

Can the Quiet Revolution Continue?

The Effects of Childcare Subsidies on Female Labor Force Participation in a Life-Cycle Model

Jonas Nauerz

Abstract

The uninterrupted rise in U.S. female labor force participation throughout the 1970s and 1980s stalled in the 1990s and has fallen since. Women in other western countries, especially in Scandinavia, where childcare is heavily subsidized, are more likely to participate in the labor market. Even though working mothers in the U.S. face substantially higher childcare costs, they receive less public support. I build a structural, life-cycle model of heterogeneous households, family labor supply, and intra-household bargaining which allows me to examine whether increasing U.S. public spending on childcare to Scandinavian levels can promote growth in the women's labor supply. While financing larger public spending with higher payroll taxes has distortionary effects, raising child-related transfers reduces mothers' disincentives to work. I find that the policy increases long-run labor force participation among married women by 3.8 percent. Households with young children benefit substantially from the reform, but overall welfare falls as there are many households that lose marginally.

Keywords: Household Labor Supply, Childcare, Family Policy

JEL classification: E62, H24, H31, J13, J31

1 Introduction

Over the last 50 years, the U.S. labor market has been transformed by what [Goldin \(2006\)](#) calls the quiet revolution—the dramatic increase in the labor force participation of women, particularly married women. Data from the Panel Study of Income Dynamics (PSID) shows that from 1967 to 1997, hours worked by single and married women rose by 68 percent and 125 percent, respectively. This latter result meant that the fraction of double income households (DIH) among married couples increased substantially. The macroeconomic impacts of this revolution are sizable. [Heathcote et al. \(2017\)](#) estimate that half of the growth in U.S. earnings per capita from 1967 to the early 2000s came from an expansion in female labor supply.

A number of scholars have noted that the forces that drove this increase may have largely run their course as the female labor force participation rate (LFPR) has flattened and even reversed direction over the past 20 years—a trend commonly observed in most industrialized countries. Today, the level of U.S. female LFPR, especially among married women, is below OECD average. Despite facing higher income tax rates, women in Scandinavia have a higher LFPR than women in Continental Europe and the United States. One explanation for higher LFPR among women in Scandinavia in the face of strong disincentives to work is government spending on family policies. Child-related transfers in Scandinavia are much larger than in other industrialized countries: for instance, the share of public spending allocated to childcare subsidies as a fraction of GDP in Sweden is about ten times higher than in the United States. To make matters worse for working mothers, childcare in the U.S. is extraordinarily expensive. The net cost of childcare for a family with two earners¹ is more than 22 percent of average net income in the U.S.—almost twice the OECD average.

I build a structural model of family labor supply to test whether increasing public spending on childcare to Scandinavian levels can encourage growth in the LFPR in the United States. I consider whether such a policy is likely to extend the quiet revolution, and how changes to U.S. female LFPR relate to welfare. Ex ante, it is not clear how increased public spending on child-related transfers will affect women’s LFPR. Additional childcare subsidies reduce disincentives to work for mothers. On the other hand, additional payroll taxes increase disincentives to work for men and women.

To address these questions quantitatively, I simulate a life-cycle version of the standard heterogeneous-agent incomplete markets model. My model features a matching

¹measured as the cost of full-time care in a typical childcare center minus discounts and tax breaks

stage where some agents get married and negotiate a resource sharing rule and others stay single. Some households have children. During the active stage, individuals are subject to childcare expenditure shocks and idiosyncratic efficiency shocks and choose, i.a., whether or not to work. Finally, during the retirement stage, agents receive retirement benefits and only choose how much to consume and save. In order to preserve the marriage and enjoy the additional utility associated with it, the spouse with the highest wage is willing to transfer part of their earnings to their partner. Intuitively, a gender wage gap implies a lower outside option for women on average, thus the bargaining solution will likely favor the husband.

Since markets are incomplete, households save to finance retirement and to build a buffer stock of wealth that can be drawn upon in the event of adverse labor market shocks. Married households have a second shock absorber available: within-family insurance. If one partner experiences a negative labor market shock, their spouse is potentially able to compensate for the lost earnings by joining the labor force. However this option is less attractive for families with children due to significant childcare costs associated with the entry of a second partner into the labor force. The government provides subsidies to working households with young children which allows them to partially insure against this uncertainty as they only have to pay a share of the child-related costs out-of-pocket.

I find that expanding government subsidies for working households with children increases the participation of women by about 2 percent, with a greater impact on married women (+3.8%). Despite a 0.4 percentage point increase in the labor income tax (resulting in a slight decrease in the labor force participation of men), total participation rises by about 1 percent. While welfare for single women and married households with children increases by 2.3 percent and 1.6 percent, respectively, overall welfare falls. Given that single men and households with no childcare needs (who make up the majority of households in the model economy) do not benefit from the subsidy but have to bear the cost, this result is intuitive and highlights why it might be difficult to secure broad political support for such a reform.

There is a variety of potential reasons for the large time series changes in the female LFPR. [Goldin \(2006\)](#) argues that a cultural shift has produced a higher acceptance of the idea of a working mother and a change in female preferences have played an important role. Specifically, women today see their job more often as a career as opposed to work or a means to an end. Accompanied by a decline in the gender wage gap and a decline in marriage rates, women's bargaining power within the household has

strengthened over time. Knowles (2013) highlights that bargaining is critical for trends in female labor supply. Therefore, wage differentials, marriage, costs for childcare, and womens' within-household bargaining power play a central role in this study.

Attanasio et al. (2008) and Heathcote et al. (2017) study the rise in female labor supply in the United States and investigate what accounts for the increase in participation observed from the 1960s to early 2000s. Greenwood et al. (2016) develop a model of marriage, divorce, and married female labor force participation to study how technological progress in the household sector and shifts in the wage structure have shaped female labor supply. While these papers make important contributions by offering potential explanations for the observed time series changes in the female LFPR (and quantify their macroeconomic impacts), my study focuses on a different question and asks whether policy intervention can promote growth in women's labor supply today.

My work is related to an empirical literature documenting that subsidized day care encourages labor supply of mothers.² In one of the earliest contributions, Heckman (1974) finds evidence that subsidies increase female labor supply in the U.S. and more recently, Baker et al. (2008) find strong evidence of a shift into new childcare use after the introduction of subsidized childcare in Quebec. Using a structural model calibrated to Germany, Domeij & Klein (2013) find that a 50 percent subsidy for day care expenses leads to substantial increases in labor supply of mothers. Bick (2016) and Laun & Wallenius (2017) find similar results, emphasizing that a large fraction of part-time working mothers would work full-time if they had greater access to subsidized childcare. Related work by Guner et al. (2020), examines the macroeconomic effects of child-related transfers in the context of a life-cycle model calibrated to the U.S. economy. They find that an expansions of transfers that are conditional on market work have positive effect on female labor supply, while unconditional transfers have the opposite effects.

The literature on child-related transfers does not consider intra-household decision making and most papers (with the exception of Bick (2016)) abstract from income uncertainty. However, Blundell et al. (2016) show that female labor supply plays an important role in insuring households against labor market shocks. The contrast provided by my paper is that I take this idiosyncratic uncertainty into account. Since markets are incomplete, agents in my model cannot fully insure against idiosyncratic income and childcare shocks. Therefore, I can contribute to the literature by accounting

²See for a Olivetti & Petrongolo (2017) for a survey of studies on the effect of family policies on labor market outcomes.

for the welfare effects of the insurance mechanism implicit in the government support for working mothers. My model is tailored to the United States thus allowing me to study the expansion of existing programs in the U.S. in an environment with rich heterogeneity and where policies affect female labor supply decisions over the entire life cycle. Finally, I highlight the reality that household decisions are the outcome of a bargaining process, motivated by [Goldin \(2006\)](#) and [Knowles \(2013\)](#), who emphasize that intra-household bargaining is critical for understanding female labor supply.

2 Empirical Facts

In this section, I briefly document some key empirical facts about the female LFPR, childcare costs, and child-related transfers in the United States and other western countries. [Figure 1](#) plots the time series of LFP rates for single and married women in the U.S. age 16 years and over. For single women, the LFPR is roughly the same today as in the late 1970s. This series increases by about 9 percentage points then declines by the same amount over the next two decades. For married women, the LFPR increased by about 18 percentage points from 1975 to 1995, was flat for about 15 years, then has declined slowly since.³

Today, the level in the female LFPR in the U.S. is similar to some other western countries but lower compared to Scandinavian countries. [Figure 2](#) reports the LFPR for married women, net childcare costs for a DIH (measured as the costs of full-time care in a typical childcare center, net of rebates/tax reductions, as a share of average earnings), and public childcare-related spending as a share of GDP across countries in Europe and North America.⁴ The LFPR rate for married women in the United States ranks below average in the sample. [Fuchs-Schündeln & Bick \(2014\)](#) argue that differences in the tax treatment of married couples can account for a part of the differences in hours worked of married couples across countries, while [Rogerson \(2007\)](#), [Ragan \(2013\)](#), and [Wallenius \(2013\)](#) highlight differences in social security programs and family policies.⁵

For instance, child-related transfers in the U.S. are much smaller in relation to other developed countries: The U.S. spends less than 0.1 percent of output on childcare

³The decline in male LFP has been the subject of many recent studies, see e.g. [Acemoglu & Autor \(2011\)](#), [Krueger \(2017\)](#), [Tuzemen \(2018\)](#), and references therein.

