

Explaining Changes in Female Labour Supply in a Life-cycle Model*

Orazio Attanasio,[†] Hamish Low[‡] and Virginia Sanchez-Marcos[§]

July 17, 2007

Abstract

In this paper we study the life cycle labour force participation of three cohorts of American women: those born in the 1930s, 1940s and 1950s. We first document the large shifts in labour supply behaviour and the patterns of observed wages among these three cohorts. Our focus is on the increase in the labour supply of mothers between the 1940s and 1950s cohorts. We construct a life cycle model with endogenous female labour force participation, consumption and saving choices to search for an explanation. The dynamics of labour supply depends on child costs (relative to earnings), returns to experience and the level of female wages. We calibrate the model to match the behaviour of the middle cohort and investigate which changes in the main determinants of labour supply could have accounted for the substantial increase in labour supply in the early part of the life cycle observed for the youngest cohort. We conclude that a combination of a reduction in the cost of children alongside a reduction in the wage gender gap is needed to match the observed changes.

*We would like to thank participants at the NBER Summer Institute Workshop on the Aggregate Implications of Microeconomic Consumption Behavior, and seminar audiences at the Macro Seminar at the University of Pennsylvania, the Stockholm School of Economics, MadMac at CEMFI, the University of California at San Diego, the Federal Reserve Bank of Chicago, the University of Wisconsin, and the conference on the Macroeconomics of Imperfect Risk Sharing at the University of California at Santa Barbara (May, 2006). We are very grateful for particularly valuable comments and suggestions from three anonymous referees and the Editor. Funding from the Spanish MEC and the ESRC in the UK is gratefully acknowledged.

[†]University College London, IFS and NBER

[‡]University of Cambridge and IFS

[§]University of Cantabria and Fedea

1 Introduction

Female labour force participation and labour supply, in the US, as in many other developed countries, has changed dramatically over the last 30 years. If one contrasts the labour supply behaviour of the cohorts of women born in the 1930s (such as Elizabeth Dole), 1940s (Hillary Clinton) and 1950s (Oprah Winfrey) in its various dimensions, two main features emerge. First, comparing the Elizabeth Dole cohort to the Hillary Clinton one, we can see a substantial shift of the age profile of labour supply: the Clinton cohort worked more than the Dole cohort. However, the shape of the age profile does not change much. In particular, in both profiles we observe low participation (relative to other ages) corresponding to child rearing years. When comparing the Hillary Clinton cohort with the Oprah Winfrey one, we see that the low participation rates associated with the “fertility years” are no longer present. The aim of this paper is to construct a life cycle model of labour supply and saving that could explain these dramatic changes and, in particular, the difference between the Clinton and Winfrey cohorts. We explore the extent to which realistic changes to some specific parameters and exogenous variables of this model can generate the patterns observed in the data. Or, to use a different perspective, we want to quantify the size of changes in these variables that would be needed to explain the observed patterns.

We build a model of the participation and savings decisions of households across the life-cycle and calibrate this model by matching simulated participation and wage profiles to observed participation and wage profiles. The main change in female labour supply behavior in the data is on the extensive margin. For this reason, we focus on participation choices. Households are able to save and borrow and women choose whether or not to work. Decisions are taken at an annual frequency. In our life cycle model households face uncertainty about wife’s wages and husband earnings; maternity is exogenously given and children impose some monetary fixed cost when mothers decide to work. In the model returns to experience are a result of participation and of the depreciation of human capital when labor market interruptions are made. There is no exogenous aggregate productivity growth, nor are wages determined in equilibrium. Human capital is accumulated only through participation in the labour market. However, observed average wages are endogenous, as they result both from the selection process that induces some women to work and others to stay out of the labour force in the current period and the selection process that operated in the past.

