11 and one child 11-16, respectively) are calculated with denominators lower than 100. In the age range 35-44, the transitions from category 7, and in the age range 45-54, the transitions from category 20 (two children under 5) are also computed with small denominators.

The ergodic distribution of the transition matrix corresponding to all ages is fairly close to the distribution of the sample population across categories of K reported in Table 2, which makes sense since households mostly enter and exit the sample with no children under 16, i.e. the 0 category.

B Dynamics of wage quintiles

Five-year transition matrices between wage quintiles.

Full time				
0.34	0.29	0.15	0.14	0.08
0.14	0.27	0.24	0.25	0.09
0.08	0.14	0.21	0.40	0.17
0.05	0.05	0.09	0.39	0.41
0.06	0.04	0.04	0.13	0.74
Part time				
0.39	0.30	0.11	0.10	0.11
0.30	0.33	0.14	0.12	0.11
0.17	0.19	0.20	0.24	0.19
0.13	0.14	0.11	0.25	0.38
0.10	0.08	0.07	0.12	0.64
Non-participation				
0.77	0.13	0.04	0.03	0.03
0.11	0.80	0.04	0.03	0.02
0.08	0.08	0.70	0.10	0.04
0.05	0.05	0.06	0.70	0.13
0.08	0.04	0.03	0.07	0.79

Table B-1: Wage quintile transition matrices

C Dynamics of household labour supply

Five-year transition matrices between the household labour market choices

$\{(lm_m < FT), (FT, NP), (FT, PT), (FT, FT)\}$.

All ages			
0.31	0.20	0.21	0.28
0.06	0.46	0.36	0.12
0.05	0.12	0.56	0.27
0.05	0.11	0.21	0.64
Ages 21-34			
0.29	0.23	0.23	0.26
0.08	0.47	0.33	0.13
0.04	0.23	0.52	0.20
0.03	0.16	0.27	0.53
Ages 35-44			
0.29	0.21	0.24	0.27
0.05	0.43	0.41	0.10
0.03	0.11	0.61	0.25
0.04	0.10	0.24	0.61
Ages 45-54			
0.36	0.15	0.17	0.32
0.07	0.50	0.28	0.14
0.08	0.09	0.51	0.32
0.06	0.06	0.11	0.77

Table C-1: Household labour supply transition matrices

D Model fit

	$(lm_m < FT)$		(FT, NP)		(FT, PT)		(FT, FT)	
K	Data	Model	Data	Model	Data	Model	Data	Model
0	0.07	0.13	0.14	0.15	0.25	0.19	0.55	0.54
1	0.07	0.12	0.21	0.23	0.31	0.28	0.41	0.36
2	0.09	0.13	0.26	0.25	0.32	0.27	0.33	0.35
10	0.08	0.12	0.20	0.17	0.36	0.45	0.36	0.26
11	0.10	0.12	0.25	0.27	0.35	0.30	0.30	0.31
20	0.10	0.13	0.23	0.20	0.38	0.42	0.28	0.25
100	0.09	0.13	0.17	0.12	0.38	0.48	0.37	0.27
<hr/>								
<i>w_m</i>								
1	0.11	0.17	0.18	0.16	0.28	0.25	0.43	0.42
2	0.08	0.14	0.19	0.15	0.26	0.27	0.47	0.44
3	0.11	0.13	0.16	0.17	0.29	0.29	0.44	0.41
4	0.06	0.12	0.16	0.18	0.31	0.29	0.47	0.42
5	0.07	0.11	0.20	0.20	0.31	0.29	0.42	0.40
<hr/>								
<i>w_f</i>								
1	0.09	0.11	0.19	0.20	0.36	0.33	0.36	0.35
2	0.08	0.12	0.23	0.18	0.31	0.33	0.37	0.37
3	0.06	0.12	0.13	0.16	0.29	0.28	0.52	0.43
4	0.07	0.14	0.12	0.16	0.27	0.22	0.55	0.48
5	0.07	0.15	0.10	0.15	0.27	0.19	0.56	0.50

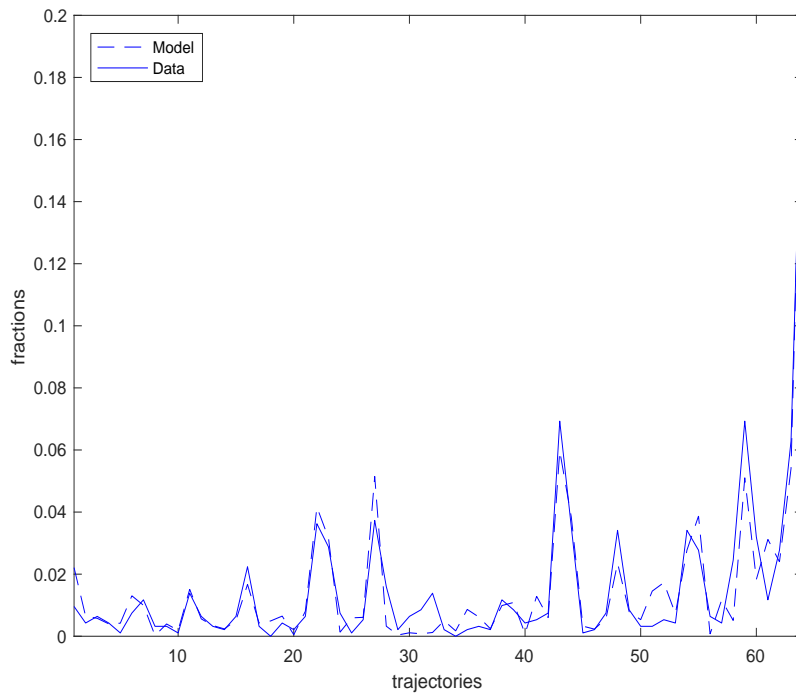
Note: NP: non-participation, PT: part-time, FT: full-time.

