Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Review of Economic Dynamics 9 (2006) 557–587
Review of Economic Dynamics
www.elsevier.com/locate/red

Changes in women's hours of market work: The role of returns to experience

Claudia Olivetti *

Boston University

Received 30 September 2005; revised 5 June 2006

Available online 24 July 2006

Abstract

Over the past several decades, married women's hours of market work increased significantly in the US. I argue that changes in behavior by married women with children account for much of this change. In particular, the pattern of married women's work hours has changed substantially over the life cycle. In the past, married women in childbearing age tended to specialize in childrearing and home production activities at the expense of engaging in market work. Now they do not curb the hours they work in the market.

What factors contribute to this change in behavior? In this paper, I focus on relative changes in returns to experience as an explanation. I use PSID data for the 1970s and the 1990s to estimate the extent to which relative returns to experience have changed. I then use a life-cycle model with human capital accumulation and home production to quantitatively assess the consequences of this increase for married women's hours of work over the life cycle.

The estimates of the human capital production function show that women's marginal returns to experience increased by 25% across decades, whereas men's increased by only 6%. I show that this relative change accounts for 96% of the observed variation in married women's hours of work. Moreover, according to the model, the increase in returns to experience accounts for roughly half of the increase in the female/male wage ratio that is found in the data. I also show that a decline in the gender wage gap, holding returns to experience constant, accounts for only 18% of the total increase in hours of work. As a consequence, it cannot explain the change in the shape of women's life-cycle profiles. Although the focus of the analysis is the labor supply behavior of women, the model also allows predictions about the behavior of men and single women. These predictions are consistent with the data.

JEL classification: D1; J22; J24

* Fax: +1 617 353 4449.
E-mail address: olivetti@bu.edu.

1094-2025/$ – see front matter
doi:10.1016/j.red.2006.06.001

Keywords: Married women's labor supply; Learning-by-doing; Returns to experience; Child care

1. Introduction

An important trend in American society in the last several decades has been the dramatic increase in married women's hours of market work. This trend has been particularly prominent for married women with young children. Furthermore, it has been documented that the returns to labor market experience have increased more for women than they have for men between the 1970s and the 1990s. In this paper, I argue that these trends are deeply interconnected: the change in relative returns to experience is a key factor in accounting for the change in hours of work for married women with young children. The logic behind my argument is straightforward. In the 1970s, married women would temporarily cut back on market work during their child-rearing years. One cost of this withdrawal is the loss in accumulated labor market experience. Since this cost became greater by the 1990s, we would expect married women to decrease the extent of their reduction in market hours.

The contribution of this paper is twofold. First, I estimate the extent to which returns to experience have changed. Second, I use theory to assess the consequences of this increase for married women's hours of work over the life cycle. In particular, I build a life-cycle model with human capital accumulation and home production in which the basic unit of analysis is a married couple with children. Parents allocate their time between market, home, and leisure activities. Childcare is produced through a combination of parental time and market goods and services. Human capital is accumulated through learning-by-doing. As a consequence, wage profiles are determined endogenously.

Using this model, I use PSID data to analyze how the change in returns to experience affects age profiles of earnings and hours worked. The basic strategy is to parameterize the model using data from the 1970s and then ask how the model's predictions change when we allow for changes in the returns to experience to the values observed in the 1990s.[1] The values of the rates of return to labor market experience used in the calibration exercise are obtained by estimating a gender-specific learning-by-doing human capital production function for each of the two decades. These estimates are corrected for the sample selection bias that arises in this context, and are based on the PSID. I also consider a second experiment where, given the baseline economy parameters, the female/male gender differential is changed to its 1990s values.

The estimates of the human capital production function show that women's marginal returns to experience increased relatively more than men's did between the 1970s and the 1990s. In particular, women's returns to experience increase by 25%, whereas men's increase only by 6%. According to my simulations, the relative change in returns to experience can account for about 96% of the observed increase in hours worked for married women over their life cycle. The model over-predicts the increase in hours worked for older women. Moreover, consistent with the data, the model predicts only modest changes in the life-cycle profiles of hours worked for men and for single women. The model also has implications for the gender wage gap. Because wages are endogenous, increases in hours of work for women will also lead to higher wages. According to my simulations, the increase in returns to experience can account for roughly half

[1] This approach is similar to the one adopted in Regalia and Rios-Rull (2001) to investigate the causal relationship between the narrowing of the gender gap and the increase in single female households.

of the decrease in the gender wage gap between the 1970s and the 1990s. This result is consistent with the estimates obtained by O'Neill and Polachek (1993).

An alternative explanation for the increase in women's average work hours that is related to changes in the wage structure is the decline in the gender wage gap (Jones et al., 2003). I also use my model to evaluate the consequences of a 13% decline in the gender wage gap, holding returns to experience constant. While I find that this change can account for 18% of the actual increase in lifetime hours of work for married women, it cannot account for the large increase observed for married women with young children. This result has an intuitive explanation. Higher relative wages by gender should cause all women to work more, irrespective of their age. On the contrary, the change in the returns to labor market experience differentially affects women of childbearing age because it makes labor market interruptions early in the life cycle relatively more costly.

This paper contributes to a recent literature that studies the causes of the increased labor market participation of US women over the past century. One set of explanations stresses the importance of technological progress in the consumer durables sector and in medicine. According to Greenwood et al. (2005), the diffusion of new household technologies allowed women to decrease the time spent in home production and released time for market work. Goldin and Katz (2002) point out that the availability and diffusion of more reliable contraceptive methods, such as the pill, allowed women to plan their fertility and the timing of births, making it easier for them to plan a career. Other developments are the shift from manufacturing to services, which employ more women in jobs such as clerical work and sales (Goldin, 1990), changes in the wage structure that favored women (Pencavel, 1998, Jones et al., 2003), changes in the cost of children relative to lifetime earnings (Attanasio et al., 2004), and "female-biased" technological progress, which increased the relative attractiveness of working in the marketplace (Galor and Weil, 1996). Finally, the rise in the female labor force participation can also be explained in terms of the change in society's attitudes towards the role that women should play in the home and in the market place. On this topic, Fernández et al. (2004) examine the role of preferences formation within the family.

Since the focus of the analysis is the change in life-cycle behavior of married women with children, I abstract from modeling fertility and marriage decisions.[2] The study of the effects of a changing wage structure on these decisions is the focus of a recent paper by Caucutt et al. (2002). They use a dynamic general equilibrium model of family formation and investment in children to study the determinants of women's timing of births and labor supply decisions. They find that the rate of return to labor market experience is an important factor in the decision of women to delay child birth through its impact on wage growth. Their results are complementary to and reinforce the results reported in this paper.

The idea that the link between labor market participation and wages is important for understanding gender differences in earnings growth dates to Weiss and Gronau (1981), who identify the basis of this link in the process of human capital accumulation that occurs on-the-job. Subsequent papers estimate structural models of married women's labor force participation decisions in the presence of learning-by-doing human capital accumulation. Most notably, Eckstein and Wolpin (1989) construct and estimate one of the first dynamic models of married women's labor force participation decision in the presence of endogenous wage determination. Altug and Miller (1998) use a different methodology to estimate a model that incorporates both labor force partic-

[2] For a survey of life-cycle models of fertility, see Hotz et al. (1997). More recently Greenwood et al. (2003) built an overlapping generation model that incorporates both marriage and fertility decisions.

ipation and labor supply decisions in the presence of learning-by-doing and habit persistence in the agents' preferences for leisure.[3] This paper uses insights from this literature as the building blocks of its theoretical framework.

The paper is organized as follows. Section 2 documents the dynamics in hours of work by married women. I present the life-cycle model in Section 3. Section 4 describes the estimation of the learning-by-doing human capital production function that is used in the quantitative exercise. Section 5 discusses the calibration strategy. The results of the exercise and the sensitivity analysis are presented in Section 6. Finally, I provide some concluding remarks in Section 7.

2. Changing patterns in women's market work

One of the striking changes in the labor market over the last three decades has been the dramatic increase in hours worked for women. In particular, average hours worked per adult female increased by 43% between 1970 and 1990. At the same time, average hours worked per adult male remained essentially constant (they increased by 0.43%) between the two periods. Over this period, several characteristics of the female population changed substantially, such as educational attainment, fertility, marital status and out-of-wedlock childbearing. These facts have been extensively documented in the empirical literature.[4] The purpose of this section is to document that a key component of the increase in women's hours worked between 1970 and 1990 is the drastic change in the behavior of married women with young children. The extent of this change is reflected in the change in cross-sectional age profiles of hours worked for married women between 1970 and 1990.

Figure 1 uses Census data from Mc Grattan and Rogerson (1998) to plot the age profiles of weekly hours worked per person across decades. As Fig. 1 displays, in the 1970s, the age profiles of hours worked per person are double-peaked for married women. Women work more hours at the beginning of their adult life, then temporarily withdraw from the labor market while they have children, and eventually engage again in market activities as the children grow older. This is the "typical" shape of women's profiles, as first pointed out in the work on women's earning and labor force participation by Mincer (1974), and Mincer and Polachek (1974). As shown in Fig. 1, the shape of women's profiles is quite different in the 1990s. In particular, the average number of hours worked per woman increase for every age group, and this increase is largest for 25 to 44 year old women. As a consequence, the 1990 profile has a single-peaked shape, which typically characterizes men.[5]

The comparison between hours profiles for different demographic groups further clarify the extent of the change. To this effect, Fig. 2 compares Census data in 1970, panel A, and 1990, panel B, for married men and women, married mothers, and single women. Married women and married mothers of young children show a behavioral profile different from that of men and single women in the 1970s. In the 1990s, the gap between average number of hours worked by

[3] In both models, fertility choices are taken as given. Francesconi (2002) extend the Eckstein and Wolpin framework to incorporate fertility decisions. See Blundell and MaCurdy (1999) for an extensive review of the literature on dynamic models of labor supply.

[4] See Blau (1998) for a thorough account of the trends in labor participation and wages of US women between 1970 and 1995.

[5] Note that the data on average hours worked per person include changes in both the extensive and the intensive margin. The age profiles of labor force participation are almost flat (slightly decreasing) in the 1970s, but they are concave in the 1990s. The age profiles for average hours worked per worker (the intensive margin) are similar to the ones per person: more women are working in the 1990s, and they are working more hours conditional on their labor force participation.

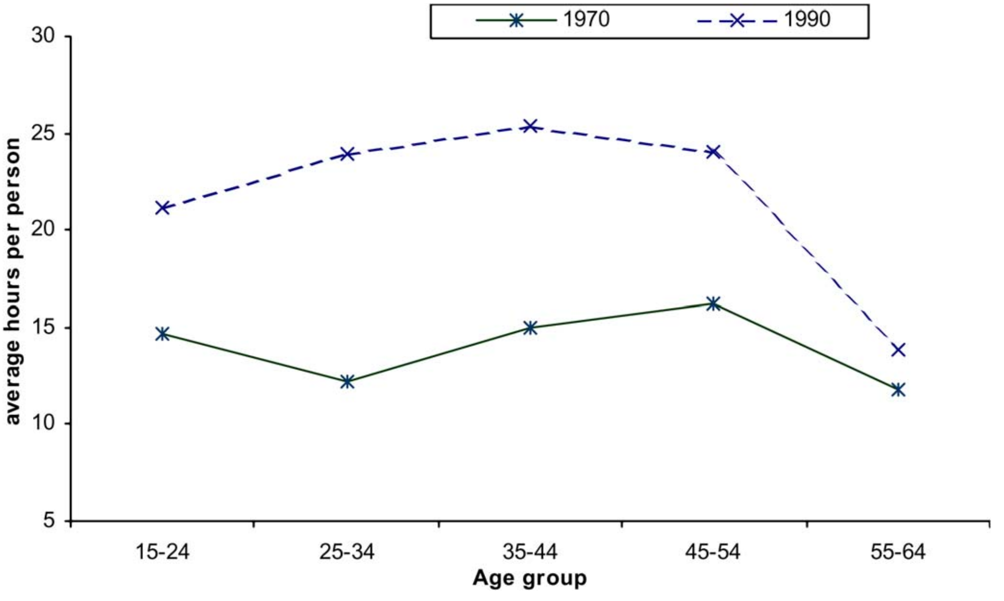

Fig. 1. Weekly average hours worked per married woman.

married women and by married men (or single women) decreases. The largest increase in work hours is experienced by 25–44 year old mothers with children of preschool age, with the average number of hours worked per person more than doubling. At the same time, for the same age group, average hours worked by single women and married men remained constant or decreased slightly for every age group. As a consequence, married women's and married mothers' age profiles have the same shape as men's and single women's profiles in 1990.[6]

Between 1970 and 1990 the demographic distribution of young women changed substantially. In particular, the fraction of married women with children of preschool age declined, whereas the fraction of single women in the population doubled.[7] Simple calculations show that the change in the demographic composition of the female population between 25 and 44 years of age can account for about 24% of the total increase in women's hours of market work. The increase in hours of market work by married women with children explains 70% of the residual. At the same time, the increase in hours of work by single women can account for only about 3% of the change.[8] This fact reinforces the observation that the increase in average hours of work by young women is due mainly to a change in the working behavior of married women with children.