The ability of households to save and to borrow makes our model different from Eckstein and Wolpin (1989), van der Klaauw (1996) and Francesconi

(2002) who impose that consumption coincides with income and estimate structural models of females' employment decision in the first case, females' employment and marital status decisions in the second case, and female employment and fertility in the third case. Without the saving choice, the only way to substitute consumption intertemporally would be through changing labor supply and hence, in a model with returns to experience, the future wage rate. Saving is potentially a more flexible tool for intertemporal substitution and therefore ignoring it might overstate the importance of labor supply choices in life-cycle smoothing. More recently, French (2005) and van der Klaauw and Wolpin (2005) estimate models with a joint saving and labour supply choice, focusing on the retirement decision, while Rendon (2006) estimates a joint saving and labour supply / search model.

To calibrate the parameters of our model, we use observed profiles from the cohort born at the start of the 1940s (cohort 2). We explore the role of different factors in shaping the participation age profile and the age wage profile. The latter will be shaped by the fact that a selected set of women work. Having built a model that represents well the behaviour of the Clinton cohort, we experiment with changes in the basic determinants of labour supply to determine which are more likely to yield the profiles of the cohort born in the 1950s. We consider a number of possible determinants of changes in participation. First, the fixed cost of participating for mothers might have fallen relative to female wages. For example, the costs of child-care may have fallen. Second, on-the-job learning or the return to experience may have increased. As argued by Olivetti (2006), this increases the opportunity cost of reduced labor supply. Third, the level of female wages might have increased, either due to a level shift or to faster growth. Our structural model of life-cycle behavior attempts to evaluate these alternative explanations.

The facts on employment are not in dispute and have been described before. For instance, Pencavel (1998) and Coleman and Pencavel (1993) report age profiles for participation similar to those we present, and McGrattan and Rogerson (2004) show similar year and cohort effects. More controversial is understanding the data on wage profiles and on the process of human capital accumulation, as well as the underlying question of why participation has changed.

Obviously, wages are likely to be an important determinant of female labour supply. However, by looking at observed wages alone, it is difficult to disentangle the return to experience, the depreciation rate of human capital and the extent to which observed average wages are affected by participation bias (selection). Moreover, the interactions of these effects with other important determinants (such as fertility patterns, the cost of children, uncertainty, and so on) even in a simple life cycle model can be quite complex and difficult to quantify.

Olivetti (2006) suggests that changes in wage profiles across cohorts reflect a change in the return to experience. While it appears that the Oprah Winfrey cohort has indeed achieved faster wage growth from early in the life-cycle, it is less clear to what extent this is due to year effects, and it is less clear to what extent the cohort effects reflect a return to experience or general productivity growth. For example, wage growth seems to have benefited those who have worked only intermittently as well as those who have worked full time.

There is now a substantial literature addressing the underlying question of why participation has changed. For example, Olivetti (2006) uses a four period model and estimates of the returns to experience to show the effect that increases in the returns to experience have on hours worked by women. Greenwood and Seshadri (2002) measure the impact of technological progress on the increase in women's participation, while Greenwood et al. (2005) focus on the role played by changes in home production (with the development and diffusion of many household appliances) in explaining the increased labour force participation of women. Jones, Manuelli and McGrattan (2003) investigate the effect on average hours worked by women of the decrease in the wage gender gap as well as the effect of technological progress. Caucutt, Guner and Knowles (2001) explore the interaction between wage inequality, marriage, fertility and employment decisions of young women emphasising the importance of the timing of children for participation decisions.

The contribution of the current paper is primarily to use a realistic life-cycle model of saving and participation to compare alternative explanations. Relative to other contributions, we focus on the changes in the labour supply behaviour of mothers belonging to the cohort of women born in the 1950s relative to that born in the 1940s. In this sense, our paper can be seen as a complement to some of the contributions above. Our main conclusion is that a decrease in child care cost plays an important role in explaining the observed changes in participation rates. However, we need to combine this decrease with an increase in the level of female wages to match the facts concerning the dynamics of observed average wages. Neither factor alone can explain the observed changes in participation and wages for plausible changes in parameters.

The rest of the paper is organized as follows. In Section 2, we describe the data that motivates the paper, in particular female employment behavior and patterns of observed wages. We describe the behaviour of three cohorts: those of women born in the 30s, 40s, and 50s. In Section 3, we describe the model and we compare it with the literature. In section 4, we discuss how we calibrated the parameters of the model and report results for our 'baseline' simulations, meant to describe the behaviour of the cohort of women born in the 1940s. In section 5 we carry out comparative statics exercises for the variables we think are the

main candidates for explaining the changes in female labour force participation. In Section 6, we discuss the implications of our exercise and conclude the paper.