Table D-1: Fit of labour supply choices

Data			
0.28	0.18	0.24	0.29
0.06	0.44	0.36	0.14
0.04	0.12	0.58	0.25
0.04	0.12	0.23	0.61
Model			
0.25	0.21	0.28	0.25
0.11	0.40	0.47	0.02
0.13	0.10	0.54	0.23
0.12	0.15	0.22	0.51

Table D-2: Fit of labour supply transitions (all ages)

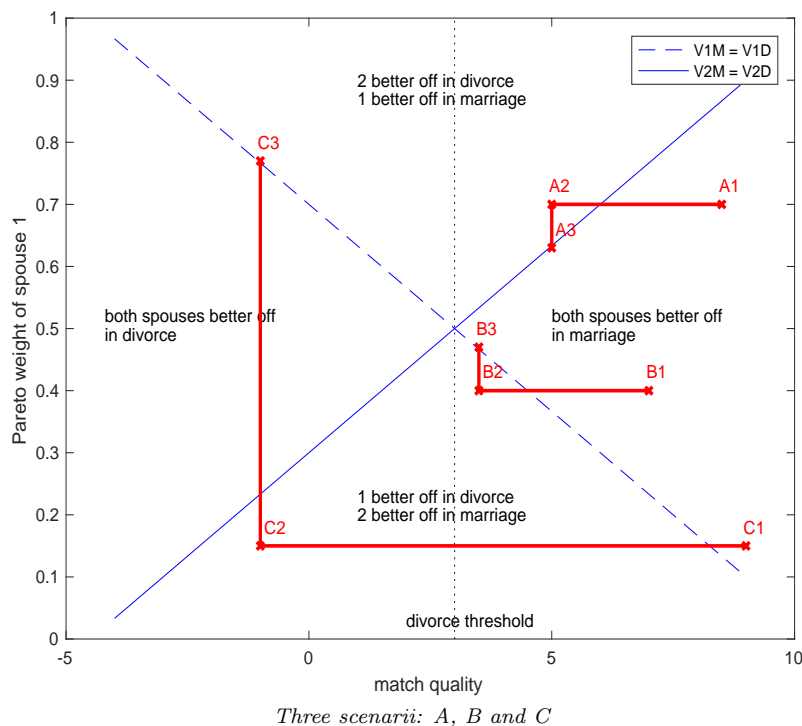
Figure D-1: Fit of labour market trajectories



All possible trajectories, indexed 1 to 64.

E Dynamics of Pareto weights

Figure E-1: Illustration of limited commitment model



The above figure illustrates the workings of the limited commitment model in our context. The lines showing the combinations of (ϵ, μ) which make either spouse indifferent between being in the marriage and divorce are straight only for illustration purposes, since the relationship $V_i^M(\mu) + \epsilon = V_i^D$ is not necessarily linear. Continuing marriages are located in the triangle on the right of the diagram, where both spouses are better off in the marriage than in divorce. The three red paths illustrate three possible scenarios.

In scenario A, the couple starts off in A1, then is hit by a negative match quality shock which takes the couple to A2. Here, spouse 2 would rather divorce than stay in this marriage at these values of (ϵ, μ) . She is thus in a position to force renegotiation towards point A3 where she is now indifferent between being in this marriage and divorce. This corresponds to a decrease in μ .

In scenario B, the reverse happens. When the negative shock happens (B2), spouse 1 has a credible threat to divorce and force renegotiation to increase μ and move the couple to B3.

In scenario C, the couple are initially in C1. The negative shock takes them to C2 where spouse 1 forces renegotiation towards C3. In C3, spouse 2 can force renegotiation since she would be better off in divorce. However, spouse 1 is only just indifferent between staying in this marriage and divorce so there is no room for any adjustment of μ downwards. The couple thus divorces. In fact, any value of the match quality ϵ below the divorce threshold line will yield a divorce.

Note that a positive shock to ϵ does not give rise to a renegotiation of Pareto weights since it increases both parties' surplus over divorce, giving no-one a credible threat to walk away from the marriage.