Although the facts presented in this section are based on cross-sectional data, the same pattern is observed for life-cycle histories. Goldin (1989) uses historical data to show that women's labor participation at different stages in the life cycle started to change with the 1880s to 1910s

6 This drastic increase is still observed when we control for education. For instance, the largest increase in average hours worked per married woman with young children occurs for 25 to 34 year old high school graduates, with the average hours worked increasing by 185%. For the same demographic group, the average hours worked by college graduates increased by 120%.

7 In particular, the fraction of married women with children of preschool age declined by 30% for the age group 25 to 44, whereas the fraction of single women in the same age group more than doubled.

8 The percentages for single women with and without children are 2 and 4%, respectively.

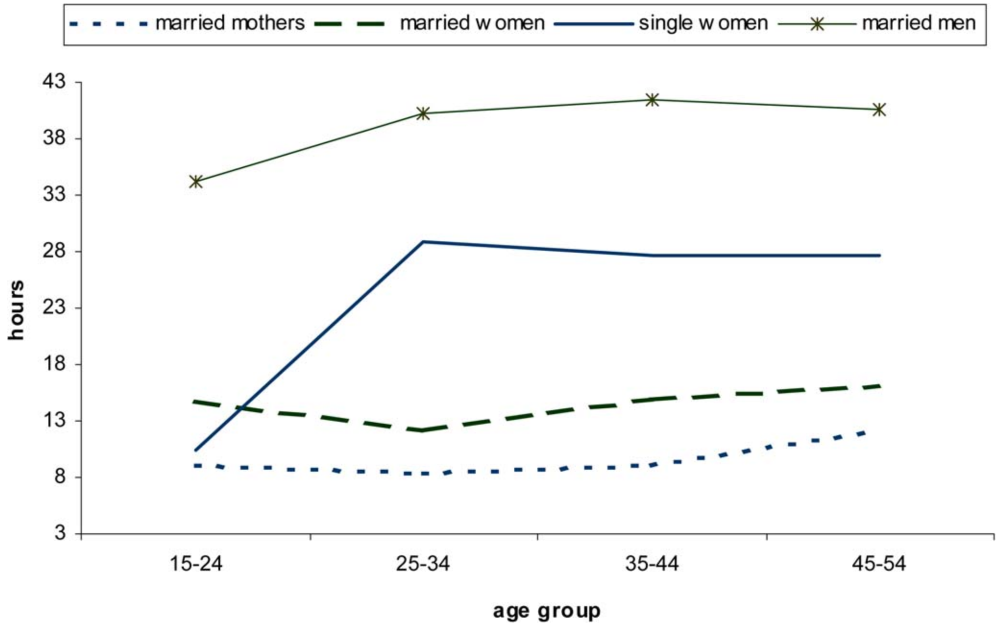

Panel A: Census 1970

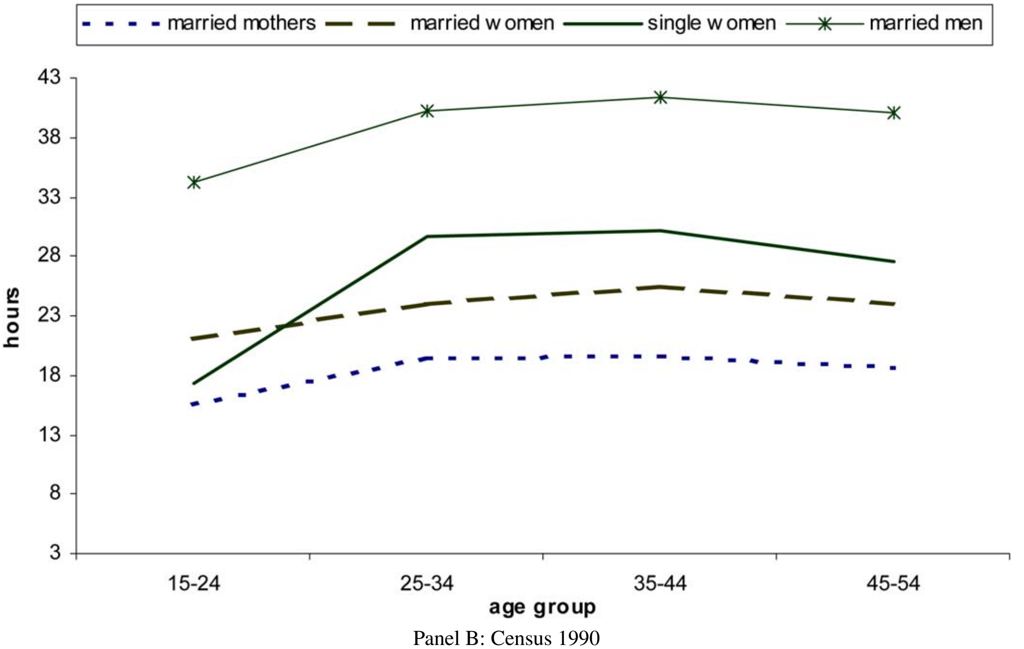

Panel B: Census 1990

Fig. 2. Weekly average hours worked per person.

cohorts. Mc Grattan and Rogerson (1998) use Census data for the years 1950–1990 to construct complete life-cycle profiles for several consecutive cohorts of men and women. Their data show that profiles for average hours worked are still double-peaked for women belonging to the 1926–1935 cohort, which corresponds to women between 35 and 44 years of age in 1970. The profile is distinctively single-peaked starting with the 1946–1955 cohort, that is women who are 35 to 44 years old in 1990. Attanasio et al. (2004) show a similar change in labor force participation profiles across two cohorts of women—the 1940s cohort and the 1950s cohort. This considerable time reallocation by age across cohorts is an important determinant of the change observed at the cross-sectional level.

The evidence reported in this section suggests that the 1990s pattern of work hours over the life cycle is not caused simply by women's postponing their decision to get married or have children. Rather, marriage and motherhood seem to play a much smaller role in women's labor decisions. In what follows, I show that the change in relative returns to experience is a crucial determinant of the observed change in the shape of work hours age profiles for married women.

Of course, additional factors might have contributed to this change—changes in women's schooling and fertility and the decline of the gender wage gap. Although these changes have contributed to the increase in overall women's labor force participation, they seem, at least to a first approximation, to play a lesser role in explaining the change in the life-cycle hours profile for married women. For instance, the change in women's profiles is observed for married mothers of children of preschool age and within education groups. Thus, the change in fertility and in educational attainment seems to play a secondary role. With regard to the decline in the gender wage gap, I will show later in the context of my model that although this decline induces an increase in average hours worked for every age group, it cannot explain the change in the shape of the work hours age profiles.

My argument considers one direction of causality—the effect of a change in wage structure that favors women on their labor supply decisions and on their accumulated labor market experience. It is also true, however, that the increase in women's wages can, at the same time, be explained by the increase in the average level of experience and education of the female workforce.[9] A few empirical studies have shown that the first direction of causality is quantitatively relevant. Smith and Ward (1985) show that the increase in real wages can explain about 60 per cent of the increase in women's labor force participation between 1950 and 1980. Their calculations take into account the change in fertility that also occurred in response to the wage rise. Pencavel (1998) finds that the increase in women's wages accounts for between one-quarter to one-half of the increase in women's labor supply between 1975 and 1994 (depending on the cohort).[10] Moreover, starting with the work by Katz and Murphy (1992), several researchers have documented how returns to education and experience, and within-group wage inequality, have been rising for both men and women over this period. This evidence indicates that these changes in the wage structure can be taken, at least to some extent, as exogenous with respect to women's prior work behavior.

[9] Alternatively, women might have entered occupations in which experience is more important, or they might face less discrimination in the labor market. See Blau et al. (2002) for a detailed discussion of these arguments and the empirical literature on this topic.

[10] Pencavel (1998) also finds that husband's wages have a very minor role in explaining a married woman's work decisions. This result is consistent with the work of Juhn and Murphy (1997), which shows that the slowdown of married men's earnings growth in the 1980s and 1990s cannot explain the increase in married women's labor force participation.

3. The life-cycle model

In order to quantify the contribution of the relative change in returns to experience to the change in married women's life-cycle profiles of hours worked, I develop a life-cycle model with human capital accumulation and home production where the basic unit of analysis is a married couple with children. Parents allocate their time between market, home and leisure activities and human capital is accumulated in the form of learning-by-doing. Childcare is produced through a combination of parental time and market goods and services. Thus, the parents' time allocation decision is key to the analysis. There is no uncertainty in the model, individuals have perfect foresight about the length of their lifetime. I assume price taking behavior. Individuals take as given both the rental rate of human capital, that differs by gender, and the real interest rate. Due to the learning-by-doing nature of the process of human capital accumulation, earning profiles are endogenously determined in equilibrium. The labor supply decision in one period not only affects current earnings, but also determines wages at later ages. This framework differs from other models of married women's life-cycle labor force participation with endogenous wages mainly on two dimensions: the specification chosen for the process of human capital accumulation, and the fact that labor supply and earnings are jointly determined over the life cycle both for husband and for wives.[11] In what follows I describe the model in detail and define the equilibrium.

Demographics The basic unit of analysis is a married couple. Fertility is exogenous. A couple is assumed to have two children (one male, one female) in the second period of life. Individuals live for six periods, where one period in the model corresponds to ten years. I do not explicitly model the first two (ten year) periods of agents' life when children live with their parents. They become adults in the third period as a married couple. Every couple dies at the end of period six.

According to this framework the adult life-span is forty years long (four 10-year periods). I would refer to adults as individuals in the age range 20 to 60. Hence, the first adulthood period corresponds to the age range 20–29, the second period (i.e. the parenthood period) corresponds to age 30 to 39, and so forth. Kids leave their parents household at age 20 and become adult as a married couple. Agents "die" at age 60.

Time allocation Each agent is endowed with one unit of time in every period of his/her life. In the first and fourth period he/she decides to allocate his/her time between market and leisure activities. In the second and third period, when children are present in the household, time is also allocated to the (home) production of childcare.

Preferences I assume that each couple has preferences over joint consumption, husband's and wife's leisure, and children's quality. Preferences are given by:

$$\sum_{t=1}^{N} \beta^{t-1}\big(U(c_t, \ell_{mt}, \ell_{ft}) + V(x_t)\big)$$

[11] For structural dynamic models of married women's labor force participation with endogenous wages, see Altug and Miller (1998), Eckstein and Wolpin (1989), and Moffitt (1984a).

where c_t represents family consumption, ℓ_{mt} is husband's leisure, ℓ_{ft} is wife's leisure, x_t represents children's "well-being" and β $(0 < \beta < 1)$ is the discount factor.[12]

Childcare production Child quality is produced using a combination of mother's and father's time and services purchased on the market. A key element in the production of childcare is the degree of substitutability between parental time and market services. More formally I assume:

$$x_t = X_t(h_{ft}, h_{mt}, s_t) \quad t = 2, 3$$

where h_{mt} is the fatherly time spent in childcare, h_{ft} represents motherly time spent in childcare and s_t is the input of market goods and services. As mentioned above, the production of childcare takes place only during the second and third period of a married couple's life. The function X is allowed to change over time. In particular, I assume that as children grow older the amount of parental time needed in the production of children's quality decreases whereas more market inputs are needed.[13]

Human capital accumulation I assume that agents accumulate human capital via a process of learning-by-doing. In particular, human capital next period depends on the amount of human capital accumulated up to the current period and on the number of hours worked in the market in the current period. That is,

$$\theta_{gt+1} = H_g(\theta_{gt}, n_{gt}) \quad g \in \{f, m\}$$

where n_{gt} is hours of market work by an individual of gender g at time t, and θ_{gt} is the human capital stock of an individual of gender g at time t. The human capital production function may differ for males and females. A married couple is indexed by its initial stock of human capital $(\theta_{m0}, \theta_{f0})$. In this framework the wage of an individual of gender g at time t (ω_{gt}) is given by the product of the individual current human capital stock (θ_{gt}) times the gender-specific rental rate of human capital (π_g) which is common across all agents.