2 Facts to explain

The aim of this section is to illustrate the main facts about female labour supply and about a number of variables that are likely to be important determinants of labour supply choices. Clearly, as we discuss below, some of these variables could be jointly determined with labour supply either at the individual level - such as fertility - or in a general equilibrium setting, such as wages.

The main data sources we use is the PSID. In particular, we use the PSID core sample, including the SEO low income sample. In all our computations we use the PSID weights. As the focus of this paper is a life cycle model, we follow three different cohorts of women over the observed part of their life cycle. The first cohort is made up of women born between 1934 and 1938 and is therefore observed between (median) ages of 35 and 60. The second cohort contains women born between 1944 and 1948 and is observed between ages 25 and 50, and the third cohort contains those women born between 1954 and 1958 and is observed between ages 25 to 40. Sample sizes are reported in Table ?? in the Appendix. While we do not observe the complete life cycle profiles for each cohort, each cohort overlaps, at some ages, with the others. With the important caveat that different cohorts are obviously observed at the same age at different points in time, these overlaps can be informative about possible differences in life cycle profiles. On the other hand, we should keep in mind the impossibility of disentangling, without additional information, year, age and cohort effects.

We concentrate on married women. It is well known that the key factor underlying the increase of female employment in the US over the last decades is the change of married women's behavior. The main issue is whether the increase in the proportion of women who are single or single mothers and the trend towards marrying later might affect or bias our results: changes in participation across cohorts may reflect composition changes rather than changes in behaviour. This selection into marriage (and out of marriage) is not part of our model and its importance remains an open question. We start our descriptive analysis with labour supply variables. We then move on to wages and to other variables, such as fertility, child care arrangements and so on, that might be relevant for labour supply choices.

2.1 Employment

We start our analysis by looking at the life cycle profiles of hours worked. In Figure 1, we plot average hours worked for the three cohorts, averaging over both workers and non workers. In this figure we first see a pattern that we observe repeatedly. Two features stand out. First, there is a large increase across cohorts in the number of hours worked by women, especially if we compare the first and third cohort. Second, the difference in the life cycle profile between the second and third cohort: whilst the difference between the two cohorts is quite large early in the life cycle when participation is much higher for the third cohort, by age 37 the difference between the two cohorts is minimal.¹

In Figure 2, we report average hours worked by women who work. We observe that differences across cohorts are much smaller now, suggesting that the main change in women's labour supply behaviour is in participation decisions. This supposition is confirmed in Figures 3 and 4 which refer to employment rates and full time employment rates respectively. In the former case a woman is classified as employed if she works at least 100 hours per year, while in the latter she is considered working full time if she works at least 1,500 hours per year. Both figures show large differences in employment rates across the different cohorts. Again, it is interesting to note that the main differences between cohorts 2 and 3 are observed from age 25 to age 35. For cohort 2, employment rates are low but increasing from age 25 to age 35, corresponding to child rearing years. However, for the youngest cohort, participation rates are less correlated with women's age.

Next, we focus on the early part of the life cycle and relate labour supply behaviour to fertility behaviour. In Figures 5 and 6, we plot employment rates and full time employment rates for mothers of children younger than 3. As from age 35 on there are very few observations, we restrict our comparison to ages 25 to 35 so we only have observations for cohorts 2 and 3.

The difference between the two cohorts is remarkable. In cohort 3 as many as 66% of mothers with a child less than 3 are working, while the same figure for cohort 2 is only 47%. This evidence is consistent with some facts reported by the US Census Bureau, reproduced in figure 7. They consider women who were mothers in four different periods 1961-65 (cohort 1), 1971-75 (cohort 2), 1981-85 (cohort 3) and 1991-95 and look at employment decisions before and after childbirth. Figure 7 shows that the first two cohorts were unlikely to have

¹It is possible that the increases in participation that we show here are due to year effects, with participation simply being higher in more recent years. This explanation would suggest that wages were higher for the youngest cohort in the early years of working life than for older cohorts. We show in figure 8 below that real wages are, if anything, lower for the younger cohort.