⁴Source: [OECD Family Database](#), Table PF3.1 and [ILO Data for LFPR](#). [Figure 4](#) in [Appendix A.2](#) reports the LFPR of single women and net childcare costs for a single HH.

⁵According to the Current Population Survey in March 2020, married women with children younger than 5 years old are about 11 percent less likely to be part of the labor force compared to their childless counterparts, supporting the argument that factors unrelated to the tax code matter.

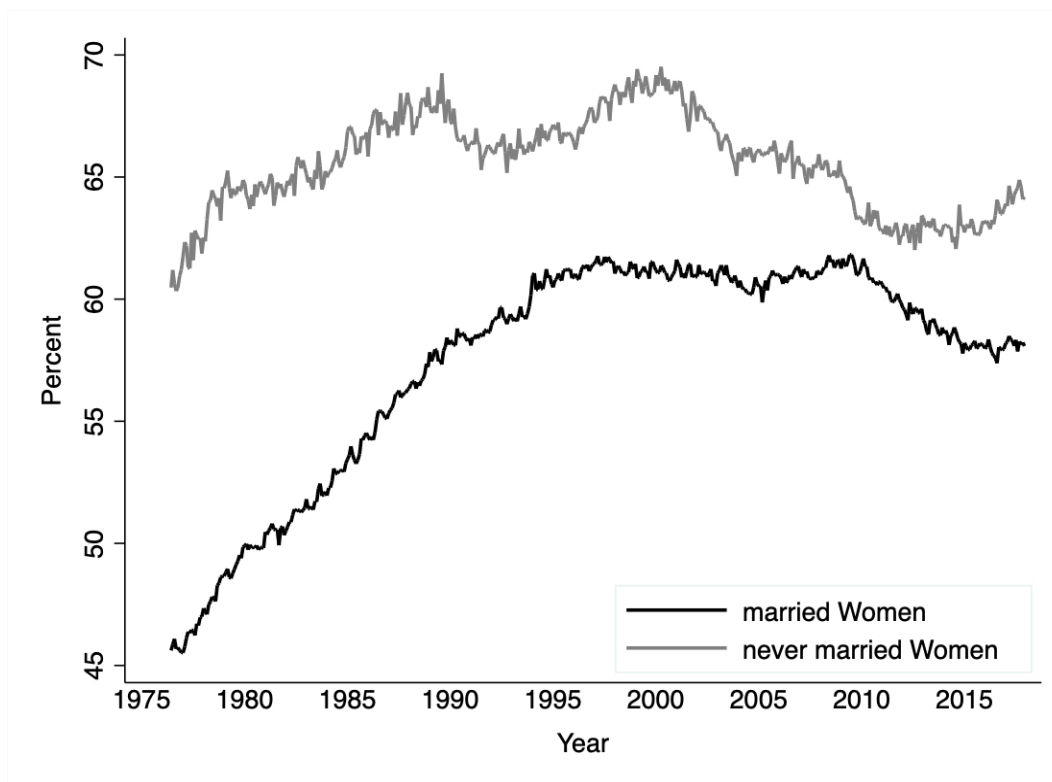


Figure 1: Female LFPR by Marital Status

subsidies (Sweden spends approximately 10 times as much) and subsidies per child in formal childcare amount to less than U.S. \$700 in 2015. Rogerson (2007) argues that differences in the form of government spending can potentially explain why both tax rates and hours worked are significantly higher in Scandinavia than they are in Continental Europe and the United States.

In addition to the less generous public support, childcare costs are much higher in the U.S. compared to other western economies. The net childcare cost to a dual-earner family in the United States is more than 22 percent of net income, compared to an average of 12.6 percent across the countries in the OECD (Thévenon 2011). The net childcare costs to a single household in the United States are even larger: Full-time care in a typical childcare center costs more than 50 percent of average earnings for a single household.⁶ Recent micro level data from Census' 2014 Survey of Income and Program Participation (SIPP) reveals that, conditional on paying for childcare, single households in the United States spend between roughly 22 percent and 26 percent (with single women spending a larger share) and married households approximately 15 percent of their household income on the service.

⁶See Figure 4 in Appendix A.2




Figure 2: LFPR Married Women, Childcare Costs, and Public Spending

To estimate the quantitative importance of childcare costs on female labor supply, I use my calibrated structural model. In that environment, I can assess whether changes in net childcare costs result in long-run increases in the participation of women and whether potential welfare gains for mothers outweigh the social costs of higher distortionary taxes.

3 The Model

The model, based on [Heathcote et al. \(2010, 2017\)](#), is a life-cycle version of a heterogeneous agent incomplete-markets framework in which agents are ex-ante heterogeneous only with respect to gender. I extend their work to incorporate expenditure shocks related to childcare similar to [Attanasio et al. \(2008\)](#) and [Guner et al. \(2020\)](#) to study child related transfers in a framework with rich heterogeneity and individual uncertainty. The model considers an economy populated by a continuum of finitely lived individuals with measure one. Time is discrete and a period represents one year.

Individuals enter the model at age 25 and either match with a spouse or stay single. Married households negotiate a resource sharing rule according to a Nash bargaining

framework. Both spouses gain utility from their marriage independently of gender which captures in reduced form both monetary and non-monetary benefits such as e.g. psychological, taxation, and legal benefits. I do not model fertility choices. However, I calibrate the arrival of children to approximate what is observed in the data. While this is a limitation of my study, it is not clear that including parental choices will significantly alter my quantitative findings. Childcare costs for young children are only a small fraction of the total monetary costs of children and changes in these costs due to an increase in governmental support are unlikely to affect the decision of a couple to have a child.⁷

Active individuals are subject to idiosyncratic efficiency shocks. They choose whether or not to work and how much to save. During retirement, agents enjoy their leisure and only consume and save until death. The assumption that agents either work full time or not at all is motivated by the well-known fact that hours worked are clustered around 2,000 hours (corresponding to a 40 hour work week) and zero hours in the data—a feature that standard utility functions cannot replicate (French 2005). Furthermore, Attanasio et al. (2008) find that extensive margin of labor supply is particularly important for women.

Before describing the bargaining and the decision problems in more detail, I highlight central assumptions and modeling choices.

3.1 Preferences and Technology

Individuals discount future utility using a common discount factor β . Agents' preferences over consumption and leisure are represented by

$$u(c^g, l^g; \psi^g) = \frac{(c^g)^{1-\gamma}}{1-\gamma} + \psi^g \frac{(1-l^g)^{1-\sigma}}{1-\sigma}, \quad (1)$$

where $l^g \in \{0, 1\}$ is the labor supply of an individual with gender $g \in \{f, m\}$ and ψ^g is a scale parameter that determines the relative value of leisure. The parameter γ measures the relative risk aversion and the parameter σ influences the Frisch elasticity of labor supply.

In the corporate sector, large firms operate in a perfectly competitive market. The production function for the representative firm is a constant return to scale Cobb-Douglas technology which takes as inputs the aggregate capital and efficiency units of

⁷For example, Bick (2016) finds negligible effects on overall fertility.

labor used in the corporate sector, i.e., the capital and labor supplied by agents in the economy. In the production process, male and female efficiency units are perfect substitutes. The production function for the representative firm of the corporate sector is

$$F(K, N) = K^\alpha N^{1-\alpha}, \quad (2)$$

where K is aggregate capital, $N = N^f + N^m$ the sum of the male and female efficiency units of labor used in the corporate sector, and α is the capital share. Firms in the corporate sector hire workers in a perfectly competitive labor market and rent capital from financial intermediaries.

3.2 Heterogeneity and Uncertainty

Similar to [Aiyagari \(1994\)](#), agents of the same gender are ex-ante identical and there is no aggregate uncertainty. However, ex post, agents are heterogeneous in terms of their marital status, child status, labor efficiency, and financial wealth. Markets are incomplete as there is a limit to how much agents can borrow (so they do not have access to a state-contingent insurance) and as a result, individuals accumulate assets to partially self-insure. Uncertainty enters the model through two shocks, the arrival of children in the beginning of an agent's working life and an idiosyncratic efficiency shock while working.