Note that, as pointed out in the pioneering work by Weiss (1972), the learning-by-doing specification implies that past work experience has a direct effect on the determination of market wages. This consideration significantly changes the nature of the agents' labor supply decisions. In particular, it introduces a trade-off between the increase in utility that can be obtained by substituting market hours with home hours devoted to child care activities and the increase in future wages that can be obtained through on-the-job learning. As a consequence the current wage rate is no longer the only determinant of the return to working. There is an additional term that is included in the first order conditions for the optimal labor supply decisions. This term accounts for the increase in future wages resulting from the accumulation of human capital (i.e. experience) while working on the labor market, and is the driving force for the results reported in this paper.

[12] Here I consider a unitary framework where the two, symmetric, members of the household maximize joint utility under the assumption of income pooling. The main results of the analysis would not change if we were to consider an efficient non-unitary model of household decision making.

[13] Hotz and Miller (1988) make similar assumptions for the childcare production function. Also in Weiss and Gronau (1981) as children grow older women's productivity at home decreases.

Household decision problem Given prices π_g and R and the initial stock of human capital a household solves the following decision problem:

$$\max \sum_{t=1}^{N} \beta^{t-1}\big(U(c_t, n_{mt}, n_{ft}) + V(x_t)\big)$$

subject to:

$$\sum_{t=1}^{N} \frac{(c_t + s_t)}{(1+R)^{t-1}} \leqslant a_0 + \sum_{g\in\{m,f\}} \sum_{t=1}^{N} \frac{\omega_{gt} n_{gt}}{(1+R)^{t-1}}$$

$$\theta_{gt+1} = G(\theta_{gt}, n_{gt}) \quad g \in \{f, m\}$$

$$n_{gt} + \ell_{gt} + h_{gt} = 1 \quad g \in \{f, m\}$$

$$a_0 = a_{N+1} = 0$$

where $\omega_{gt} = \pi_g \theta_{gt}$ is the hourly wage of a worker of gender g at time t, π_g is the gender-specific efficiency wage, R is the rental rate of physical capital, and the adult life-span N is equal to 4 (10-year) periods. I assume that married couples begin their life with no asset and they consume everything they have during the last period of their life (no bequest).

3.1. Functional forms

This section describes the functional forms used in the simulation exercises.

Preferences As it is standard in the literature on labor supply, I assume that preferences are separable in consumption and leisure and also across time. In particular, I assume period utility functions of the form:

$$U(c_t, n_{mt}, n_{ft}, h_{mt}, h_{ft}) = \ln c_t - A\left(\frac{(n_{mt} + h_{mt})^{\alpha_m}}{\alpha_m} + \frac{(n_{ft} + h_{ft})^{\alpha_f}}{\alpha_f}\right)$$

$$V(x_t) = b \ln x_t$$

where $A > 0$ and $\alpha_g > 1$, $g \in \{f, m\}$. To be noticed, in this model $\sigma_g = 1/(\alpha_g - 1)$ cannot have the same interpretation as in a model with exogenous wage profiles since here the intertemporal elasticity of substitution changes over the life cycle. As mentioned above x_t represents children welfare. The parameter b represents the weight that parents give to the enjoyment of children welfare.

Childcare production The literature provides little guidance in modeling the childcare production function. There is no evidence on the shape of this function, although the degree of substitutability between maternal time and market produced goods and services should play an important role.[14] I assume a constant elasticity of substitution production function, since it provides a parsimonious representation of childcare production. I consider the following functional form:

$$x_t = \big(\gamma_t h_{ft}^{\rho} + (1-\gamma_t) s_t^{\rho}\big)^{1/\rho}$$

[14] Moffitt (1984b) explicitly consider the substitution possibilities between parental time and market goods and services in the production of childcare.

where $1/(1-\rho)$ is the elasticity of substitution between maternal time, h_{ft}, and market goods and services such as day care, s_t, in the production of childcare.[15] Weights on inputs of production, γ_t, are allowed to change over time. That is, as the children grow older we expect that a relatively bigger amount of market goods and services (i.e. college tuition and fees, etc.) as compared to parental time is used as an input in the production of children quality.[16]

The next section discusses the functional form chosen for the human capital accumulation equation and the estimation strategy used to estimate its parameters.

4. The human capital production function

A number of studies have focused on measuring the change in the rates of return to labor market experience by gender and its contribution to the decline in the gender wage gap between the mid 1970s and the late 1980s.[17] In particular, O'Neill and Polachek (1993) and Blau and Kahn (1997) show that both female and male returns to experience increased over this period. The estimates in these papers are obtained by running wage regressions for samples of men and women from the PSID. Both studies find that women's returns to full-time experience doubled, whereas men's returns did not increase as much over this period. According to Blau and Kahn's estimates, the extent of the increase for women is so large that women's returns to full-time experience are 25% higher than are men's in 1989.

Although both studies provide evidence that supports my hypothesis, there are two reasons why their estimates cannot be directly used to calibrate the human-capital technology parameters and their change over time in the context of my model. First, their specifications do not take into account dynamic aspects of the working decision, such as the strong complementarity existing between current hours of work and wages, in determining future wages over the life cycle that is at the core of the mechanism highlighted in this paper.[18] Second, these estimates do not take into account the problem of non-random selection into the sample. As it is well known in the literature, this problem introduces an upward bias in the estimates of earnings equations for women. As a consequence, the estimates tend to over-predict the growth in labor market returns for women. In my quantitative exercise this would lead to over-predict women's labor supply response to the change in labor market returns.

In order to circumvent these problems, I use PSID data to measure how the relationship between past wages and hours worked and current wages evolved between the 1970s and the 1990s. I use standard econometric techniques to obtain gender-specific, bias-corrected estimates of the

[15] We can easily consider a more general childcare production function that also include fatherly time. Although we would need to make assumptions on the degree of substitutability between motherly and fatherly time. In this case the function would generalize to:

$$x_t = \left(\gamma_t\left(h_{ft}^{\sigma} + h_{mt}^{\sigma}\right)^{\rho/\sigma} + (1-\gamma_t)s_t^{\rho}\right)^{1/\rho}$$

where $1/(1-\sigma)$ is the elasticity of substitution between motherly and fatherly time.

[16] Hotz and Miller (1988) show that the childcare intensity in motherly time declines as the children ages, $a_t = \delta^{t-1}a$. The assumption that $\delta \neq 0$ cannot be rejected. δ is significantly different from 1.

[17] These studies show that the increase in return to experience and the increase in actual experience for women explain a large portion of the decreasing gender wage gap over this period.

[18] The importance of past labor supply decisions and wages as determinants of current wage levels has been documented in the literature. Using PSID data, Altug and Miller (1998) find a significant learning-by-doing effect for married women. Imai and Keane (2004) document a similar finding for men based on NLSY data.

human capital technology parameters, and hence, of the returns to labor market experience, for the two decades.

I consider the following specification for the human capital production technology:

$$\theta_{gt+1} = (1-\delta_g)\theta_{gt} + \eta_g n_{gt}^{\psi_g}\theta_{gt} \qquad (*)$$

where θ_{gt} is the human capital stock of an individual of gender g at time t and n_{gt} is hours worked by an individual of gender g at time t, $g \in \{f, m\}$. Human capital depreciates at a constant rate δ. According to this specification, the stock of human capital at time $t+1$ is given by the previous period's human capital stock minus depreciation, plus an increasing function both of the individual's stock of human capital and of the total number of hours worked in the previous period. This functional form implies that current human capital is a sufficient statistic for an individual's past work history.[19] The marginal rate of return to extra hours of work is given by the first derivative of this function, and is proportional to $\eta\psi$. The parameter η can also be generalized to include a pure age-effect. This assumption would capture the fact that incentives to accumulate human capital may change as individuals age. The parameters of the production function are estimated separately by gender for the 1970s and the 1990s.

Note that I assume the human capital technology to be linear in previous period human capital stock, θ_{gt}. Under this assumption, the rate of growth of human capital does not depend on the level of human capital stock in the previous period.[20] As discussed by Weiss and Gronau (1981), this assumption makes it possible to separate the effect of labor supply on the level and the rate of growth of human capital accumulation. An analogous human capital production function is estimated by Imai and Keane (2004) in the context of a structural dynamic model that allows for both human capital accumulation and savings.[21] Their estimates are based on a sample of prime-age men during the 1980s and the early 1990s. Unfortunately, given the focus of my analysis, their estimates cannot be used in my calibration exercise.

4.1. Estimation methodology

Since the human capital stock is not observed, Eq. $(*)$ cannot be estimated directly. However, we do observe wages. As is typical in the literature on human capital, wages can be defined as the product of the individual human capital stock times the rental rate of human capital in the economy (i.e. the efficiency wage). That is, the wage rate for the ith individual of gender g is given by: $\omega_{igt} = \theta_{igt}\pi_{gt}$, where the rental rates, π_{gt}, differ by gender. In what follows the subscript g is dropped for expositional purposes. However, the parameters of the production function are allowed to vary by gender and across decades in the estimation.

The empirical counterpart of Eq. $(*)$ is obtained by multiplying both terms of the equation by π_{t+1}, dividing by π_t, substituting $\omega_{it} = \theta_{it}\pi_t$, taking logs and rearranging. We obtain:

$$\ln\omega_{it+1} - \ln\omega_{it} = \ln\left((1-\delta) + \eta n_{it}^{\psi}\right) + \ln\frac{\pi_{t+1}}{\pi_t} + \varepsilon_{it+1} \qquad (**)$$

[19] Alternatively, as in Altug and Miller (1998), it could also be assumed that today's human capital is determined by the individual's entire past history of hours worked on the market. Unfortunately, the data available do not allow us to follow an individual over his/her entire life cycle. This specification circumvents this problem.

[20] Imai and Keane (2004) and Altug and Miller (1998) also make a similar assumption.

[21] Shaw (1989) estimates a similar life-cycle model for prime-age men under the assumption of translog preferences and of a quadratic law of motion for human capital.

where $\ln(\pi_{t+1}/\pi_t)$ represents the rate of growth of rental rates of human capital and ε_{it+1} is an individual-specific productivity shock that is assumed to be correlated over time. The individual-specific productivity shock can be interpreted as a health shock, or as a job-related shock that changes an individual's chance to accumulate skills on the market. In this context, an individual's wages are observed only if he/she participates in the labor market, and the labor supply decision is one of the inputs in the production of human capital. Both decisions may depend upon individual-specific productivity, ε_{it+1}. This generates a selection bias in estimating the parameters of the production function. The selection bias arises because only individuals characterized by a higher level of productivity will be on the labor market. Moreover, in this dynamic framework, an individual is included in the sample only if he/she worked for two consecutive periods. This implies that the conditional expectation of $\ln\omega_{it+1} - \ln\omega_{it}$ (conditional on past and current period participation and information available at time $t+1$) includes the term $E[\varepsilon_{it+1} \mid \chi_{it} = 1, \chi_{it+1} = 1]$ where χ_{ij} is an indicator function that equals 1 if the individual worked at time j, $j = t, t+1$. Hence the sample selection term is a function of two indices as opposed to the standard, single index model that dates back to Heckman (1974). Following the discussion in Vella (1998) I use a simple modification of the Heckman (1979) two-step selection correction model to estimate the human capital production function in this case.