Figure 1: Average Hours Worked




Figure 2: Average Hours Worked, Employed Women Only




Figure 3: Employment Rate




Figure 4: Full-Time Employment Rate

returned to the labor market a year after birth. However, for the two youngest cohorts employment rates three months after childbirth are similar to those two months before.

The women belonging to the three cohorts we are studying are very different in many dimensions. A very important one is their education achievements. The members of the youngest cohort are much more educated than their predecessors: in the Current Population Survey, only 20% of the women belonging to our cohort 1 had more than high school education. This percentage increased to 26% for the second cohort and to 41% for the third cohort. It could be the case that part of the observed increase in women employment rates was due to a composition effect. However, when re-doing the exercise for different education groups, we find that the trends just discussed can be observed at different education levels. This does not imply that changes in education attainment are not part of the explanation through a composition effect, but it does suggest that there is behavioural change for both education groups.²

Figures 5 and 6 showed the importance of the changes in the behaviour of mothers in explaining the different participation rates of cohorts 2 and 3. To complement that evidence, in Table 1 we compute, for cohorts 2 and 3, the percentage of women who exit from the labour market for each age between 26 and 32. The table shows that a smaller fraction of cohort 3 women exit the labour market at each age. Once out of the labour market, women belonging to the two cohorts also differ in terms of the amount of time they stay out of the labour market. In Table 1 we report the median duration of time out of the labour market by age of exit for women who return to work by age 40 and also for all women who exit (ie including those who we do not observe returning). We also include duration statistics pooling together all women who exit at or before age 32. Finally, we report the fraction of women who return to the labor market before age 40. Conditioning on exit, women in cohort 3 return to work somewhat faster than cohort 2, but the difference is small. The main difference between cohorts is in the rate of exit.

2.2 Wages

The price of human capital is determined in equilibrium by the interaction of demand and supply of the relevant factors. For an individual, however, it could be argued that the path of wages is given. As we discuss further below, current wages are not the only important determinant of the participation decisions in a life time framework. Dynamic aspects, such as the return to experience and

²We report the evidence on labour supply (and wages) by education in an earlier (working paper) version of this paper, Attanasio et al. (2004).




Figure 5: Employment Rates of Mothers of Children less than 3 Years





Figure 6: Full-Time Employment Rates, Mothers of Children less than 3 yrs

Women Who Worked Before and After Their First Birth by Month Before and After Birth: Selected Birth Cohorts, 1961-65 to 1991-94



Source: U.S. Census Bureau, Survey of Income and Program Participation (SIPP), 1996 Panel, Wave 2.

Figure 7: Speed of Return After Childbirth

Table 1: Rate of Exit and Duration of Exit

Age	Fraction of workers that exit	Cohort 2			Cohort 3			
		All Median	Duration Mean	Reentering Median	Fraction of workers that exit	All Median	Duration Mean	Reentering Median
26	0.27	3	3.0	2	0.16	2	2.6	1
27	0.15	4	3.2	2	0.12	1	2.1	2
28	0.15	7	4.6	3	0.07	3	3.2	2
29	0.16	3	3.9	2	0.10	4	2.0	1
30	0.20	3	3.2	2	0.14	2	1.6	2
31	0.14	7	4.0	2	0.10	2	2.6	1
32	0.15	4	2.9	2	0.10	4	2.1	1
All		3	3.5	2		2	2.3	1

We report numbers up to age 32 only in order to keep sample size large enough. As cohort 1 is only observed from age 35 onwards, the interesting comparison is between cohorts 2 and 3.

the depreciation of human capital when not participating in the labour market are also likely to be important determinants. For this reason, in this section we look at the life cycle profile of wages for our three cohorts of women.

While life cycle profiles for wages are informative about the return to human capital for women who work, two important caveats should be kept in mind when looking at these pictures and thinking about the role wages could play in determining participation. First, it is not clear whether the observed profile was actually rationally predicted by the decision makers at the time the labour supply decision was made. We already mentioned the existence of macro effects: a future increase (or stagnation) in wages for a given cohort is not necessarily anticipated. Second, the pictures we construct do not necessarily reflect the average (or median) offer wage, which is the one relevant for the decision: selection into employment is not random and can induce important biases.