I estimate the stochastic processes for idiosyncratic productivity directly from the data. The individual endowment of efficiency units is denoted by

$$\epsilon(g, j, k, y) = \exp(L(g, j) + k + y_j^g) \quad (3)$$

and depends on the realization of the idiosyncratic labor productivity shock y_j^g at age j and $L(g, j)$, a gender specific deterministic quadratic polynomial in age.⁸ The indicator $k \in \{0, 1\}$ denotes whether an individual has children (with $k = 1$ indicating children).

In contrast to some of the literature, I allow that men and women face different experience profiles and stochastic processes for their idiosyncratic productivity.⁹ There

⁸In my model, for computational convenience and following [Heathcote et al. \(2017\)](#), I assume a labor market return to age rather than experience. Since most men participate in the labor market their entire active working life, this simplification will not matter for them. This might be different for women: while [Attanasio et al. \(2008\)](#) find small effects, [Olivetti \(2006\)](#) argues that increasing returns to experience have an effect on women's hours worked.

⁹Given that I only use data after the year 1997 to estimate the stochastic processes, the selection

is an extensive literature examining the explanations for gender differences in wages (see [Blau & Kahn \(2017\)](#) for an excellent review) and a growing literature that documents a large and persistent income penalty experienced by women after having children (e.g. [Angelov et al. \(2016\)](#), [Adda et al. \(2017\)](#), and [Kleven et al. \(2019\)](#)). Therefore, in the estimation of the productivity profiles, I explicitly take gender into account and whether an individual has children. An alternative approach is to calibrate initial human capital levels (and assume that human capital evolves endogenously) and that there is a permanent depreciation in human capital associated with non-participation, see e.g. [Attanasio et al. \(2008\)](#) and [Guner et al. \(2020\)](#).

Idiosyncratic productivity for single individuals follows an AR(1) process. As in [Bick \(2016\)](#), the idiosyncratic productivity of spouses in a married household is correlated,

$$\log(y_j^g) = \rho^g \log(y_{j-1}^g) + \xi_j \quad \text{where } \xi_j \sim N(0, \sigma_\xi^2), \quad \sigma_\xi^2 = \begin{bmatrix} \sigma_\xi^{f2} & \sigma_\xi^{f,m} \\ \sigma_\xi^{f,m} & \sigma_\xi^{m2} \end{bmatrix}.$$

Households face no uncertainty with regards to their marital status. I abstract from divorces in the model because the divorce rate in the United States has fallen from 4.5 in 1995 to 3 per 1,000 people in 2017.¹⁰ When married women face uncertainty about their future marital status, they have an additional incentive to participate in the labor market to accumulate human capital as an insurance, as [Peters \(1986\)](#) and [Sen \(2000\)](#) point out. However, even though the threat of divorce affects intra-household decision-making, the implication for female labor supply depends among other things on the regime of divorce laws ([Voena 2015](#)). She argues that even though women have an incentive to work while being married when they expect fewer resources, divorce laws that favor women may mute this incentive as women can expect more tangible assets as insurance upon divorce. Finally, it is not clear that the threat of divorce changes the effect of the policy reform considered in this study.

Some households have children. I assume that single men do not have children.¹¹ The arrival of children is exogenously determined after the marital status is decided. As in [Attanasio et al. \(2008\)](#) and [Guner et al. \(2020\)](#), households with young children in which all adults work (single parents and double income households with kids) face a fixed cost in their budget constraint which captures in a parsimonious way the idea that

bias for women is expected to be small compared to e.g. [Heathcote et al. \(2017\)](#) and [Blau & Kahn \(1997\)](#), who use longer time series.

¹⁰Source: [OECD Family Database](#) Chart SF3.1.C

¹¹In the 2014 wave of the Survey of Income and Program Participation, less than 5 percent of single (never married) men were identified as reference parent to any household children under 18 years old.

childcare is costly, i.e., single working mothers have to pay the fix costs and married households can avoid the cost in case one spouse stays at home.¹² I assume that only households with children younger than 6 years incur these costs. This is consistent with the fact that in the majority of states compulsory school attendance starts at the age of 6. Furthermore, in the 2014 SIPP, the share of households that paid for childcare was almost twice as large among households with children younger than 6 years compared to households with children age 6–18.

Let $\mathcal{M} \in \{S, M\}$ be the marital status and let $\Phi^S(l^g)$ and $\Phi^M(l^f, l^m)$ be the childcare costs for a single household and a married couple, respectively. Since only households in which all adults work face these costs, $\Phi^S(l^g) = \phi^S$, if $l^f = 1$ and 0 otherwise. Similarly, we have $\Phi^M(l^f, l^m) = \phi^M$, if $l^f = l^m = 1$ and 0 otherwise. Since the competitive price of childcare services is the wage rate w , the total cost of childcare for a single and married couple household with children is given by $w k \Phi^S(l^g)$ and $w k \Phi^M(l^f, l^m)$, respectively.

3.3 Timing

At the beginning of each period the labor productivity shock y is realized. After observing the outcome of the shock (or shocks, in case of married households) and conditional on their available resources, agents make their optimal consumption, labor supply, and savings choices for the current period. Firms decide how much capital to rent from a financial intermediary (who in turn makes zero profits) and labor to hire. The intermediary also pools savings at the end of each period and returns them proportionately to all agents who are still alive at the start of the next period, at actuarially-fair rates.¹³ After production takes place, firms sell their production and compensate their workers according to the prevailing wage per unit of effective labor and repay their loans. Finally, workers pay their payroll and capital income taxes.

3.4 Government

The government does not have access to lump sum taxation and instead finances itself with distortionary taxes on labor and financial income. It provides subsidies to working households with childcare responsibilities, pays out lump sum pension benefits and balances its budget every period. The size of the pension transfer depends on the

¹²Alon et al. (2020) for example model childcare as a time cost as agents can not work full-time while they care for children

¹³As Heathcote et al. (2017) point out, this assumption is unrealistic but allows to abstract from bequests (which is why it is commonly used in the literature).

household type, the computation of these values is discussed in more detail in the calibration section below.

Each household with children, married or single, receives a non-income-contingent subsidy of η percent of the childcare payments. As a result, effective childcare expenditures for a single household with children is given by $(1 - \eta) w \Phi^S(l^f)$. For a married couple household, the effective expenditures for childcare are given $(1 - \eta) w \Phi^M(l^f, l^m)$.

3.5 Matching and Intra-Household Bargaining

Individuals enter the economy at the matching stage. Some of them marry, others remain single. Marriage yields a utility gain M to each spouse. Individuals are selected in to marriage according to exogenous probabilities. Let μ^g be the probability of marrying for an individual of gender g . Since the measure of married men equals the measure of married women¹⁴, the following constraint must hold: $\mu^m = \mu^f = \mu$. Note, while the total measure of individuals is one (a measure 0.5 of men and a measure 0.5 of women), the total measure of households is given by: $\frac{\mu}{2} + (1 - \mu) = 1 - \frac{\mu}{2}$, which is equal to one when all individuals stay single and minimized (at 0.5) if all individuals were married. Matching (and thus the chance of marriage) only takes place once in an agent's lifetime as I abstract from divorce and marriages last until the couple dies together.

The resource sharing rule is determined as the outcome of a Nash bargaining problem. Denote by $\mathbb{M}^0(g, \kappa)$ the expected lifetime value for an individual of gender g who finds a spouse in the matching stage with sharing rule κ . Let $\mathbb{S}^0(g)$ denote the value for unmatched singles. The bargaining problem of a married household is given by

$$\begin{aligned} \max_{\kappa} \quad & [\mathbb{M}^0(m, \kappa) + M - \mathbb{S}^0(m)]^{\frac{1}{2}} \times [\mathbb{M}^0(f, \kappa) + M - \mathbb{S}^0(f)]^{\frac{1}{2}} \\ \text{s.t.} \quad & \mathbb{M}^0(g, \kappa) + M \geq \mathbb{S}^0(g) \quad \text{for } g \in \{m, f\}. \end{aligned}$$

Intuitively, a positive gender wage gap implies a lower outside option for women on average, thus a Nash bargaining solution will favor the husband. Therefore, the burden of childcare falls heavier on women as it is financially more attractive for the husband to work in the case that one spouse stays home.