An accurate description of the estimation methodology requires additional notation and parametric assumptions on the structure of the error terms. Let us define $y^*_{it+1} = \ln\omega_{it+1} - \ln\omega_{it}$ as the latent endogenous variable with observed counterpart y_{it+1}—the wage growth between two consecutive periods. Let us also define the latent variable that captures sample selection as $\chi^*_{ij} = Z'_{ij}\gamma_j + v_{ij}$ where Z_{ij} is a set of worker's characteristics for individual i at time j, $j = t, t+1$, γ is a vector of unknown parameters and v is a zero mean error term that is correlated with ε. χ^*_{ij} is associated with the function $\chi_{ij} = I(Z'_{ij}\gamma_j > -v_{ij})$ where $I(\cdot)$ is an indicator function that is equal to 1 if $\chi^*_{ij} > 0$ and is equal to 0 otherwise. The two functions, χ_{it} and χ_{it+1} reflect whether the primary dependent variable is observed. The four variables are related in the following way: $y_{it+1} = y^*_{it+1}\chi_{it}\chi_{it+1}$, that is, the wage growth between two consecutive periods is observed only if the individual worked in both periods. The method of estimation of Eq. (**) in this case relies crucially on the assumption made about the relationship between the error terms v_{it} and v_{it+1}. I assume that the stochastic disturbances v_{it} and v_{it+1} are jointly normally distributed and denote by ρ the correlation coefficient between the two variables.[22] Under this assumption we can re-write the control function as:

$$E[\varepsilon_{it+1} \mid \chi_{it} = 1, \chi_{it+1} = 1] = \frac{\phi(Z'_{it}\gamma_t)\Phi\left(\frac{Z'_{it}\gamma_t - \rho Z'_{it+1}\gamma_{t+1}}{\sqrt{1-\rho^2}}\right)}{\Phi_2(Z'_{it}\gamma_t, Z'_{it+1}\gamma_{t+1}, \rho)} + \frac{\phi(Z'_{it+1}\gamma_{t+1})\Phi\left(\frac{Z'_{it+1}\gamma_{t+1} - \rho Z'_{it}\gamma_t}{\sqrt{1-\rho^2}}\right)}{\Phi_2(Z'_{it}\gamma_t, Z'_{it+1}\gamma_{t+1}, \rho)}$$

where ϕ and Φ denote, respectively, the density and cumulative distribution functions of a standard normal distribution and Φ_2 denotes the bivariate normal distribution. I will refer to the two

[22] This amounts to assuming that the stochastic disturbances $\{v_{ij}\}_{j=1}^{T}$ follow an AR(1) process. Alternatively, one could use semi-parametric procedures to correct for the bias (see Vella, 1998 for a survey.) This exercise, although interesting, is outside the scope of this paper.

correction terms as to $\lambda_1(Z_{it}; \gamma_t, \rho)$ and $\lambda_2(Z_{it+1}; \gamma_{t+1}, \rho)$, respectively. They are the bivariate equivalent of the inverse Mills' ratio in a standard Heckman two-step selection correction model.

I then follow the two-step procedure to adjust the estimates of Eq. $(**)$ for the sample selection bias. In the first step I estimate the two correction terms based on a probit specification.[23] The resulting terms, $\lambda_1(Z_{it}; \hat{\gamma}_t, \hat{\rho})$ and $\lambda_2(Z_{it+1}; \hat{\gamma}_{t+1}, \hat{\rho})$, are added to Eq. $(**)$ to obtain:

$$\ln \omega_{it+1} - \ln \omega_{it} = \ln\big((1-\delta) + \eta n_{it}^{\psi}\big) + \ln \frac{\pi_{t+1}}{\pi_t} + \lambda_1(Z_{it}; \hat{\gamma}_t, \hat{\rho}) + \lambda_2(Z_{it+1}; \hat{\gamma}_{t+1}, \hat{\rho}) + \eta_{it+1} \qquad (***)$$

where η_{it+1} are i.i.d. standardized Normal stochastic disturbances. In the second step equation $(***)$ is estimated by least squares.

I estimate $(***)$ by gender and by decade, the 1970s and the 1990s, using PSID data. The 1970s samples, one for men and one for women, include individuals from the 1970 to 1977 waves of the PSID. The 1990s samples include individuals from the 1990 to 1997 waves. The four samples include men and women aged 20 to 60. Summary statistics are reported in the Appendix. Real hourly wages, ω_{it}, are obtained by dividing annual real earnings by annual hours worked, n_{it}.[24] The year-to-year variation in the rates of return to human capital, π_t, is captured by introducing time dummies in $(***)$.

4.2. Estimation results and the change in returns to experience

In this section I discuss the estimation results and document the increase in returns to labor market experience between the 1970s and the 1990s. I estimate Eq. $(***)$ under the assumption that the depreciation rate δ and the exponent ψ are the same across genders and decades. Thus the variation in returns to experience, across genders and over time, will be captured by the estimated values for the parameter η.

The estimates must be constrained in order to generate numbers that are consistent with interior solutions for both men's and women's working hours in the calibration exercise.[25] I also specify $\eta(\cdot)$ as a quadratic function of age: $\eta(a_{it}) = \eta_0 + \eta_1 a_{it} + \eta_2 a_{it}^2$.[26] Both assumptions are needed to obtain realistic age-earning profiles when I use the estimates of the human capital production function to calibrate the model.

Table 1 presents the results when I set $\psi = 0.4$ and $\delta = 0.2$, and estimate the remaining parameters for the sample of full-time workers (which is defined as individuals who worked at least 1500 hours during the year). Imposing $\delta = 0.2$ (a 20% depreciation rate for human capital) is consistent with the estimates reported in Imai and Keane (2004) for prime age men.

---

[23] Z_{it} includes the following explanatory variables: age, two education dummies, marital status, number of children less than 18. The latter two variables provide exclusion restrictions. The education dummies correspond to high school completed and at least some years of college. The omitted variable refers to individuals who did not complete high school. The dependent variable is an indicator function that equals one if we observe the individual's wage in a given year and it is equal to zero otherwise.

[24] Reported nominal wages are deflated by the Consumer Price Index, base year = 1992.

[25] In particular, the estimates must be constrained to obtain a small value for the ratio $(\delta/\eta)^{1/\psi}$.

[26] Since the youngest individuals in the sample are 20 years old, I use $a_{it} = age_{it} - 19$ in the estimation.

Table 1
Estimates of the human capital production function

	1970s		1990s	
	women	men	women	men
η_0	0.0114	0.0149	0.0143	0.0158
	(0.0015)	(7.3e–04)	(0.0012)	(8.4e–04)
η_1	−1.19e–04	−2.67e–04	−1.89e–04	−2.55e–04
	(1.3e–06)	(6.0e–05)	(1.13e–05)	(7.2e–06)
η_2	−6.29e–07	4.3e–06	7.43e–07	3.51e–06
	(1.9e–07)	(1.3e–06)	(1.04e–08)	(2.3e–07)
constant	−0.22	0.70	−0.64	−1.58
	(0.04)	(0.06)	(0.051)	(0.155)
Observations	7432	11,037	10,554	11,982

Standard errors are in parentheses.

According to my 1970s estimates, one additional hour of work at age 20 increases the rate of growth of hourly wages by 0.0073 percentage points for men, and by 0.0056 points for women. These statistics are computed by substituting the estimated parameters in:

$$\frac{\partial \ln(\frac{w_{it+1}}{w_{it}})}{\partial n_{it}} = \frac{\psi \eta(a_{it}) n_{it}^{\psi-1}}{(1-\delta) + \eta(a_{it}) n_{it}^{\psi}},$$

which I refer to as the rate of return to labor market experience. Note that since the estimated η_1 and η_2 are fairly small, the key parameter in the estimation is η_0, which acts like TFP in the production of human capital investment. The resulting female/male ratio in returns to experience for the 1970s is 77%.

In the 1990s, one additional hour of work increases the rate of growth of hourly wages by 0.0077 points for men and by 0.007 points for women, which results in a female/male ratio of the rates of return to labor market experience of approximately 91%, a considerable increase compared to the 1970s estimate. Women are more similar to men in terms of their returns to labor market activities. These estimates imply an increase by 6% in men's returns between the two decades, and by 25% for women.

I will use the numbers reported in Table 1 for the calibration exercise. More specifically, I use the estimates for the 1970s to calibrate the benchmark economy. The estimates for the 1990s are used in the first experiment.

In the literature, estimates of the returns to labor market experience are obtained from the coefficients of a Mincer log-earnings equation that includes experience and its square.[27] My definition of returns to labor market experience refers to the impact of additional working hours on the rate of growth of earnings. Simple calculations show that my estimates of the human capital production function parameters imply returns to one extra year of full-time work experience ranging from 3 to 5%. These values are of the same order of magnitude as the estimates reported in Blau and Kahn (1997) based on a Mincer-log earnings specification.[28]

[27] Given the earnings equation $\ln w_{it} = \alpha_0 + \alpha_1 X_{it} + \alpha_2 X_{it}^2 + \alpha_3 Z_{it}' + \varepsilon_{it}$, where X_{it} is year of experience and Z_{it} is a vector of additional controls, the returns to labor market experience are given by: $\partial \ln w_{it} / \partial X_{it} = \alpha_1 + 2\alpha_2 X_{it}$.

[28] The following transformation can be used to compare the estimated coefficients. Based on the Mincer specification for the wage regression we can obtain: $\partial \ln(\frac{w_{it+1}}{w_{it}}) / \partial P_{it+1} = \alpha_1 + 2\alpha_2 (X_{it} + P_{it+1})$ where X_{it} represents accumulated labor market experience at time t and P_{it+1} is an indicator function that is equal to 1 if the individual worked full-

My estimates for men's returns to experience and for their growth over time are also consistent with those documented by Blau and Kahn (1997). According to their estimates, men's returns to full-time experience increased by 9% whereas women's returns more than doubled. Hence, as one would expect, the correction for non-random selection into the sample has a large impact on the estimates of the returns for female workers, whereas it has no effect on the estimated values for men.

The results described in this section are robust to the choice of the values of the fixed parameter δ and ψ as well as to the choice of an alternative specification for the human capital production function. In all these cases I find that women experienced a relatively larger increase in returns to labor market experience than men, with estimated returns ranging from 16 to 27%.[29]

5. Calibration

The basic calibration strategy is to parametrize the model using data from the 1970s, and then ask how the model's prediction change when we allow for changes in the returns to experience to the values estimated for the 1990s. The baseline economy, the 1970s, is calibrated to the PSID data in the following way. First, the parameters of the learning-by-doing human capital production function are set to their 1970s estimated values for men and women. The remaining parameters are chosen to match the 1970s work hours and hourly wage age profiles for married women and married men. Given the parameters obtained for the baseline economy, I change the parameters of the human capital production function to their 1990s estimated values and study the model's prediction in terms of age profiles for hourly wages and hours worked. I also consider a second experiment where, given the baseline economy parameters, the female/male gender differential is changed to its 1990s values. Finally, I compute and compare the contribution of the relative change in returns to experience, and of the decline in the male/female wage gap, to the observed 1970s–1990s change in the life-cycle pattern of work hours for married women and men.

The model is parametrized to the annual hours worked and hourly wage age profiles for married men and married women as in the PSID data for the 1970s. This is the subset of married individuals from the PSID samples used for the estimation of the human capital production function.[30] The actual 1970s and 1990s age profiles for average annual hours worked per person, and hourly wages are reported in Table A.2 in the appendix.

Initial condition The initial levels of human capital, θ_{m0} and θ_{f0}, are both set equal to one. Thus I assume that men and women in the 1970s do not differ in terms of their initial level of human capital. The female/male differential is captured by the market premium to female and male human capital, π_g. I also assume that households start their adult life with no assets.

time between time t and time $t+1$. Roughly speaking, this corresponds to my measure of the returns to labor market experience.

[29] The results of this sensitivity analysis are available at http://people.bu.edu/olivetti/workingpapers.htm. I have also performed the estimation without imposing any constraints on the parameters of the human capital production function. While these estimates cannot be fed into the model they imply changes in returns to experience, by gender and over time, that are consistent with the ones reported in this section.

[30] Note that the estimates of the learning-by-doing production function refer to the working population as a whole. In my interpretation, they represent the market returns that a married individual would obtain if he/she were to work on the labor market.