Figure 8 plots for each of the three cohorts, the median female hourly wage against age. Cohort 3 faces a much steeper wage profile in the early part of the life-cycle than cohort 2. This can be interpreted as an increase in the return to experience. However, it is difficult to separate out cohort and year effects. This difficulty can be appreciated if we plot, as in Figure 9, the median wages for the three cohorts against time, rather than age. If we consider separately the three decades of 70s, 80s and 90s, Figure 9 shows real wages being flat in the 1970s for cohort 1 and 2 (cohort 3 has not yet entered). In the 1980s, cohort 3 experiences faster wage growth than cohort 1 and 2, but in the 1990s all three cohorts experience fast real wage growth. These patterns suggest there is a mix of cohort and aggregate time effects driving wages. Moreover, it should be remembered that the observed wages are affected by selection.³

3 Model

In this section we describe the model we use to explain the changes in female labour supply. We assume that unitary households maximize expected lifetime utility. The utility function is intertemporally separable and instantaneous utility depends on household consumption per adult equivalent and the labour supply choice of the wife. We assume that all households have two adults who remain married and that husbands always work and receive earnings that are determined by a stochastic process introduced below. We do not model fertility choices. However, we calibrate the arrival of children to make it similar to

³ In Figure 8 we average the wages of women with different levels of labour market experience at each age. As women select into and out of the labour market and paid work, the composition of working women (by ability and levels of human capital) changes over time.




Figure 8: Log of Median Hourly Wage by Age




Figure 9: Log of Median Wage by Year

what we observe in the data. We assume that there are three types of women in our model. The first type never has any children, the second and third type both have two children, with the second child arriving two years after the first. However, the second type of women have their first child at 24, while the third type have theirs at 29. We will label these two groups as 'young' and 'older' mothers. We assume that there is no correlation between productivity (male or female wages) and fertility types. Children do not have a direct effect on utility (except for deflating consumption by their adult equivalent). However, they do affect the fixed cost of work.

In particular, we consider an individual household with an instantaneous utility function of the form

$$u_t = u(c_t, P_t, e_t)$$

where P_t is a discrete $\{0, 1\}$ female labour supply choice, c_t is total household consumption and e_t is the number of adult equivalents in the household. The household is assumed to maximize lifetime expected utility,

$$\max_{c, P} V_t = E_t \sum_{s=t}^T \beta^{s-t} u(c_s, P_s; e_s)$$

where β is the discount factor and E_t the expectations operator conditional on information available in period t . We consider a retirement period after which neither of the members of the household participate in the labor market and the household receives a pension proportional to the husband's earnings after retirement in period $T - r$. One of the reasons for including the retirement savings motive in our model is to have a realistic amount of savings so that the potential role of female labor supply as a private insurance mechanisms is not overestimated. We use a utility function of the form

$$u(c_t, P_t; e_t) = \frac{\left(\frac{c_t}{e_t}\right)^{1-\gamma}}{1-\gamma} \exp(\psi_1 P_t) - \psi_2 P_t \quad (1)$$

As we use values of $\gamma > 1$, we constrain $\psi_1 > 0$ so that participation reduces the utility of consumption. The functional form in 1 purposefully deviates from the homothetic specifications typically used in macro growth models that generate balanced path. Consistent with estimates from micro studies, such as Browning and Meghir (1991), utility is not homothetic. Moreover, given that the focus of this paper is in understanding the reaction of labour supply to different types of incentives, we did not want to build into the specification of the utility function a restriction on the response of labour supply to changes in wages. Consumption is equalised by the factor e_t which depends on the age and number of children.