¹⁴I abstract from same-sex marriages. Same-sex households make up only 0.8 percent of all U.S. households, while married same-sex couples made up about 59.4 percent of all same-sex households. Source: [U.S. Census](#)

3.6 Single Household's Problem

The value of staying single (that determines the outside option in the bargaining problem above) is the solution to the problem of a single household. The problem of a single working age household in recursive form is:

$$\begin{aligned} \mathbb{S}(g, j, k, a, y^{S,g}) &= \max_{\{c, l, a'\}} u(c, l; \psi^g) + \beta \pi^j \mathbb{E} \left[\mathbb{S}(g, j+1, k, a', y') \mid y \right] \\ \text{s.t. } c + \pi^j a' &= [1 + (1 - \tau^a)r] a + (1 - \tau^w) w l^g \epsilon(g, j, k, y^g) - (1 - \eta) w k \Phi^S(l^g) \\ a' &\geq \underline{a}, \quad a_{J+1} \geq 0, \quad c \geq 0, \quad l, k \in \{0, 1\}, \quad \Phi^S(l^g) = \begin{cases} \phi^S & \text{if } l^g=1 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

where π^j is an age dependant survival probability, \underline{a} a borrowing limit, and $w\phi^S > 0$ is the childcare cost incurred when being a working parent ($l^g = k = 1$). The inequality $a_{J+1} \geq 0$ implies that agents cannot die in debt. The outside option in the bargaining problem is then given by $\mathbb{S}^0(g) = \mathbb{E} \left[\mathbb{S}(g, 1, k, 0, y) \right]$, assuming that agents enter the model without wealth, which is consistent with the absence of bequests.

3.7 Married Household's Problem

The problem of a married working age household is a little bit more involved:

$$\begin{aligned} \mathbb{H}(j, a, k, y^m, y^f) &= \max_{\{c^m, c^f, l^m, l^f, a'\}} u(c^m, c^f, l^m, l^f; \kappa, \psi^m, \psi^f) \\ &\quad + \beta \pi^j \mathbb{E} \left[\mathbb{H}(j+1, a', y^{m'}, y^{f'}; \kappa) \mid y^m, y^f \right] \\ \text{s.t. } c^m + c^f + \pi^j a' &= (1 - \tau^w) w \left[\epsilon(m, j, k, y^m) l^m + \epsilon(f, j, k, y^f) l^f \right] \\ &\quad + [1 + (1 - \tau^a)r] a - (1 - \eta) w k \Phi^M(l^f, l^m) \\ a' &\geq \underline{a}, \quad a_{J+1} \geq 0, \quad c^m, c^f \geq 0, \quad l^m, l^f, k \in \{0, 1\}, \quad \Phi^M(l^m, l^f) = \begin{cases} \phi^M & \text{if } l^m=l^f=1 \\ 0 & \text{otherwise} \end{cases}, \end{aligned}$$

where $u(c^m, c^f, l^m, l^f; \kappa, \psi^m, \psi^f) = \kappa u(c^m, l^m; \psi^m) + (1 - \kappa) u(c^f, l^f; \psi^f)$.

Since agents do not work during retirement, the state space and the choice set are reduced. Otherwise, their decision problems are identical to working age individuals except for the fact that retired households receive retirement benefits $b^{S,g}$ and b^M that depend on the average earnings for each individual (and thus on gender).

3.8 Equilibrium

The representative firm in the corporate sector maximizes its profits, and as usual, the equilibrium wage is determined by the marginal product of labor in the corporate sector: $w = F'_N(K, N)$. I assume that economy is open and the interest rate r is internationally determined which implies a constant capital to aggregate effective labor ratio. Capital depreciates at the rate δ , thus the zero-profit condition of the financial intermediaries implies that the rental rate of capital is $R = r + \delta$.

An equilibrium is defined such that households take wages and the interest rate as given and the decision rules for consumption, labor supply, savings, and the associated value functions solve the household problem in [Section 3.6](#) and [Section 3.7](#) (and the analogous problems during retirement), respectively. For more details about the algorithm used to solve the household problem, see [Appendix A.1](#)

Capital and labor inputs are allocated optimally. The labor market clears, i.e., $N = L^m + L^f$, where aggregate labor supply is found by integrating the labor supply in efficiency units of each worker. The policy rules which solve the consumption-saving decisions, together with the labor supply decisions, and the transition probabilities of efficiency units induce a distribution of agents in this economy $\Lambda(j, a, k, y^g)$.

The government budget constraint satisfies

$$G + \int \left[(1 - \tau^l)B + \eta w k (\Phi^S(l^g) + \Phi^M(l^f, l^m)) \right] d\Lambda = \int (\tau^w w l^g + \tau^r r a) d\Lambda,$$

where the left-hand side represents the government's spending on transfers to agents which receive pension benefits $B = b^{S,g} + b^M$ or a childcare subsidy and on final consumption G , while the right hand side shows total tax revenues from payroll, and capital income taxes.

4 Calibration

In this section, I discuss the values I have assigned to the parameters in the model. I calibrate the model to mimic the U.S. economy. Since agents enter the model at age 25, the maximum model age, $J = 75$ implies that individuals die at the age of 100 with certainty. The age-specific survival rates π^j are taken from the U.S. Life Tables. The statutory retirement age is 67, thus $j_r = 42$. I follow [Domeij & Heathcote \(2004\)](#) and set the labor income tax rate to 27 percent. The ad-hoc borrowing constraint is

roughly 20 percent of mean annual after-tax income, thus $\underline{a} = -4.76$. For the technology parameters, I use NIPA data and set the depreciation rate δ to 0.06 and the capital share of output α to 0.33. [Table 1](#) summarizes the exogenous and [Table 4](#) the internal calibration.

4.1 Preference Parameters

The discount factor β , the coefficient of relative risk aversion γ , the relative weight on consumption versus non-market time ψ^g , as well as the parameter σ which affects both the extent to which within-household differentials in bargaining power and preferences for non-market time translate into differences in non-market time and the Frisch elasticity of labor supply¹⁵ are the preference parameters that need to be assigned values. I choose the value of the discount factor to avoid any life cycle effects on consumption. From the consumption Euler Equation, follows that consumption is stable over time if $\beta = \frac{1}{1+r(1-\tau_a)}$, where $r = 0.03$ is the annual interest rate and $\tau_a = 0.4$ is the withholding tax as in [Domeij & Heathcote \(2004\)](#). This implies an after-tax return to saving of 1.8 percent and hence, β is set to 0.982 in the model. The resulting effect is referred to as consumption smoothing, i.e. individuals' desire for a stable path of consumption, which implies that they shift their consumption from periods with higher income to periods with lower income. Because retirement, a low income period, follows the earning period, consumption smoothing leads to a humped-shaped age path of asset holding. For the remaining parameters, I follow [Heathcote et al. \(2017\)](#) and [Rogerson & Wallenius \(2013\)](#) and set $\gamma = 1$ and $\sigma = 0.5$, respectively. The relative weight on consumption for men ψ^m and for women ψ^f , are calibrated to match the average male and female labor force participation for individuals 19 and older in 2019 in the United States of 71.61 percent and 57.87 percent, respectively, as reported by the U.S. Bureau of Labor Statistics.

4.2 Idiosyncratic Efficiency Units

To estimate the efficiency units in [Equation \(3\)](#), I use data from the Panel Study of Income Dynamics. The PSID is a longitudinal study of a representative sample of U.S. individuals and the families in which they live. I focus on data from 1997–2017 in the main sample, the so-called SRC (Survey Research Center) sample. It includes 28,839 prime age households, and since it is representative of the U.S. population it does not require weights. I exclude observations before the year 1997 and for individuals not of

¹⁵The Frisch elasticity then is $\frac{1}{\sigma} \frac{1-l}{l}$

Table 1: Summary Table Exogenous Calibration

Parameter	Description (Source)	Value
$\{\pi^j\}$	age-specific survival rates (U.S. Life Tables)	see text
$b^{S,g}, b^M$	pension system (Social Security Administration)	see text
j^r, J	statutory retirement age, maximum age (SSA, U.S. LT)	42, 75
γ	inverse of IES (Heathcote et al. 2017)	1
σ	curvature parameter leisure (Rogerson & Wallenius 2013)	0.5
α	capital share of output (NIPA)	0.33
δ	depreciation rate (NIPA)	0.06
r	annual risk-free pre-tax interest rate 2010–2020 (OECD)	0.03
β	discount factor (see text)	0.982
τ^a, τ^w	capital, labor income tax rate (Domeij & Heathcote 2004)	0.40, 0.27
\underline{a}	borrowing constraint (see text)	-4.76

working age, to reduce sample selection issues. I also trim the income distribution to remove outliers. This is important as I use the variance of log wages as a measure of dispersion.¹⁶ Specifically, I restrict my sample to individuals between 25 and 55 years and trim the top and bottom 5 percent of the total household income and earnings distribution. I construct hourly wages as annual earnings divided by annual hours worked, excluding households with missing observations and those who receive the majority of their income from a business. I only include individuals that work at least 1,500 hours a year, my definition of full-time work. The final sample includes 24,452 male and 20,619 female individual-year observations.