Rental rates In the simulation I set the real interest rate, r, to 5% annually. The compounded interest over the ten years period is therefore equal to 0.62. The gender-specific rental rates for human capital are set so that wages in the first period are equal to the average hourly wage for young men and women aged 20–29 as found in the PSID. That is, $\pi_f = 9.1$ and $\pi_m = 13.1$. The ratio $\pi_f/\pi_m = 0.696$ represents the "pure" gender wage differential. I will refer to this ratio as to the 1970s gender wage differential.[31]

Preferences The parameter α_g and the disutility of labor A are set to mimic the actual age profiles for average hours worked and hourly wages for both men and women in the 1970s. The discount rate β is set to be equal to $1/(1+R)$ where $R = (1+r)^{10} - 1$ is the compounded real interest rate. That is, I assume the rate of time preferences and the real interest rate to be equal over the (10 year) time period. This is consistent with the data when we control for changes in family size over the life cycle.[32]

Childcare production To my knowledge there is no evidence about a correct range of values for the parameter representing the elasticity of substitution between maternal time and market produced goods and services. In the following, I set the parameter representing the elasticity of substitution, ρ, and the relative weight on motherly time and market good and services, γ_t, in order to match the hours and wages age profiles for both men and women in the 1970s.

Law of motion for human capital The estimates are set to the 1970s values reported in Table 1. They correspond to a female/male differential in the rate of returns to experience of approximately 0.77. Since the estimates are obtained using yearly data, I rescale them to the 10-year period using the following procedure. Given that the number of hours worked is constant over the (10-year) period, I simply iterate the yearly law of motion over the ten year period for the (constant) number of hours worked.[33]

6. Results

In this section the results of the calibration are presented, and the 1970s predicted age profiles for hourly wages and work hours are compared to the 1970s data both for married men and for married women. In Section 6.1.1. I present the results obtained when the returns to experience are set to their 1990s values, keeping the remaining parameters of the baseline economy constant. The predicted age profiles for wages and hours worked for 1970s and 1990s are compared to the corresponding actual age profiles. In Section 6.1.2. I study the effect of a decline of the gender wage gap, holding returns to experience constant. Then, I quantify and discuss the contribution of these two competing explanations to the change in the pattern of work hours over the life

[31] I also run simulations under a different parametrization for the rental rates and the initial stock of human capital. In this case, the rental rates of human capital are normalized to one, whereas the θ_{g0}s are set to the average hourly wage for men and women aged 20–29 as found in the PSID (i.e. ω_{g1}). In this case θ_{f0}/θ_{m0} represents the "pure" gender wage differential. The results obtained under this parametrization are equivalent to the ones described in the next section.

[32] As in Attanasio (1994).

[33] More in detail, the human capital stock at the end of period t is given by: $\theta_{gt} = [\prod_{j=t}^{t+9}((1-\delta_g)+\eta_g(a_{gj})n_{gt}^{\psi_g})]\theta_{gt-1}$, $t = 1, \ldots, 4$, where n_{gt} represents the average hours worked during the (10-year) period t. Also, since the estimates are obtained for annual hours of work the model prediction on n_{gt} is rescaled by multiplying by 5000 (assuming that the total time endowment of an individual is 5000 hours a year).

cycle for married women. The model also gives a prediction about the percentage increase in the female/male differential observed in the data that can be imputed to the relative change in returns to experience. Finally, in the last two sections, I briefly discuss the results obtained when the economy is calibrated to match the data for single female- and male-headed households and perform sensitivity analysis.

6.1. The baseline economy

Table 2 represents the parameters that characterize the baseline economy. The values chosen for α_m and α_f would imply an intertemporal elasticity of substitution around 0.49 for men, and 0.53 for women, although these parameters do not have a straightforward interpretation as intertemporal elasticity of substitution in this model since wage age profiles are endogenous. This is consistent with the finding, typical in the micro literature, that the female elasticity of substitution is generally bigger for women than for men.[34] The values chosen for π_m and π_f give raise to a female/male gender wage differential equal to 0.696. As mentioned above, I refer to this ratio as to the 1970s "pure" gender gap.

Figure 3 compares the model prediction and the actual data for the 1970s age profiles. Panel A displays the work hours age profiles for married men and married women in the data and in the model. The crossed solid line and thick solid line represent the actual 1970s annual hours worked, respectively for married women and married men, the dotted and dashed lines are generated by the model. The predicted profiles closely approximate the actual profiles. For what concerns hourly wages (see panel B), the model does quite well, although men's predicted hourly wages are slightly higher than the actual ones for the 40–49 and 50–59 age groups. Also, women's predicted age profiles are decreasing, whereas in the data they slightly increase for the age group 50–59. This is due to the fact that the system is over-identified. In fact, given the assumptions made for the real interest rate, the discount factor, the rental rates for male and female human capital, the initial condition for asset and human capital stock, and the parameters of the human capital production function, we are left with seven free parameters and 14 data points. Thus the predicted age profiles are slightly different from the actual ones.

The model also predicts concave profiles for total household consumption. The profiles are flat when we restrict our attention to consumption goods and services. This is not surprising since the

Table 2
Model parameters

Preferences		Rental rates	
α_m	3.02	r	5%
α_f	2.88	π_f	9.08
A	25	π_m	13.05
b	0.35	Initial conditions	
Childcare production		a_0	0
ρ	0.75	θ_{f0}	1
γ	0.56, 0.23	θ_{m0}	1

[34] In the micro literature the estimates of the elasticity of substitution for men range between 0.1 and 0.45 (see, for example, Altonji, 1986 and Pencavel, 1986). For women there is a wide array of estimates for the intertemporal elasticity of substitution ranging from negative values to large positive values (see Mroz, 1987, and Killingsworth and Heckman, 1986).

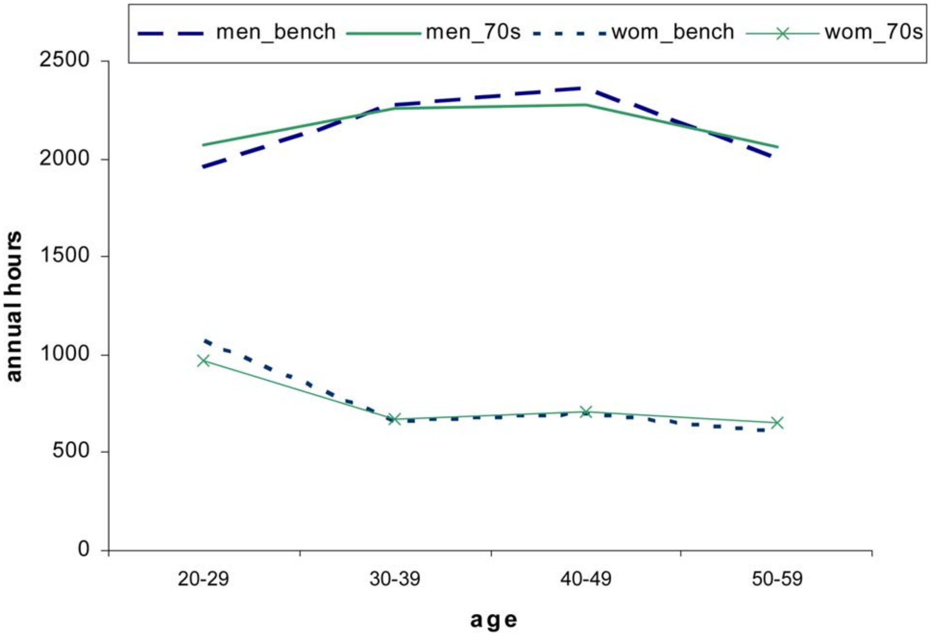

Panel A: Annual hours worked

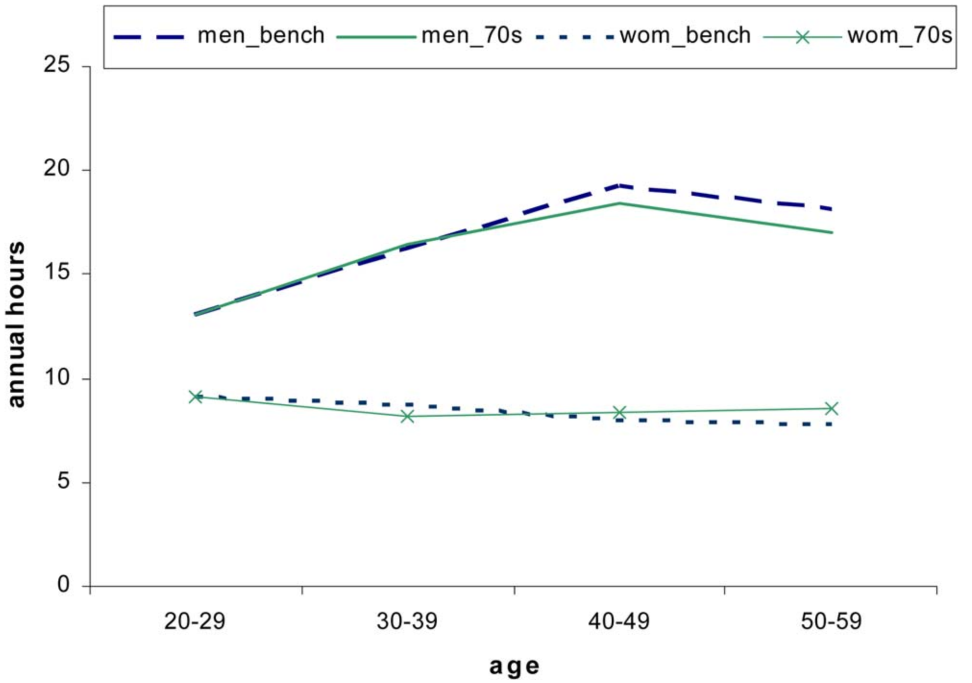

Panel B: Hourly wages

Fig. 3. The baseline economy and the data.

rate of time preferences is assumed to be equal to real interest rate. Moreover, the model predicts a childcare consumption share of total household income around 6% in the first (10 year) period of children's life and of 7.5 percent in the second period. On average the model predicts a childcare expenditure share of 6.77%. Women spend 32% of period 2 working hours taking care of their children. In the second period the fraction of maternal time devoted to childcare decreases to around 3.7% of total hours worked.

6.1.1. Experiment 1: Change in returns to experience

In the following experiment the parameters of the human capital production function are changed to their 1990s estimated values for both men and women while keeping the remaining parameters constant to the baseline economy. I compare the predicted change in age profiles for married men and women to the actual change observed in the data. The 1990s actual profiles for work hours and hourly wages are reported in Table A.2 in the appendix.

In this experiment the estimates of the human capital production function are set to the 1990s values reported in Table 2. They correspond to a female/male differential in the rates of return to experience equal to 0.91. That is, men's returns at age 20 increase by 6% with respect to the 1970s estimates presented in Table 2, whereas women's returns increase by 25%. This implies an 18% increase in the female/male returns to labor market experience (from 0.77 to 0.91). Note that since the estimates for η_1 and η_2 are very small this experiment basically amounts to changing a single parameter, η_0, for both men and women.

The results are described in Fig. 4. Panel A displays the age profiles for annual hours worked for the benchmark economy, the "1970s", as well as the predicted profiles under this experiment, the "1990s." In all the graphs the dashed line represents "1990s" men, the solid line "1970s" men, the dotted line "1990s" women, and the crossed solid line "1970s" women. As panel A shows, given the parameters of the baseline economy, the increase in rates of return to experience alone generates the change from the single-peaked to the double-peaked work hour age profile for married women that this paper aims to explain. At the same time, consistent with the data, the model predicts only modest changes in the life-cycle profile of hours worked for men. This is mainly due to the fact that men are already working forty percent of their time in the baseline economy.

For what concerns wages (see panel B), the change in profiles generated by the model goes in the same direction as the actual change, although the 1990s profiles do not perfectly match what is observed in the data. Both men's and women's profiles are steeper in the 1990s than in the 1970s. Moreover, married women's age profile for hourly wages increases with age as compared to the decreasing profile of the baseline economy. These results show that changes in returns to experience alone can generate substantial changes in women's labor force participation and wages over the life cycle.