We use the McClements scale to determine e .⁴

The intertemporal budget constraint has the form

$$A_{t+1} = R \left(A_t + \left(y_t^f - F(a_t) \right) P_t + y_t^m - c_t \right) \quad (2)$$

where A are beginning of period assets, R is the interest rate, F the fixed cost of work which depends on a_t , the age of the first child born to the female. Female earnings are given by y_t^f , and husband earnings are given by y_t^m . In any period, individuals are able to borrow against the minimum income they can guarantee for the rest of their lives. Notice that this feature differentiates our model substantially from those used by Eckstein and Wolpin (1988) and van der Klaauw (1996) who rule out any borrowing or saving. More recently, Rendon (2006), French (2005) and van der Klaauw and Wolpin (2005) estimate labor supply models jointly with savings.⁵

We denote the child care units needed by a family whose first child is age a_t by $G(a_t)$ and the price of each unit of child care by p . Therefore, the total child care cost faced by a household when females participates in the labor market is given by

$$F(a_t) = pG(a_t) \quad (3)$$

The function $G(a_t)$ will be estimated from expenditure data of households with children of the relevant ages.

Female earnings are given by

$$\ln y_t^f = \ln y_0^f + \ln h_t^f + v_t^f \quad (4)$$

where h_t is the level of human capital at the start of the period and v_t^f is the permanent productivity shock.

Human capital evolves with employment decisions in the following way

$$\ln h_t^f = \ln h_{t-1}^f + (\eta_0 + \eta_1 t) I(P_{t-1} = 1) - \delta I(P_{t-1} = 0)$$

$$\eta_0 > 0, \eta_1 < 0$$

We think of δ as the permanent depreciation in human capital associated with non-participation, as discussed and estimated by Mincer and Polachek

⁴According to the McClements scale, a childless couple is equivalent to 1.67 adults. A couple with one child is equivalent to 1.9 adults if the child is less than 3, to 2 adults if the child is between 3 and 7, 2.07 adults if the child is between 8 and 12 and 2.2 adults if the child is between 13 and 18. As we mention in the text, we assume that each couple has two children who arrive at a predetermined age and leave at age 18.

⁵The introduction of borrowing and saving combined with the heterogeneity in types and the persistent of shocks means that the model becomes difficult to estimate structurally because of the computational time in each iteration.

(1974) and Mincer and Olfek (1982). We do not model direct investment in human capital (such as schooling decisions or on the job training), which are extensively discussed in Mincer and Polachek (1974) and Mincer and Olfek (1982). However, we impose a floor on how far human capital can decline so that human capital will not fall below its initial value. This means that the marginal loss of human capital associated with a year out of the labour force becomes zero for a sufficiently long unemployment spell.

The process of human capital accumulation is important to our model. One issue is whether the increase in human capital associated with working diminishes with the level of human capital. A related issue is whether this increase depends on age. Eckstein and Wolpin (1989) and van der Klaauw (1996) assume that the increase in human capital diminishes with the level of human capital. Olivetti (2006) assumes the increase in human capital diminishes with age, but not with the level of human capital. As in Olivetti, we assume that the increase in human capital depends on age only, with the increase in human capital decreasing with age if $\eta_1 < 0$. We calibrate the depreciation rate and the human capital function parameters so that our model match certain moments of the data that we discuss below.

Since we assume men always work, male earnings are given by

$$\ln y_t^m = \ln y_0^m + h_t^m + v_t^m \quad (5)$$

$$h_t^m = \alpha_1^m t + \alpha_2^m t^2 \quad (6)$$

We estimate directly from the data the parameters of the human capital accumulation function for males. As for married men labor market participation is very high, selection bias in this estimation is expected to be small.

Both female and male earnings, y_t^f and y_t^m , in the household are subject to permanent shocks, v_t^f and v_t^m , that are positively correlated. In particular we assume

$$\begin{aligned} v_t^f &= v_{t-1}^f + \xi_t^f \\ v_t^m &= v_{t-1}^m + \xi_t^m \quad \text{where } \xi_t = (\xi_t^f, \xi_t^m) \sim N(\mu_\xi, \sigma_\xi^2) \end{aligned} \quad (7)$$

$$\mu_\xi = \left(-\frac{\sigma_{\xi^f}^2}{2}, -\frac{\sigma_{\xi^m}^2}{2} \right) \quad \text{and} \quad \sigma_\xi^2 = \begin{pmatrix} \sigma_{\xi^f}^2 & \rho_{\xi^f, \xi^m} \\ \rho_{\xi^f, \xi^m} & \sigma_{\xi^m}^2 \end{pmatrix} \quad (8)$$