Let $w_{i,j,t}$ be the hourly wage of individual i of age j in year t . I regress the natural logarithm of hourly wages on a year dummy, a dummy which is equal to one if individual i has children, and a quadratic polynomial in age:

$$\ln w_{i,j,t} = \beta_0 + \beta_1 \text{age}_{i,t} + \beta_2 \text{age}_{i,t}^2 + \beta_3 k_i + y_{i,j,t}.$$

This specification is consistent with Equation (3) in the structural model and thus I can use the regression residuals as an estimate of the stochastic labor productivity

¹⁶see Katz & Autor (1999) for a discussion of the importance of truncating the income data at the lower end of the distribution.

component as having kids ($k_i = 1$) is predetermined with respect to realizations of $y_{i,j,t}$.

Figure 3 plots estimated, deterministic life-cycle profiles for log wages of individuals with and without children, by gender. I find that men earn more than women and that mothers earn less than childless women. These estimates are consistent with findings by Angelov et al. (2016), Adda et al. (2017), and Kleven et al. (2019) who document a wage penalty associated with children. Between the age of 25 and 30, the level of wages for mothers is above that for childless women. By age 35, however, mothers experience a decline in wages which persists almost over the entire life cycle.




Figure 3: Life Cycle Profile of Wages

Using the estimates for $y_{i,j,t}$ from the previous regression, I estimate an AR(1) process for idiosyncratic productivity. The estimates for persistence, standard deviation, and co-variance of the idiosyncratic productivity of spouses are given in Table 2 below. Idiosyncratic productivity is more persistent for men but relatively less volatile than for women. As Attanasio et al. (2008) and Heathcote et al. (2010), I find that the co-variance of productivity shocks for spouses is positive.

Table 2: Idiosyncratic Productivity

	Men	Women
Persistence (ρ_ξ^g)	0.729	0.642
Standard dev. (σ_ξ^g)	0.363	0.417
Co-variance ($\sigma_\xi^{f,m}$)	0.149	0.149

4.3 Marriage and Children

According to the U.S. Census Bureau’s Current Population Survey (CPS), in 2019, 48.2 percent of all households were married.¹⁷ I use the utility gain M that each spouse receives in the Nash Bargaining framework to calibrate the model such that it matches the ratio of labor force participation rates of married and single women in the data.

I use data from the 2017 CPS June supplement, which contains information on the total number of live births, to calculate the fraction of households with children. The sample consists of 23,980 households. Using 40 to 44 years old single (never married) females, I calculate the fraction of women with zero live births, which is 41.9 percent. The corresponding fraction for married females is 8.5 percent. Therefore, the probability of receiving children as a single women is $\mathbb{P}(k = 1|g = f, \mathcal{M} = S) = 0.581$ and the probability for a married household to have children is $\mathbb{P}(k = 1|\mathcal{M} = M) = 0.915$.

To calibrate parameters related to childcare costs, I use data from the 2014 wave of the Survey of Income and Program Participation (SIPP)—a household-based survey consisting of a series of nationally representative panels, each containing information over a period lasting approximately four years. For comparison with the PSID sample, I exclude observations for individuals not in working age, to reduce sample selection issues and trim the income distribution to remove outliers and households with a majority of non-labor income. Table 3 shows the percentage of income paid on childcare by households where the youngest child is of age 5 or younger.

Using the share of income on childcare from the SIPP, I calibrate the childcare costs ϕ^S and ϕ^M such that, in the model, single mothers spend 24 percent of their net income on childcare and married households with kids 15 percent. The government’s final consumption expenditure G is chosen to satisfy the government’s budget constraint and the childcare subsidy, η , is calibrated to match U.S. public spending on child-related transfers as a 0.1 percent share of output, as reported by the OECD.

¹⁷Source: CPS [America’s Families and Living Arrangements](#)

Table 3: Households with children age < 6 years (SIPP)

	Single Female Households		Married Households	
	No Paid Care	Paid Care	No Paid Care	Paid Care
Observations	3,458	1,988	7,511	7,938
Income (\$ monthly)	3,813	3,980	5,864	7,703
Share of income on childcare	–	23.93%	–	14.51%

In the data, a majority of married households pays for childcare, while only 36.5 percent of single mothers pay for the service. Naturally, the question arises: who takes care of children in households that do not or cannot pay for childcare services? Many mothers rely on family members for childcare. Among single mothers who do not pay for childcare, almost 80 percent rely on relatives such as grandparents and about 18 percent care for the child while working. Interestingly, among married households that do not pay for childcare, the share that relies on relatives for childcare is also almost 80 percent.

Table 4: Summary Internal Calibration

Target Moment	Data	Model	Parameter
Ratio LFPR married/single women	0.905	0.879	$M = 7.2$
Avg male LFPR	71.6%	72.4%	$\psi^m = 0.12$
Avg female LFPR	57.9%	58.6%	$\psi^f = 0.59$
Childcare/income single	23.9%	23.5%	$\phi^S = 0.46$
Childcare/income married	14.5%	13.7%	$\phi^M = 0.63$
Public childcare spending	0.1%	0.1%	$\eta = 0.006$

4.4 Pension Benefits

The pension is a lump sum transfer conditional on household type. These values are computed as follows. I first simulate the economy and compute the model’s equivalent of average indexed monthly earnings (AIME) for each individual. Up to 35 years of earnings are needed to compute average indexed monthly earnings. Following the U.S. Social Security Administration, I sum earnings in the years with the highest earnings and divide the total amount by the total number of years. Using the 2020 Social Security formula for benefits, I then compute pension benefits for agents in the model.

In particular, I use the 2020 bend-points and replacement rates: 90 percent of AIME up to a first threshold (bend point) equal to 22 percent of average individual earnings, plus 32 percent of AIME from this bend point to a higher bend point equal to 133 percent of average earnings, plus 15 percent of the income exceeding the second bend point. For married households, I apply the so-called spousal benefit rule: I compute household benefits as the maximum between the sum of the two benefits and 1.5 times the higher of the two benefits. Finally, for each household type I compute the lump sum transfer as the average pension benefit of that type.

5 Policy Experiment

I use the model to simulate the impact of the expansion of childcare subsidies on female labor supply and welfare. The government supports households with childcare responsibilities by paying for a small share of the childcare expenditures (which means households must cover most child-related expenses out-of-pocket). Thus, before the reform, public spending on childcare subsidies is only 0.1 percent of GDP, as currently in the United States. In the initial equilibrium, roughly 59 percent of all women work, while the labor force participation rate among married women is 55 percent. Since markets are incomplete, agents cannot fully insure against the childcare shock and the related out-of-pocket expenditures. Strengthening the insurance mechanism implicit in the government support for working mothers, is therefore likely to increase welfare for this group. I consider the policy of raising public spending on childcare to levels observed in Scandinavian countries, i.e., increasing public spending tenfold. I assume that the government raises the childcare subsidy η such that public spending increases to 1 percent of GDP while leaving spending for general government consumption unchanged.

Currently, the United States government supports households with children using multiple programs: *i*) means-tested child care subsidies that are conditional on work through the Child Care Development Fund (CCDF), *ii*) universal (not means-tested) childcare credits through the Child and Dependent Care Tax Credit (CDCTC) which is also conditional on work, and *iii*) means-tested but unconditional child credits through the Child Tax Credit (CTC).¹⁸

The CCDF provides assistance to low income families by providing childcare vouchers so they can work or receive training (Lynch 2010). However, only a small portion of

¹⁸Furthermore, the Earned Income Tax Credit (EITC), which provides transfers for low income households through a means-tested wage subsidy.

eligible families actually receive the subsidy: [Blau & Tekin \(2007\)](#) state that, in 2000, the CCDF served less than 15 percent of eligible children and in 2017 about 1.3 million children were served.¹⁹ According to the Federal Interagency Forum on Child and Family Statistics, there were 73.3 million children age 0–18 (23.8 million age 0–5) in the United States in 2018.²⁰ Families that receive the subsidy, must make a co-payment which is increasing in income and varies by state. On average, the co-pay in 2017 was 6 percent of household income.²¹

The CDCTC provides working households with the opportunity to deduct a percentage of the amount of work-related expenses (up to \$6,000) paid to a care provider from their tax liability.²² The percentage depends on adjusted gross income and is between 35 percent for households with income of \$15,000 or less and 20 percent for households with an income of above \$43,000 ([Guner et al. 2020](#)). In order to receive the tax credit, both parents must work and the households needs to have a positive tax liability. As a result, households in the upper end of the income distribution receive the majority of the benefits, with an average of \$500 per receiving household ([Maag 2014](#)).