Although the 1990s model has the same prediction as the baseline model for consumption profiles, household consume more now since they have more income. Moreover, the model predicts a childcare consumption share of total household income around 7.3% in the first (10-year) period of children's life, and of 7.4% in the second period. The model generates a total childcare expenditure share of 7.38% in the 1990s. Overall the total childcare expenditure share is less then one percentage point (of total household income) higher in the 1990s than in the 1970s. The childcare expenditure share is higher in the first period of children's life in the 1990s as compared to the 1970s, whereas it is lower in the second period. The fraction of time women spend in childcare drops dramatically. Women spend only 6% of their period 2 working hours

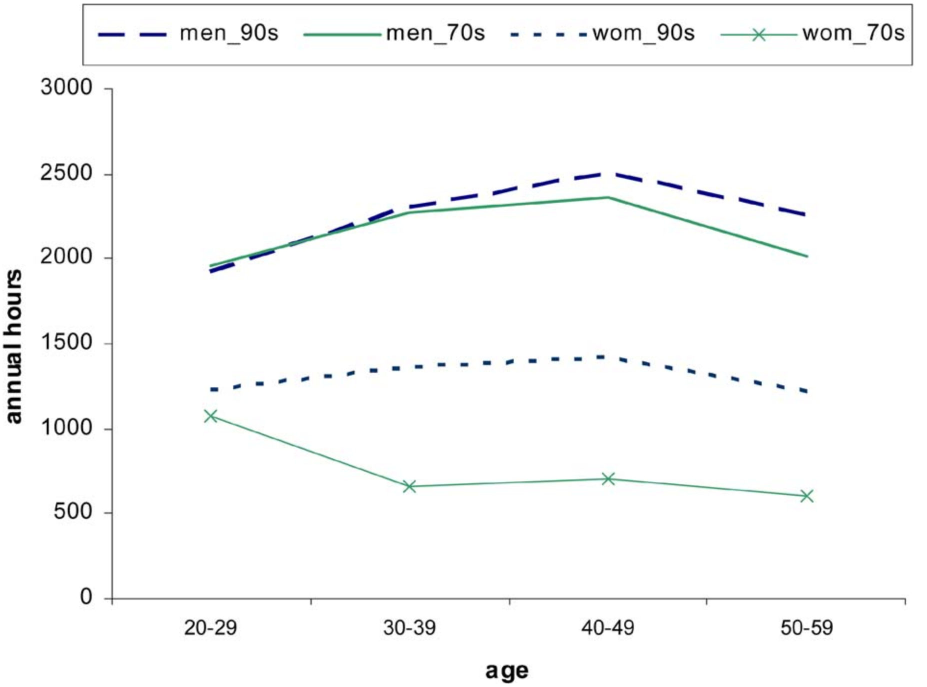

Panel A: Annual hours worked

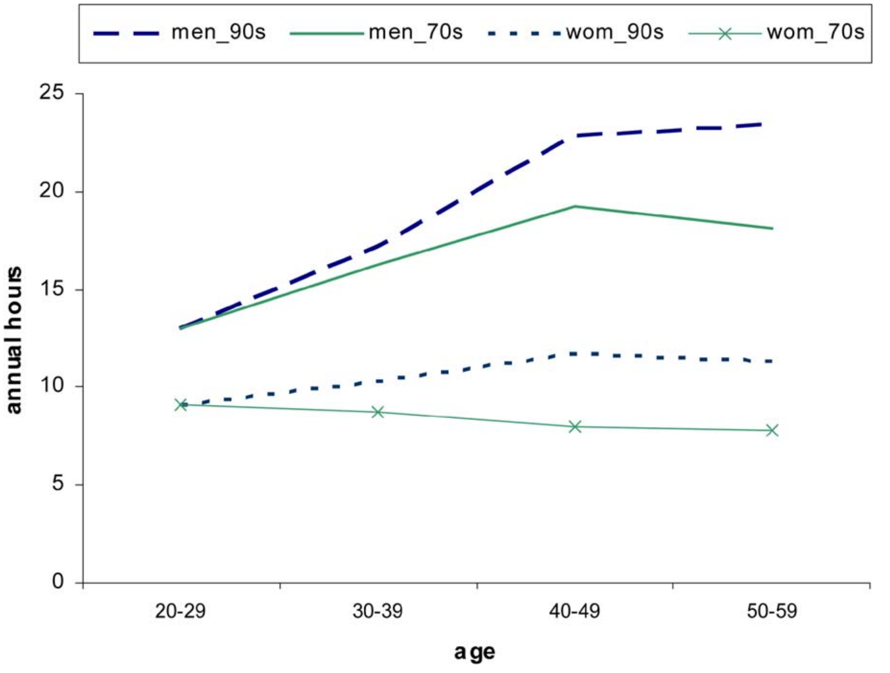

Panel B: Hourly wages

Fig. 4. Experiment 1—change in returns to experience.

providing childcare (from a 32% in the 1970s). In the second period of children's life the fraction of maternal time spent in childcare activities drops to 0.003% of total hours worked.

6.1.2. Experiment 2: Change in the "pure" gender wage gap

In the second experiment, I change the rental rates of female and male human capital to their 1990s values, holding the returns to experience and the remaining parameters constant. In particular, I set $\pi_f = 10.5$ and $\pi_m = 13.3$, thus the "pure" gender wage differential, that is the ratio π_f/π_m, is now equal to 0.79. This is equivalent to an increase in the female/male ratio by 13%.

Figure 5 presents the results for this experiment. As described in the previous section, panel A displays the predicted change of work hours age profiles and panel B presents the predicted change for hourly wages. The change in the pure wage differential generates an increase in average hours worked per person for every age group but it cannot explain the change in the shape of the age profile. Moreover, although women's hourly wages increase, the wage age profile is still declining in age. In this case, married men's behavior does not change much between the 1970s and the 1990s. For this group, both the work hours and the earnings age profile are almost unaltered.

Household consumption and savings profiles are almost identical to the 1970s profiles. The childcare expenditure share slightly increases. The model predicts a childcare consumption share of total household income around 6.4% in the first (10-year) period of children's life, and of 7.5% in the second period. The model generates a total childcare expenditure share of 6.96% in the 1990s as compared to 6.77% in the 1970s. The fraction of time women spend in childcare drops. Women spend only 19% of their period 2 working hours providing childcare (as compared to 32% in the 1970s). In the second period of children's life the fraction of maternal time spent in childcare activities drops to 0.7% of total hours worked. Although in the simulation the gender-specific efficiency wage are simply set in order to reproduce the "pure" female/male wage differential for young workers in the 1990s, the model generates an average female/male wage ratio over all age groups that closely matches the one observed in the data.

To conclude, I also consider the joint effect of an increase in both the gender wage gap and the returns to experience to their 1990s value. In this case, the model generates larger average work hours than those observed in the data for every age group. This suggests that the two explanations explored in this paper interacts to generate a further increase in average hours worked by women over the life cycle. The intuitive explanation is that in this case not only women's learning-by-doing profile is steeper than in the 1970s, but also they are initially better off in terms of their human capital. Men slightly decrease their average hours worked.

6.2. Discussion

The change in relative returns to experience has a strong impact on married women's hours worked age profiles. On the contrary, at least under this specification, the decline in the gender wage gap seems to play a much smaller role in explaining the increase in women's hours of work. In order to assess quantitatively the contribution of these two competing explanations, I compare the percentage change in average hours worked per person observed in the data to the one generated by the model. Table 3 summarizes the results.

For 20 to 59 year old women actual annual hours worked per person increased by 75.2% between the 1970s and the 1990s. The model generates a 72.4% increase in hours worked in the first experiment and a 13.5% increase in the second experiment. Hence, the change in returns to

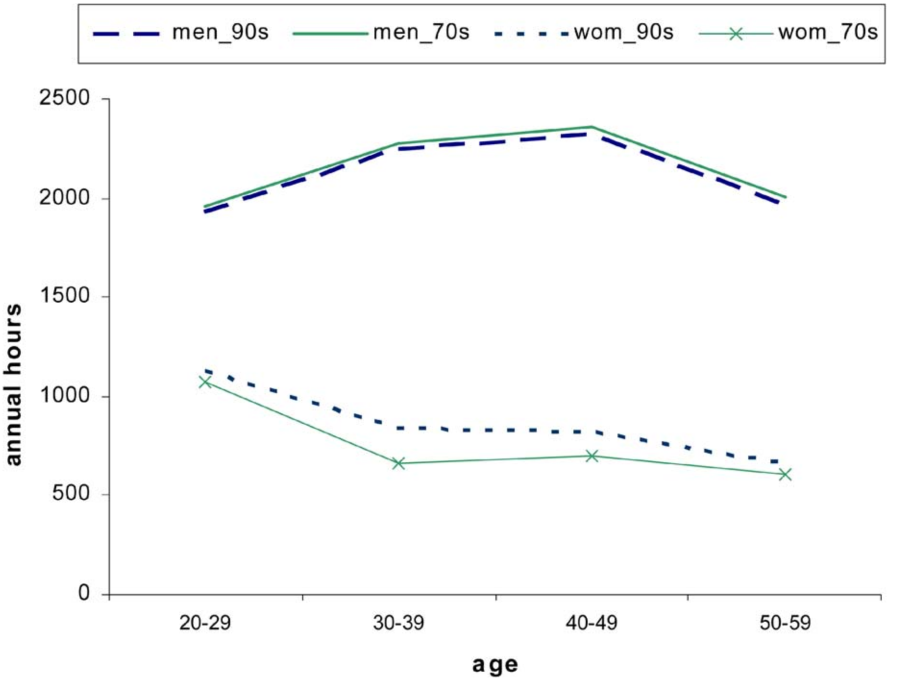

Panel A: Annual hours worked

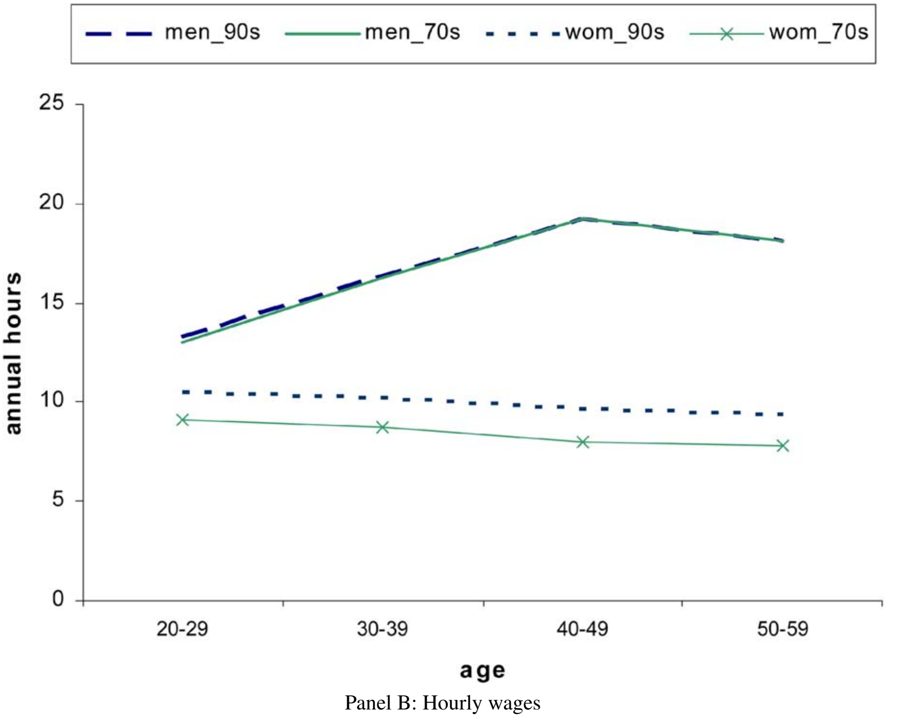

Panel B: Hourly wages

Fig. 5. Experiment 2—change in gender wage gap.

Table 3
Change in hours worked per married women: Data vs. Model (percentage values)

	20–29	30–39	40–49	50–59	20–59
Data	35	117	102	68	75.2
Experiment 1: returns	14	112	102.6	97	72.4
Experiment 2: gender gap	3.8	20.6	12.2	12	13.5

labor market experience can account for 96% of the increase in lifetime average hours worked per married woman. The change in the gender wage gap can account for only 18% of the increase. Moreover, in the first experiment the model can account for 40% of the increase in average hours worked per person observed in the data for the age group 20–29, and for approximately 95% of the increase observed for the age group 30–39. For the age groups 40–49 and 50–59 the generated increase in average hours worked over-predicts the one observed in the data. On the other hand, the change in the gender wage gap accounts for only a small increase in hours worked in each age group and for 18% of the increase in average hours worked by 30–39 year old married women—the group whose change in work behavior determines the change in the shape of work hours age profiles.