This is the only uncertainty that households face. They are assumed to have perfect foresight regarding fertility, childcare costs, the process for human capital accumulation *and the fact that they will remain married*.⁶

⁶When married women face uncertainty about their future marital status and the accumulation of labor market experience has a return in terms of higher future wages, they have an additional incentive to participate in the labor market. Several papers support this link

In each period, if the woman chooses to participate, the value function is given by

$$\max_{c_t} \left\{ u(c_t, P_t = 1; e_t) + \beta E_t \left[\max \left\{ \begin{array}{l} V_{t+1}^0 \left(A_{t+1}, v_{t+1}^m, v_{t+1}^f, h_{t+1}^f \right) \\ V_{t+1}^1 \left(A_{t+1}, v_{t+1}^m, v_{t+1}^f, h_{t+1}^f \right) \end{array} \right\} \right] \right\} \quad (9)$$

If she chooses not to participate, the value function is given by,

$$\max_{c_t} \left\{ u(c_t, P_t = 0) + \beta E_t \left[\max \left\{ \begin{array}{l} V_{t+1}^0 \left(A_{t+1}, v_{t+1}^m, v_{t+1}^f, h_{t+1}^f \right) \\ V_{t+1}^1 \left(A_{t+1}, v_{t+1}^m, v_{t+1}^f, h_{t+1}^f \right) \end{array} \right\} \right] \right\} \quad (10)$$

The decision of whether or not to participate in period t is determined by comparing $V_t^0 \left(A_t, v_t^m, v_t^f, h_t^f \right)$ and $V_t^1 \left(A_t, v_t^m, v_t^f, h_t^f \right)$. The participation choice and the consumption choice in t determines the endogenous state variables (assets and human capital) at the start of the next period. The non-concavity in the value function induced by the discrete participation decision is smoothed out by the presence of sufficient uncertainty. We check that this holds in the numerical solution of the problem, as discussed in the appendix.

One of the main differences between our model and those estimated in the literature by Eckstein and Wolpin (1988) and van der Klaauw (1996) is the inclusion of saving and borrowing as a choice variable. While allowing individuals to save adds an element of realism to our model, the assumption that they can borrow up to the present discounted value of their minimum earnings can be an unrealistic assumption. It is possible that individual households are prevented from borrowing large amounts at the current interest rate because of imperfections in credit markets. The presence of such borrowing constraints is likely to have a particularly strong effect on young households. In Attanasio et al. (2005), we use the model in this paper to show how much borrowing constraints matter for female participation.

between marital risk and employment of married women, Peters (1986), Parkman (1992), Sander (1985), Sen (2000) and Sánchez-Marcos (2002). According to OECD figures, divorce rates in the US have increased from 2.2% in 1960 to 4.8% in 1990. In this paper, we do not consider at all the effects of divorce risk.

4 Baseline Parameters and Simulations

In this section, we discuss the choice of the parameters we use in our baseline model. We then show the life-cycle profiles and other statistics implied by these parameters. We calibrate the model parameters to fit the life cycle profile of cohort 2. In the next section, we carry out comparative statics exercises and discuss the implications of changing some key baseline parameters for female participation.

4.1 Baseline parameters

We take some of the parameters of our model directly from pre-existing estimates and direct estimates from the data. We calibrate the remaining parameters to match some moments of the data. We start by discussing the first set.

External parameters In Table 2, we report the set of ‘exogenous’ parameters we use in our baseline simulations. We fix the interest rate to equal the average real return on three monthly T-bill at 0.015. We assume a discount factor equal to 0.98, which implies a discount rate slightly higher than the interest rate. In the utility function (1), the coefficient of relative risk aversion, γ , is set to 1.5. This value is consistent with the evidence on the elasticity of intertemporal substitution in the US provided by Attanasio and Weber (1995).