The CTC provides families with a refundable credit of up to \$1,000 per child which is independent of actual childcare costs. However, the CTC is income-dependent. The amount of the CTC begins to reduce or phase out at \$200,000 of income, or \$400,000 for married couples filing jointly.²³

While the model is not able to mimic all facets of the child-related transfers currently in place in the United States, the policy experiment captures the essential characteristics of parts of the U.S. governmental support programs for families, i.e., the policy experiment studies the effects of an significant increase in conditional childcare subsidies. Given the data on international labor force participation and the focus in the literature, a natural benchmark therefore are Scandinavian countries, where we observe both, strong governmental support for working mothers (and high income taxes) as well as high female labor force participation.

¹⁹Source: Office of childcare FY 2017 Final Data Table 1

²⁰<https://www.childstats.gov/americaschildren/glance.asp>

²¹Source: Office of childcare FY 2017 Final Data Table 17

²²Source: IRS [Child and Dependent Care Credit](#)

²³Source: IRS [The child tax credit](#)

5.1 Results

Table 5 summarizes the results of extending the conditional childcare subsidies for working households with young children. In my model, as a response to the policy change, some agents find it optimal to change their long-run labor force participation decision as the disincentive to work for mothers is reduced. As a result, the labor force participation rate for married women increases by 3.8 percent, a change entirely driven married women with children. Single mothers increase their labor supply by 0.8 percent which results in a 0.3 percent increase in overall labor force participation for single women. These results are in line with previous literature, though the quantitative estimate of the increase in the labor force participation rate for mothers is slightly smaller than what [Bick \(2016\)](#) finds.²⁴

Table 5: Expansion of Child-related Transfers

Moment	Before	After
Participation of married females	54.7	56.8
With children	54.5	56.7
Without children	56.9	56.9
Participation of single females	62.2	62.4
With children	61.9	62.4
Without children	62.5	62.4
Participation of females	58.6	59.7
Participation of males	72.4	72.1
Labor income tax rate	27.0	27.4
Childcare subsidy	0.6	5.2

Note: all values in percent

The government raises the labor income tax rate by 0.4 percentage points to finance the increase in child-related spending while keeping general government spending constant. After the reform, the government covers 5.2 percent of households' childcare costs (an almost tenfold increase) which implies that it now spends 1 percent of GDP on child-related transfers. The adjustment of the payroll tax depresses male labor supply by 0.4 percent while also slightly decreasing the labor force participation of single women without children. These declines do not offset the 1.8 percent growth in female labor force participation and thus overall participation appreciates. However,

²⁴Interestingly, a naïve policy experiment in which the government adjusts general consumption (which is not valued by households) instead of the payroll tax yields results very similar to [Bick \(2016\)](#)

the overall changes are fairly modest given that mothers with children younger than 5 years (the main benefactors of the reform) constitute only a small fraction of the entire population.

Unsurprisingly, the reform benefits households with children. Welfare for single mothers and married households with children increases by 2.3 percent and 1.6 percent compared to the pre-reform economy as they have to cover a smaller share of the child-related costs out-of-pocket. The government essentially provides a partial insurance against the otherwise uninsurable income shock. This insurance is welfare improving but comes at the cost of a higher distortionary tax. The improvements for mothers cannot compensate the cumulative costs of this distortion for households without children and as a result overall welfare falls.

6 Conclusions

In this paper, I evaluate the macroeconomic impact of the expansion of child-related transfers in a comprehensive dynamic model with several features that make it suitable for policy analysis. I find that the participation of women rises when the government increases its support for working mothers—despite the tax hike that is necessary to finance the reform. In terms of welfare, I find that those who gain (single mothers and married couples with children), gain substantially. However these improvements are not large enough to offset the cumulative loss of households who do not benefit from the additional governmental support. A general lesson from this result is that it might be difficult to secure broad political support for such a reform since the group that loses marginally constitutes a majority in the economy.

Limitations of my analysis include the fact that fertility decisions and childcare costs are treated as exogenous and couples cannot divorce. However, it is not clear that the inclusion of endogenous parental or divorce choices would alter my quantitative findings in a significant way. While the threat of divorce might increase the general incentive for women to work while being married, the interaction between the effect of divorces and the effect of the expansion of childcare subsidies on female labor supply is less obvious. Regarding fertility, child-related interventions that lead to higher participation rates are unlikely to change couples' decisions about having children. Childcare costs are only a small fraction of the lifetime cost of raising children and furthermore, any reduction in these costs will be partly offset by increases in the taxes necessary to finance the expansion of childcare subsidies. [Bick \(2016\)](#) finds supporting evidence for

this argument: the childcare subsidy expansion in Germany lead to a negligible change in the fertility rate.

Looking ahead, my theory emphasizes the positive gains from providing support and insurance for working mothers, a group that has suffered heavy job losses during the current recession induced by the Covid-19 pandemic and lockdown measures that followed. [Alon et al. \(2020\)](#) argue that women with children are hit hardest because their employment is concentrated in the service sector and many women are prevented from working due to the increased need for childcare caused by school and daycare closings. In future work, I would like to explore this issue by modelling household and female labor supply when the economy is hit by an aggregate shock and the ability of workers to combine work with childcare responsibilities depends on their occupation.

References

- Acemoglu, D. & Autor, D. (2011), Skills, Tasks and Technologies: Implications for Employment and Earnings, in O. Ashenfelter & D. Card, eds, 'Handbook of Labor Economics', Vol. 4 of *Handbook of Labor Economics*, Elsevier, chapter 12, pp. 1043–1171.
- Adda, J., Dustmann, C. & Stevens, K. (2017), 'The Career Costs of Children', *Journal of Political Economy* **125**(2), 293–337.
- Aiyagari, S. R. (1994), 'Uninsured Idiosyncratic Risk and Aggregate Saving', *The Quarterly Journal of Economics* **109**(3), 659–684.
- Alon, T., Doepke, M., Olmstead-Rumsey, J. & Tertilt, M. (2020), This Time It's Different: The Role of Women's Employment in a Pandemic Recession, NBER Working Papers 27660, National Bureau of Economic Research, Inc.
- Angelov, N., Johansson, P. & Lindahl, E. (2016), 'Parenthood and the Gender Gap in Pay', *Journal of Labor Economics* **34**(3), 545–579.
- Attanasio, O., Low, H. & Sanchez-Marcos, V. (2008), 'Explaining Changes in Female Labor Supply in a Life-Cycle Model', *American Economic Review* **98**(4), 1517–1552.
- Baker, M., Gruber, J. & Milligan, K. (2008), 'Universal Child Care, Maternal Labor Supply, and Family Well-Being', *Journal of Political Economy* **116**(4), 709–745.
- Bellman, R. (1956), 'Dynamic programming and lagrange multipliers', *Proceedings of the National Academy of Sciences of the United States of America* **42**(10), 767–769.
- Bick, A. (2016), 'The quantitative role of child care for female labor force participation and fertility', *Journal of the European Economic Association* **14**(3), 639–668.
- Blau, D. & Tekin, E. (2007), 'The determinants and consequences of child care subsidies for single mothers in the usa', *Journal of Population Economics* **20**(4), 719–741.
- Blau, F. D. & Kahn, L. M. (1997), 'Swimming Upstream: Trends in the Gender Wage Differential in 1980s', *Journal of Labor Economics* **15**(1), 1–42.
- Blau, F. D. & Kahn, L. M. (2017), 'The Gender Wage Gap: Extent, Trends, and Explanations', *Journal of Economic Literature* **55**(3), 789–865.
- Blundell, R., Pistaferri, L. & Saporta-Eksten, I. (2016), 'Consumption inequality and family labor supply', *American Economic Review* **106**(2), 387–435.