The increase in relative returns to experience can generate a large part of the increase in hours worked for women in childbearing age (although it over-predicts the increase for older women). Most importantly, it can explain the change in the shape of life-cycle profiles of hours worked for married women between 1970 and 1990—from double-peaked to single-peaked. On the contrary, the increase in the gender wage gap can explain only a modest fraction of the increase in hours worked for every age group, and it cannot generate the changing shape of married women's age profiles. This evidence suggests that although there are surely other factors that contributed to the change in labor force participation of married women with young children, the relative change in returns to experience seems to be a very important determinant of this change.

The model also allows us to measure the contribution of the relative change in returns to experience to the decline in the gender wage gap observed in the data. According to the model, the increase in returns to experience alone accounts for about 42% of the increase in the female/male wage ratio found in the data. Since wage age profiles are endogenous in this model, it is worthwhile to compare the female/male wage differential by age group predicted by the model for the 1990s (when only returns to experience are changed) with the values observed in the data. The results are summarized in Table 4.

Entries in the table represent the actual and predicted female/male wage ratios. Average wages by gender are weighted by the fraction of hours worked by each group with respect to the total hours worked in the population. The model matches quite well the observed age-specific female/male wage differential.

The model also makes predictions about children's welfare. Under experiment I, children's equilibrium lifetime well-being increases by 28% with respect to the baseline economy despite an eighty percent decline of the share of motherly time spent in childcare. The increase in family income allows a family to substitute motherly time with market-produced goods and services in the production of childcare. Under the second experiment, children's welfare slightly increases (by 0.3%) even if motherly time spent in childcare drops by forty percent. Thus, children's welfare does not decline as a consequence of the increase in hours worked on the market by mothers. Moreover, despite the extent to which a family substitutes parental time with market good and services, the childcare expenditure share increases only slightly in both experiments (up to 1%

Table 4
Gender wage differentials by age, actual and predicted profiles

	Data	Model
20–29	0.786	0.764
30–39	0.684	0.611
40–49	0.586	0.501
50–59	0.457	0.447

under experiment I).[35] The small increase is consistent with what is found in the PSID data when we compute the childcare expenditure share of family labor income for married household aged 30–39 with young children. The model also generates values of the childcare expenditure share that are consistent with data from the 1970s and the 1990s.[36]

6.3. Single household

I also study whether the observed change in returns to labor market experience can predict changes in the work hours age profiles for single men and women that are consistent with the data. To this aim, given the 1970s estimates of the human capital production function, I parametrize the model to match the age profiles for hours of work and hourly wages for single females and single males. As previously discussed, in 1970 single women displayed a work hours age profile similar to that of men. Moreover, lifetime average hours worked by single men and single women only slightly increased between 1970 and 1990. The experiments show that the model is capable of reproducing this feature of the data. In particular, this is true for a range of values of the intertemporal elasticity of substitution parameter (ranging from 0.84 to 5.6). The values of the intertemporal elasticity of substitution needed to match the 1970s data for single are slightly higher than those needed to match the behavior of married couples.

The results for one of these experiments are presented in Fig. 6. Panel A presents the actual change of age profiles of work hours for single men and women. Panel B presents the change predicted by the model. As in the previous experiments, the 1970s economy is parametrized in order to match the 1970s profiles for single men and single women. The model can match the shape in the work hours age profiles for both genders and the relative change in their 1990s profiles, although the predicted increase in average hours worked per person is larger than what is observed in the data.

[35] Attanasio et al. (2004) show that shifts in the cost of children relative to lifetime earnings are an important explanation for the change in labor force participation across cohorts. In the context of my model, even if child care costs do not change over time, we would observe a decline in the cost of children relative to lifetime earnings in the presence of increasing returns to labor market activities that increase hours worked by women. In addition to its direct impact on women's labor participation, a decline in the cost of childcare would reinforce the impact of the change in labor returns on women's labor supply when they have children.

[36] Anderson and Levine (1999), using 1990–1993 SIPP data, find that families with at least one child under 13 spend, on average, 7% of their family income in childcare, whereas families with at least one child under 6 spend on average 7.7% of their income. I find similar values by computing the same shares using the PSID data for 1973 and 1993.

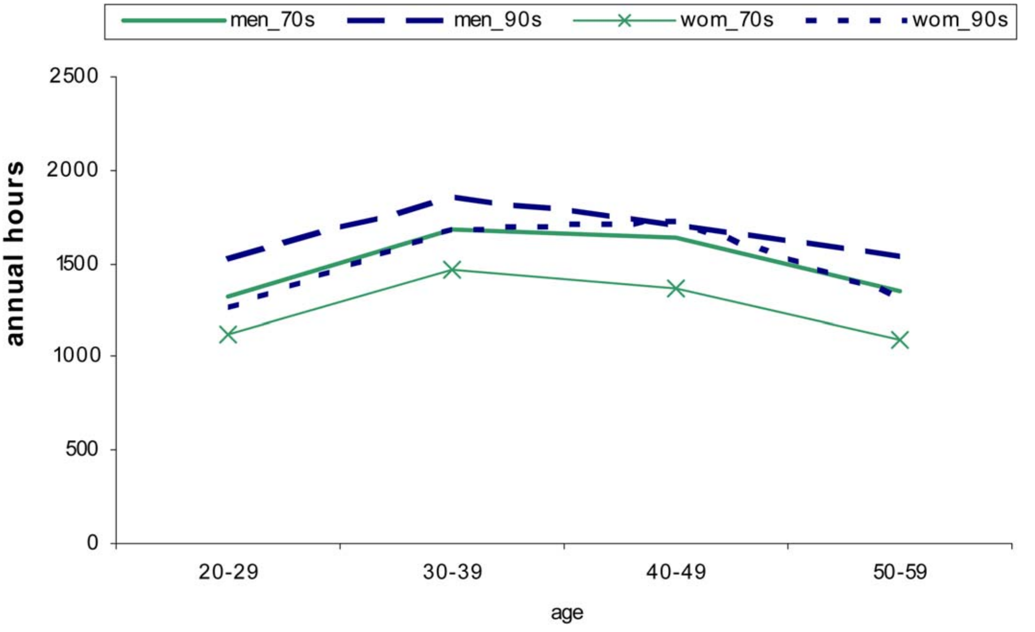

Panel A: Annual hours worked—data.

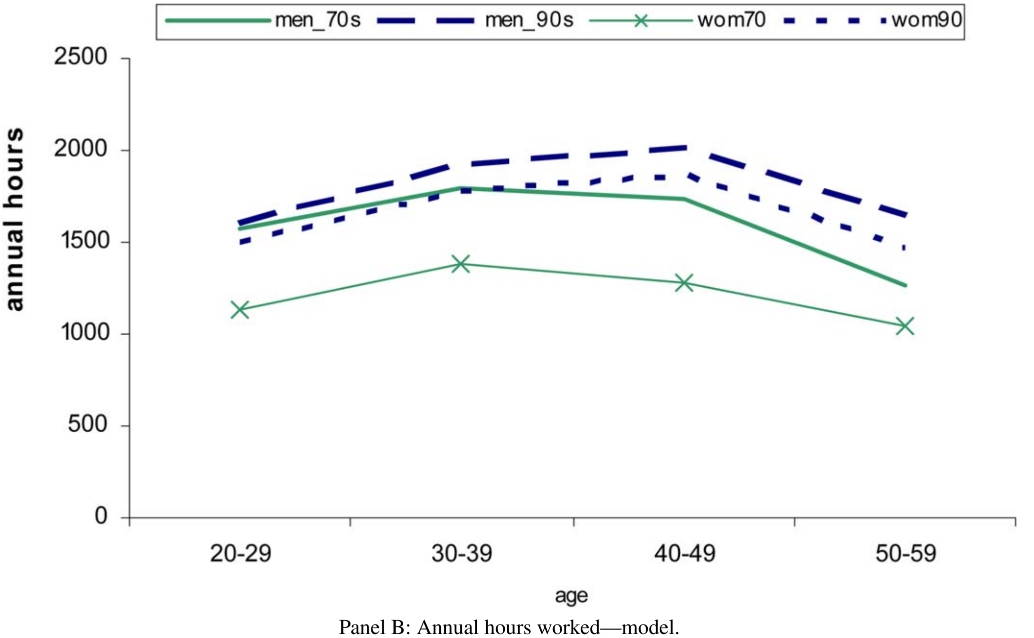

Panel B: Annual hours worked—model.

Fig. 6. Change in returns to experience: Single households.

6.4. Sensitivity analysis

In this section I study the sensitivity of the results to changes in some of the parameters used in the calibration exercise.[37] I perform the analysis by recalibrating the economy to the 1970s and then by changing the returns to labor market experience to their 1990s values. This allow to study whether the result (i.e. the strong contribution of the change in returns to labor market experience) is robust to the parameters chosen in the calibration exercise. In particular, I will study the robustness to changes in the real interest rate, the female/male wage ratio and the intertemporal elasticity of substitution for both men and women.

In the literature[38], the female/male wage ratio for the 1970s ranges from 0.6 over all age groups to 0.764 when focusing on workers aged 20 to 24 and controlling for education, race and full-time labor market participation. Therefore, I recalibrate the economy setting $\pi_f/\pi_m = 0.6$ in the first case, and $\pi_f/\pi_m = 0.764$ in the second case. Under the first parametrization the model can account for about 66% of the observed increase in average hours worked per person observed in the data for the age group 30–39. In the latter case the model accounts for 97.6% of the increase. Moreover, given the baseline economy discussed in Section 6.1, I show that in order for the increase in the female/male wage ratio to generate the change in the *shape* of the work hours age profiles observed in the data one would need a 1.9 female/male wage ratio. That is, young women should be earning almost twice as much as men. In fact, when the rental rates to human capital for men and women are the same we still observe double-peaked profiles for women's hours of market work. Hence, the increase in the female/male gender wage gap seems to play a secondary role in explaining the change in the working behavior of women over their life cycle.

I also study the sensitivity of the results with respect to the real interest rate by setting the annual real interest rate r equal to 4 and 6% respectively. The model is not very sensitive to changes in the real interest rate (indeed the values of the preferences and childcare production parameters needed to match the 1970s economy change only slightly). In both cases, I obtain results roughly equivalent to the ones obtained when I set the yearly real interest rate, r, to 5%.

For what concerns the intertemporal elasticity of substitution, in the micro literature the estimates of the elasticity for men range between 0.1 and 0.45 (see, for example, Altonji, 1986 and Pencavel, 1986). In the macro literature higher values of the elasticity are typically used to calibrate the model (for example, Prescott, 1986 uses 2 in his calibration exercise). More recently Imai and Keane (2004) obtain an elasticity estimate equal to 3.82 for prime-age men in the context of a dynamic labor supply model with learning-by-doing human capital accumulation. For women there is a wide array of estimates for the intertemporal elasticity of substitution ranging from negative values to large and positive values (see Mroz, 1987, and Killingsworth and Heckman, 1986). In general, there has been consensus that female labor supply wage elasticities are larger in absolute value than are men's. As a consequence, I perform two experiments. In the first one I set both men's and women's elasticities to "low" values (0.3 and 0.35 respectively). In the second experiment I choose "high" elasticities (1.2 and 1.3 respectively). In both cases, the model generates a change in hours of market work over the life cycle of the same order of magnitude as the ones described previously. That is, the change in returns to labor market experience

[37] I also consider a general equilibrium version of the model to study how the implications of the model change when prices are endogenous (available at http://people.bu.edu/olivetti/workingpapers.htm). I find a small general equilibrium effect.

[38] See Smith and Ward (1985), Goldin (1990).

generates an increase in life time average hours of market work around 90% whereas an increase in the female/male wage gap can only explain approximately 20% of the total increase.