- Brumm, J. & Scheidegger, S. (2017), ‘Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models’, *Econometrica* **85**, 1575–1612.
- Deisenroth, M., Rasmussen, C. & Peters, J. (2009), ‘Gaussian process dynamic programming’, *Neurocomputing* **72**(7-9), 1508–1524.
- Domeij, D. & Heathcote, J. (2004), ‘On The Distributional Effects Of Reducing Capital Taxes’, *International Economic Review* **45**(2), 523–554.
- Domeij, D. & Klein, P. (2013), ‘Should Day Care be Subsidized?’, *Review of Economic Studies* **80**(2), 568–595.
- French, E. (2005), ‘The effects of health, wealth, and wages on labour supply and retirement behaviour’, *Review of Economic Studies* **72**(2), 395–427.
- Fuchs-Schündeln, N. & Bick, A. (2014), Taxation and Labor Supply of Married Women across Countries: A Macroeconomic Analysis, 2014 Meeting Papers 321, Society for Economic Dynamics.
- Goldin, C. (2006), ‘The quiet revolution that transformed women’s employment, education, and family’, *American Economic Review* **96**(2), 1–21.
- Greenwood, J., Guner, N., Kocharkov, G. & Santos, C. (2016), ‘Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation’, *American Economic Journal: Macroeconomics* **8**(1), 1–41.
- Guner, N., Kaygusuz, R. & Ventura, G. (2020), ‘Child-Related Transfers, Household Labour Supply, and Welfare’, *The Review of Economic Studies* **87**(5), 2290–2321.
- Heathcote, J., Storesletten, K. & Violante, G. (2017), ‘The macroeconomics of the quiet revolution: Understanding the implications of the rise in women’s participation for economic growth and inequality’, *Research in Economics* **71**(3), 521–539.
- Heathcote, J., Storesletten, K. & Violante, G. L. (2010), ‘The Macroeconomic Implications of Rising Wage Inequality in the United States’, *Journal of Political Economy* **118**(4), 681–722.
- Heckman, J. J. (1974), ‘Effects of child-care programs on women’s work effort’, *Journal of Political Economy* **82**(2), S136–S163.
- Judd, K. L., Maliar, L., Maliar, S. & Valero, R. (2014), ‘Smolyak method for solving

- dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain', *Journal of Economic Dynamics and Control* **44**(C), 92–123.
- Katz, L. F. & Autor, D. H. (1999), *Changes in the Wage Structure and Earnings Inequality*, pp. 1463–1555.
- Kleven, H., Landais, C. & Sogaard, J. E. (2019), 'Children and gender inequality: Evidence from denmark', *American Economic Journal: Applied Economics* **11**(4), 181–209.
- Knowles, J. A. (2013), 'Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply', *Review of Economic Studies* **80**(3), 1055–1085.
- Krueger, A. B. (2017), 'Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate', *Brookings Papers on Economic Activity* **48**(2 (Fall)), 1–87.
- Laun, T. & Wallenius, J. (2017), *Having It All? Employment, Earnings and Children*, Working Paper Series 2017:6, Uppsala University, Department of Economics.
- Lynch, K. E. (2010), *The Child Care and Development Block Grant: Background and Funding*, Report for Congress RL30785, Congressional Research Service.
- Maag, E. (2014), *Child-Related Benefits in the Federal Income Tax*, Brief 27, Urban Institute.
- MacKay, D. J. C. (1999), 'Comparison of approximate methods for handling hyperparameters', *Neural Computation* **11**(5), 1035–1068.
- Olivetti, C. (2006), 'Changes in women's hours of market work: The role of returns to experience', *Review of Economic Dynamics* **9**(4), 557 – 587.
- Olivetti, C. & Petrongolo, B. (2017), 'The economic consequences of family policies: Lessons from a century of legislation in high-income countries', *Journal of Economic Perspectives* **31**(1), 205–30.
- Peters, E. (1986), 'Marriage and divorce: Informational constraints and private contracting', *The American Economic Review* **76**(3), 437–454.
- Ragan, K. S. (2013), 'Taxes and Time Use: Fiscal Policy in a Household Production Model', *American Economic Journal: Macroeconomics* **5**(1), 168–192.

- Rasmussen, C. E. & Williams, C. K. I. (2006), *Gaussian processes for machine learning*, Adaptive computation and machine learning, MIT Press, Cambridge, Mass.
- Rogerson, R. (2007), ‘Taxation and market work: is Scandinavia an outlier?’, *Economic Theory* **32**(1), 59–85.
- Rogerson, R. & Wallenius, J. (2013), ‘Nonconvexities, retirement, and the elasticity of labor supply’, *American Economic Review* **103**(4), 1445–62.
- Scheidegger, S. & Billionis, I. (2019), ‘Machine learning for high-dimensional dynamic stochastic economies’, *Journal of Computational Science* **33**, 68–82.
- Sen, B. (2000), ‘How important is anticipation of divorce in married women’s labor supply decisions? an intercohort comparison using nls data’, *Economics Letters* **67**(2), 209 – 216.
- Thévenon, O. (2011), ‘Family policies in oecd countries: A comparative analysis’, *Population and Development Review* **37**(1), 57–87.
- Tuzemen, D. (2018), ‘Why Are Prime-Age Men Vanishing from the Labor Force?’, *Economic Review* (Q I), 5–30.
- Voena, A. (2015), ‘Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?’, *American Economic Review* **105**(8), 2295–2332.
- Wallenius, J. (2013), ‘Social security and cross-country differences in hours: A general equilibrium analysis’, *Journal of Economic Dynamics and Control* **37**(12), 2466–2482.

A Appendix

A.1 Solution Algorithm

The household decision problems are fairly standard finite-horizon dynamic programming problems. Using a recursion working backwards from the last period of life, I can find the value functions and decision rules for all individuals. The state of an individual is given by age, gender, assets, marital status, pension status, and endowment of efficiency units. An individual knows the state in the beginning of period and decides how much to consume, how much to save, and whether or not to work.

During the retirement stage, there is no labor market income risk and individuals only decide how much to consume and save. In the last period all assets are consumed since there is no bequest motive and death is certain after period J . The value function is passed on to the next step where the decision rule and the value function for period $J - 1$ (and similarly for all other periods) are calculated. In every period other than the terminal period, agents choose consumption to maximize their flow utility given the states and the expected continuation value. Given optimal consumption, I then use the budget constraint to determine the savings choice. Solving for the decisions during the active stage is computationally more challenging since the state space and the choice set are larger. For each household type and for each age, I need to keep track of assets and the idiosyncratic endowment of efficiency units. I use Chebyshev–Gauss quadrature methods for calculating expectations.

To solve for the steady state, I guess a vector of parameters and equilibrium wages, and solve the individual’s decision problems. The solution to these problems deliver decision rules and associated value functions as well as expected start of working life values M^0 and S^0 which are inputs for the matching problem. The matching problem then determines the optimal sharing rule for married households. At this point, I simulate the economy, compute cross-sectional moments of the model, and check whether or not the matching probability satisfies the consistency condition, capital and labor inputs are optimally allocated, domestic labor and goods markets clear, and if the government budget is balanced. If any of these conditions are not satisfied at the initial guess, I update my guess using multidimensional Newton-Raphson methods and simulate again, iterating in this fashion until a convergence criterion is reached.

A.1.1 Computational Challenges

Using global solution methods to solve models with substantial heterogeneity and large state spaces is computationally very expensive as the computation time and storage requirements increase exponentially as the dimensionality of the problem grows linearly. [Bellman \(1956\)](#) calls this problem the “curse of dimensionality”. Smolyak’s method of using sparse grids with global polynomials as basis functions can attenuate the curse to some extent ([Judd et al. 2014](#)). Using an adaptive sparse grid framework, [Brumm & Scheidegger \(2017\)](#) are able to solve models with kinks in the policy functions of up to 20 dimensions.

However, [Scheidegger & Bilonis \(2019\)](#) argue that Gaussian Process Regression (GPR), a grid-free method, is not only more powerful in alleviating the curse of dimensionality, but also more flexible (as it allows for function approximation on arbitrary domains) and, perhaps most importantly, much easier to implement.

A.1.2 Gaussian Process Regression

In the following section, I will briefly introduce GPR based on the book by [Rasmussen & Williams \(2006\)](#). GPR is a supervised machine learning method, which has the purpose of learning input-output mappings from an information source (e.g. empirical data or computer simulations), the so-called training data set. Generally, given the training data, the goal is to use Bayesian inference to make predictions for new inputs that have not been observed yet, i.e. to obtain a function that makes predictions for all possible input values, giving a higher prior probability to those functions that are considered to be more likely, e.g., because of smoothness properties.

[Rasmussen & Williams \(2006\)](#) define a Gaussian process (GP) formally as a collection of random variables, any finite number of which have a joint Gaussian distribution. In that sense, a GP is a generalization of the Gaussian probability distribution because it describes random variables which are functions instead of vectors, i.e., a GP can be considered a distribution over functions.

A GP is completely specified by its mean function and co-variance function

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}; \theta), k(\mathbf{x}, \mathbf{x}'; \theta)),$$

where the random variables represent the value of the function $f(\mathbf{x})$ at location \mathbf{x} and

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})], \quad (\text{A.1})$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}\left[\left(f(\mathbf{x}) - m(\mathbf{x})\right)\left(f(\mathbf{x}') - m(\mathbf{x}')\right)\right]. \quad (\text{A.2})$$

The parameters of the mean and the covariance function θ are known as the hyper-parameters. In practice, according to [Rasmussen & Williams \(2006\)](#), the prior mean function is usually set to zero and the co-variance function will be the squared exponential.