7. Conclusion

This paper investigates the change in patterns of work hours over the life cycle for married women and married mothers of children of preschool age. In particular, it focuses on the relative increase in women's returns to experience as one possible explanation for the change. A dynamic model is built to quantitatively assess the contribution of changes in the rate of return to experience, and of the decline of the gender wage gap, to the change in married women's life-cycle profiles of hours worked. The results show that the relative change in returns to experience can account for a large fraction of the observed change in average hours worked for married women and for the change in the shape of their work hours and wage age profile. The model also allows me to quantify the contribution of the increase in returns to experience to the decline of the gender wage gap that is found in the data. The results also show that the decline of the gender wage gap cannot explain the change in the shape of women's life-cycle profiles. Consistent with the data, the model predicts only modest changes in the life-cycle profiles of hours worked for men and single women.

Although in this paper I take the change in the wage structure as exogenous, the analysis indicates it is important to go one step further and study the determinants of the relatively larger increase in women's returns to experience. In general, the increase in returns to experience can be attributed to technological change that favors more skilled workers. Technological progress favorable to women's characteristics could have contributed to the relative increase in women's returns to experience. Other possible explanations include the change in the distribution of female workers by occupation and the decline in discrimination against women, particularly against married women and married mothers. Regarding the first explanation, women earned access to jobs for which labor experience is more important, and this might have caused an increase in returns to experience. As for the second, a reduction in discrimination could have occurred as a direct result of the activity of government agencies (e.g. the Equal Employment Opportunity Commission). The investigation of the causes of the relative change in returns is left for future work.

Acknowledgments

I want to thank Richard Rogerson for insightful comments. This paper also benefitted from conversations with Marianne Baxter, Kevin Lang, and Victor Rios-Rull.

Appendix A. Summary statistics

In Table A.1, I report the summary statistics for the PSID samples used to estimate the human capital production function. Table A.2 reports the basic statistics used in the calibration exercise. I consider four samples: men/women in the 1970s and men/women in the 1990s.

The 1970s samples pools together individuals from the 1970 to 1977 waves of the PSID. The 1990s samples include individuals from the 1990 to 1997 waves. Each wave contains approximately 2000 men and women in the 20–60 age range. In the 1970s sample, the fraction of women who do not work is around 45%. This percentage drops to 25% in the 1990s sample. This fact clearly shows how the problem of sample selection is stronger for women in the 1970s. Summary

Table A.1
Summary statistics

	1970s				1990s			
	working		non-working		working		non-working	
	women	men	women	men	women	men	women	men
age	40.16	39.86	41.13	48.65	39.25	38.78	41.6	43.62
	(10.46)	(10.3)	(10.63)	(10.03)	(8.58)	(8.60)	(10.15)	(9.76)
married	0.69	0.91	0.80	0.77	0.77	0.83	0.74	0.73
	(0.46)	(0.29)	(0.40)	(0.42)	(0.42)	(0.37)	(0.44)	(0.44)
single	0.074	0.03	0.031	0.05	0.11	0.085	0.128	0.13
	(0.26)	(0.17)	(0.17)	(0.21)	(0.31)	(0.28)	(0.33)	(0.34)
widowed	0.07	0.01	0.047	0.06	0.016	0.004	0.039	0.018
	(0.26)	(0.10)	(0.21)	(0.24)	(0.126)	(0.067)	(0.19)	(0.13)
divorced	0.10	0.03	0.051	0.05	0.079	0.061	0.048	0.095
	(0.30)	(0.17)	(0.22)	(0.23)	(0.27)	(0.24)	(0.21)	(0.29)
separated	0.067	0.02	0.071	0.06	0.025	0.020	0.045	0.027
	(0.25)	(0.15)	(0.26)	(0.24)	(0.15)	(0.14)	(0.21)	(0.16)
no. children < 18	1.61	1.91	2.19	1.60	1.21	1.24	1.49	1.11
	(1.72)	(1.8)	(0.41)	(2.06)	(1.16)	(1.20)	(1.45)	(1.37)
less than HS	0.63	0.40	0.76	0.71	0.10	0.12	0.29	0.29
	(0.48)	(0.49)	(0.43)	(0.45)	(0.30)	(0.33)	(0.45)	(0.46)
HS completed	0.299	0.45	0.22	0.26	0.65	0.58	0.57	0.50
	(0.46)	(0.49)	(0.42)	(0.44)	(0.47)	(0.49)	(0.49)	(0.50)
more than HS	0.065	0.15	0.02	0.02	0.25	0.30	0.14	0.21
	(0.25)	(0.35)	(0.14)	(0.16)	(0.43)	(0.46)	(0.35)	(0.40)
annual hours	1480	2212			1728	2228		
	(677)	(657)			(650)	(625)		
ln(hourly wage)	2.10	2.53			2.39	2.71		
	(0.53)	(0.56)			(0.62)	(0.63)		
N of observations	7462	11,037	7095	638	10,554	11,982	3739	1635

Standard errors are in parentheses.

Table A.2
Average hours worked and wages by age, married samples

	Annual hours				Hourly wage			
	1970s		1990s		1970s		1990s	
age	men	women	men	women	men	women	men	women
20–29	2072.0	969.0	2202.7	1265.7	13.1	9.1	13.3	10.4
30–39	2260.2	674.0	2299.3	1466.0	16.4	8.2	19.4	13.2
40–49	2273.8	708.0	2286.9	1430.4	18.4	8.3	22.9	13.4
50–59	2057.5	654.0	2048.2	1102.3	17.0	8.6	23.7	10.8

statistics in Table A.1 are reported by gender/decade and for working and non-working individuals. The first-stage estimation of the sample selection correction terms includes both working and non-working individuals, the second-stage only consider individuals who report positive hours/wages.

Table A.2 reports the hourly wage and annual hours worked age profiles, for the 1970s and for the 1990s. These statistics are used in the calibration exercise. They are computed for a subset of married individuals from the main PSID samples described in Table A.1. The averages are computed for the pooled 1970s and 1990s samples (each consisting of 7 waves of the PSID). Hourly wages are computed as weighted averages with weights $\lambda_i = (f_i n_i)/\sum_i f_i n_i$ where f_i represents the number of people who worked n_i annual hours on the labor market.

References

Altug, S., Miller, R.A., 1998. The effect of work experience in female wages and labor supply. Review of Economic Studies 65, 45–85.

Altonji, J.G., 1986. Intertemporal substitution in labor supply: Evidence from the micro data. Journal of Political Economy 94, 176–215.

Anderson, P.M., Levine, P.B., 1999. Child care and mothers' employment decisions. Working paper No. 7058. NBER.

Attanasio, O., Law, H., Sanchez-Marcos, V., 2004. Explaining changes in female labour supply in a life-cycle model. Working paper 04/51. Cambridge University.

Attanasio, O.P., 1994. The intertemporal allocation of consumption: Theory and evidence. Working paper No. 4811. NBER.

Blau, F., 1998. Trends in the well-being of American women, 1970–1995. Journal of Economic Literature 36, 112–165. March.

Blau, F., Kahn, L., 1997. Swimming upstream: Trends in the gender wage differential in the 1980. Journal of Labor Economics 15, 1–42.

Blau, F., Ferber, M., Winkler, A., 2002. The Economics of Women, Men, and Work, fourth ed. Upper Sadle River, Prentice Hall, NJ.

Blundell, R., MaCurdy, T.T., 1999. Labor supply: A review of alternative approaches. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3A. Elsevier Science B.V.

Caucutt, E.M., Guner, N., Knowles, J., 2002. Why do women wait? Matching, wage inequality, and the incentives for fertility delay. Review of Economic Dynamics 5, 815–855.

Eckstein, Z., Wolpin, K., 1989. Dynamic labor force participation of married women with endogenous work experience. Review of Economic Studies 56, 375–390.

Fernández, R., Fogli, A., Olivetti, C., 2004. Mothers and sons: Preference development and female labor force dynamics. Quarterly Journal of Economics 119 (4), 1249–1299. November.

Francesconi, M., 2002. A joint dynamic model of fertility and work of married women. Journal of Labor Economics 20 (2), 336–380.

Galor, O., Weil, D.N., 1996. The gender gap, fertility and growth. American Economic Review 86, 374–387.

Greenwood, J., Guner, N.N., Knowles, J.A., 2003. More on marriage, fertility and the distribution of income. International Economic Review 44 (3), 827–862.

Greenwood, J., Seshadri, A., Yorukoglu, M., 2005. Engines of liberation. Review of Economic Studies 72, 109–133.

Goldin, C., 1989. Life-cycle labor-force participation of married women: Historical evidence and implications. Journal of Labor Economics 7 (1), 20–47.

Goldin, C., 1990. Understanding the Gender Gap: An Economic History of American Women. Oxford Univ. Press, New York.

Goldin, C., Katz, L.F., 2002. The power of the pill: Oral contraceptives and women's career and marriage decisions. Journal of Political Economy 100, 730–770.

Heckman, J., 1974. Shadow prices, gender differences and labor supply. Econometrica 42, 679–694.

Heckman, J., 1979. Sample selection bias as a specification error. Econometrica 47, 153–163.

Hotz, V.J., Miller, R.A., 1988. An empirical analysis of life-cycle fertility and female labor supply. Econometrica 56, 91–118.

Hotz, V.J., Klerman, J.A., Willis, R.J., 1997. The economics of fertility in developed countries. Chapter 7 in: Rosenzweig, M.R., Stark, O. (Eds.), Handbook of Population and Family Economics. Elsevier.

Imai, S., Keane, M.P., 2004. Intertemporal labor supply and human capital accumulation. International Economic Review 45 (2), 601–641.

Jones, L., McGrattan, E. Manuelli, R.E., 2003. Why are married women working so much? Research Department Staff Report 317. Federal Reserve Bank of Minneapolis.

Juhn, C., Murphy, K.M., 1997. Wage inequality and family labor supply. Journal of Labor Economics 15, 73–97.

Katz, L., Murphy, K., 1992. Changes in relative wages, 1963–1987: Supply and demand factors. Quarterly Journal of Economics, 35–78.

Killingsworth, M.R., Heckman, J., 1986. Female labor supply: A survey. Chapter 2 in: Ashenfelter, O., Layard, R. (Eds.), Handbook of Labor Economics. North-Holland.

Mc Grattan, E.R., Rogerson, R., 1998. Changes in hours worked since 1950. Federal Reserve Bank of Minneapolis Quarterly Review 22, 2–19.

Mincer, J., 1974. Schooling, Experience, and Earnings. Columbia Univ. Press.

Mincer, J., Polachek, S., 1974. Family investments in human capital: Earnings of women. Journal of Political Economy 82, 76–108.

Moffitt, R., 1984a. Profiles of fertility, labor supply and wages of married women: A complete life-cycle model. Review of Economic Studies 51, 263–278.

Moffitt, R., 1984b. Optimal life-cycle profiles of fertility and labor supply. Research in Population Economics 5, 29–50.

Mroz, T.A., 1987. The sensitivity of an empirical model of married women's hours of work to economic and statistical assumptions. Econometrica 55, 765–799.

O'Neill, J., Polachek, S., 1993. Why the gender gap in wages narrowed in the 1980s. Journal of Labor Economics 11, 205–228.

Pencavel, J.H., 1986. Labor supply of men: A survey. Chapter 1 in: Ashenfelter, O., Layard, R. (Eds.), Handbook of Labor Economics. North-Holland.

Pencavel, J.H., 1998. The market work behavior and wages of women, 1975–94. Journal of Human Resources 38 (4), 771–804.

Prescott, E.C., 1986. Theory ahead of business cycle measurement. Federal Reserve Bank of Minneapolis Quarterly Review 10 (4), 9–22. Fall.

Regalia, F., Rios-Rull, V., 2001. What accounts for the increase in single households and the stability in fertility? Manuscript. University of Pennsylvania.

Shaw, K., 1989. Life-cycle labor supply with human capital accumulation. International Economic Review 30, 431–456.

Smith, J.P., Ward, M.P., 1985. Time-series growth in the female labor force. Journal of Labor Economics 3, 59–90.

Vella, F., 1998. Estimating models with sample selection bias: A survey. Journal of Human Resources 33, 127–169.

Weiss, Y., 1972. On the optimal life-cycle pattern of labor supply. Economic Journal 83, 1293–1315.

Weiss, Y., Gronau, R., 1981. Expected interruptions in labor force participation and sex-related differences in earnings growth. Review of Economic Studies 48, 607–619.