Given input data or so-called *training inputs* $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_n]$ and corresponding observations or *training targets* $\mathbf{y}_i = f(\mathbf{x}_i) + \epsilon$, where $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma_\epsilon^2)$, the goal is to use a GP and Bayesian inference to infer a model of the unknown function $f(\cdot)$ that generated the data:

$$p(f|\mathbf{X}, \mathbf{y}) = \frac{p(\mathbf{y}|f, \mathbf{X}) \times p(f)}{p(\mathbf{y}|\mathbf{X})}, \quad \text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{marginal likelihood}}$$

According to [MacKay \(1999\)](#), optimizing the hyper-parameters θ using maximum likelihood estimation will yield a model that fits the data but also rewards simplicity of the model. Given a model of the unknown function f , the goal is to predict function values for an arbitrary input \mathbf{X}^* .

A.1.3 Gaussian Process Dynamic Programming

[Deisenroth et al. \(2009\)](#) and [Scheidegger & Bilonis \(2019\)](#) are among the first in the literature to combine dynamic programming methods and Gaussian process regression. [Deisenroth et al. \(2009\)](#) coin the term Gaussian Process Dynamic Programming (GPDP). The key idea is to model the value function (or policy functions, or both) calculated in the dynamic programming algorithm by a Gaussian process.

In that sense, GPDP is similar to conventional parametric function approximation methods that model the value function with e.g. Chebyshev polynomials or radial basis functions. The advantage over these methods is that by using non-parametric GP's and Bayesian inference, one does not have to fix the model class before observing any data ([Deisenroth et al. 2009](#)). More importantly, since GPs learn a value function based on a small sample of training inputs, the GPDP algorithm does not directly suffer from the curse of dimensionality ([Scheidegger & Bilonis 2019](#)). As a result, the GPDP al-

gorithm allows us to compute global solutions to high dimensional economic models on arbitrary state spaces and perform parameter uncertainty quantification, in a self-consistent manner (Scheidegger & Bilonis 2019). In other words, since the algorithm can handle high-dimensional state spaces, traditional model calibration algorithms become obsolete as one can simply treat parameters (such as preference or technology parameters) as additional dimensions.

Algorithm 1 outlines the GPDP I use to solve the model in Section 3. The algorithm is based on Deisenroth et al. (2009). I do not, however, model the state–action value function $Q^*(\mathbf{x})$ as a GP, as proposed by the authors. Instead, I maximize the Bellman equation using conventional methods and continuous expectations. Computationally, this turns out to be more efficient when the continuation value does not include higher order expectation, as in my case.

Specifically, let \mathcal{X} be the set of states and \mathcal{U} the set of actions or choices. The central idea and key difference between a conventional dynamic programming algorithm and Algorithm 1 is to model the state-value function $V^*(\mathbf{x})$ as a Gaussian process. The training targets (observations) for the GPR are recursively determined by the dynamic programming problem itself.

In particular, define the state-value function as follows

$$V^*(\mathbf{x}) = \max_{\mathbf{u}} Q^*(\mathbf{x}) = \max_{\mathbf{u}} \left\{ r(\mathbf{x}, \mathbf{u}) + \beta \mathbb{E}_{\mathbf{x}'} [V^*(\mathbf{x}') | \mathbf{x}, \mathbf{u}] \right\},$$

where $Q^*(\mathbf{x})$ is the state–action value function, $r(\mathbf{x}, \mathbf{u})$ the reward/utility function and $\beta \in (0, 1)$ the discount factor. Let $\pi^*(\mathbf{x})$ be the policy rule which solves the optimal control problem.

To alleviate the curse of dimensionality, we learn the value function on a low-dimensional space (the training inputs) and then use the GP model to project a high-dimensional input onto it, and finally link the projection to the output. In practice this means that \mathcal{X} and \mathcal{U} contain substantially fewer elements than their counterparts in a conventional dynamic programming algorithm which results in a significant reduction in computational time. Learning the link function is computationally substantially easier than the original problem of learning a high-dimensional function (Scheidegger & Bilonis 2019).

Another advantage of modeling the state-value function V_j^* by \mathcal{GP}_v is that we can obtain a predictive distribution for $V_j^*(\mathbf{x}^*)$ for any state \mathbf{x}^* through Equation (A.1).


```

input :  $\mathcal{X}$ 
1  $V_J^*(\mathcal{X}) = r_{term}(\mathcal{X})$  // terminal reward
2  $V_J^*(\cdot) \sim \mathcal{GP}_{V,J}$  // GP model for  $V_J^*$ 
3 for  $j = J - 1$  to 0 do
4   for all  $\mathbf{x}_i \in \mathcal{X}$  do
5     for all  $\mathbf{u}_k \in \mathcal{U}$  do
6        $Q_j^*(\mathbf{x}_i, \mathbf{u}_k) = r(\mathbf{x}_i, \mathbf{u}_k) + \beta \mathbb{E}_{\mathbf{x}'} [V^*(\mathbf{x}') | \mathbf{x}, \mathbf{u}]$  // state-action value
7     end
8      $\pi_j^*(\mathbf{x}_i) \in \operatorname{argmax}_u Q_j^*(\mathbf{x}_i, \cdot)$  // optimal policy
9      $V_j^*(\mathbf{x}_i) = Q_j^*(\mathbf{x}_i, \pi_j^*(\mathbf{x}_i))$  // state value
10  end
11   $V_j^*(\cdot) \sim \mathcal{GP}_{V,j}$  // GP model for  $V_j^*$ 
12   $\pi_j^*(\cdot) \sim \mathcal{GP}_{\pi,j}$  // GP model for  $\pi_j^*$ 
13 end

output:  $\mathcal{GP}_V, \mathcal{GP}_\pi$ 

```

Algorithm 1: Gaussian Process Dynamic Programming

Similarly, we can obtain a predictive distribution for the optimal policy for any state $\pi_j^*(\mathbf{x}^*)$, which then can be used to predict optimal behavior of all agents when simulating and calibrating the model.

A.2 Figures

Figure 4 reports the LFPR for single women, net childcare costs for a single HH (measured as the costs of full-time care in a typical childcare center, net of rebates/tax reductions, as a share of average earnings), and public childcare-related spending as a share of GDP across countries in Europe and North America Source: [OECD Family Database](#), Table PF3.1 and [ILO Data](#) for LFPR.

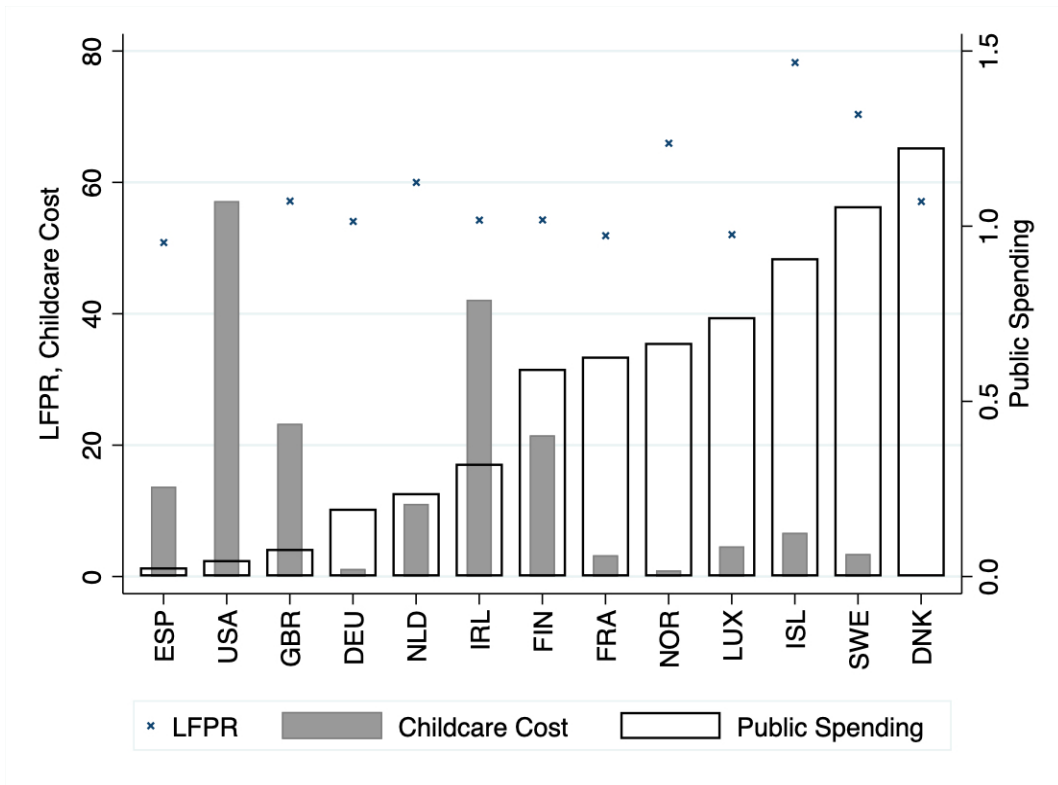


Figure 4: LFPR Single Women, Childcare Costs, and Public Spending