As mentioned above in our model women differ in the age at which they have their first child. We assume that 12% of women never have a child, 41% have the first of their two children at age 24 and 47% have their first child at 29. We have calibrated these proportions on the distribution of maternity age of women belonging to cohort 2 in the PSID. All women in our model begin life at age 23 with zero assets.

The deterministic component of the male earnings process is estimated from the PSID. Both the innovations to male earnings and those to female wages are assumed to have a unit root, consistent with the evidence on men produced by MaCurdy (1983) and Abowd and Card (1989). The degree of persistence of income shocks is important. If shocks were i.i.d. but with a high variance per period (this is necessary to keep the variance of lifetime earnings constant), participation is high across the life-cycle as individuals face large amounts of ongoing uncertainty. With persistent shocks, the uncertainty translates into heterogeneity late in life. On the other hand, our simulations show that uncertainty does not have substantial effects on the timing of participation (in contrast to the large effects that uncertainty has on the timing of consumption). We explore the implications of uncertainty for participation more fully in Attanasio et al. (2005).

Parameters		
Discount Factor	β	0.98
Interest Rate	R	1.015
Regression Log Wage on Age and Age ² (Men)	α_1^m, α_2^m	0.034, -0.00033
Husband and Wife Wage Correlation	ρ	0.25
Standard Deviation of Permanent Shock (Men)	σ	0.13
Standard Deviation of Permanent Shock (Woomen)	σ	0.13
Risk Aversion	γ	1.5
Length of Life	T	50
Length of Working Life	r	40

Table 2: Exogenous Parameters

The standard deviation of the innovation for husband’s earnings is assumed to be 0.13. This number is similar to the variances estimated using PSID data by Carroll and Samwick (1997) or Low, Meghir and Pistaferri (2006). There is not much evidence on the variability of female wages and/or earnings. Assuming that the variance of female wages innovations is the same as that for men earnings, we compute the implied coefficient of variations of female earnings, which comes out at 0.69. Obviously, this is very different from the coefficient of variation of male earnings, because of endogenous participation. However, it turns out to be quite similar to the value of 0.65 reported by Hyslop (2001) for female earnings.⁷ In addition, for our baseline cohort in the PSID, the same coefficient of variation is not very different at 0.77. We therefore settled for the value of 0.13. We assume that the correlation coefficient between the two shocks (for husband and wife) is equal to 0.25 as estimated by Hyslop (2001).

There are two components to child care cost: the function $G(a_t)$ and the price p . We estimate the function $G(a_t)$ directly from data. In particular, for households where the mother is working, we regress total childcare expenditure on the age of the youngest child, the age of the oldest child, the number of children and a dummy that equals one if the youngest child is 0. The shape $G(a_t)$ can be derived from the coefficients of this regression function, considering that in our model all women with children have two of them and at the same interval. This implies that the child care cost can be expressed as a function of the age of the oldest child. As we discuss below, we calibrate the price p using our model. However, we do not model the choice over quality of childcare and we rule out the possibility of using unpaid childcare (eg other family members) which may be important, particularly for low income families (Blau, 2001).

⁷Hyslop (2001) assumes a different process for wages, as he include individual fixed effects in the process for female wages, rather than persistent innovation.

Calibrated parameters In the top panel of Table 3 we report the set of parameters that we have calibrated to match certain statistics of the data. In the bottom panel, we report the value of the statistics we chose to calibrate in the data, along side their value in the calibrated model implied by the chosen parameters. All the statistics refer to cohort 2. Obviously, when one changes one parameter, all simulation statistics change so that the mapping suggested in the columns of Table 3 is only approximate and indicative. However, it is clear that some parameters are crucial in determining certain features of the simulations.

The parameters of the model we calibrate are: the utility cost of working (given by 2 parameters), the depreciation rate (defined as the annual decline in human capital while not working), the initial offer wage gender gap, the returns to experience and the price of child-care. The moments we are matching in the data are computed for women aged between 25 and 50 and belonging to cohort 2. We match the following moments:

1. The average participation over the life-cycle
2. The average participation of women who have no children under 18. This includes women before they have children, women after their children have left home and women who never have children.
3. The average participation rate by mothers with children aged under 3.⁸
4. The observed wage gender gap for women aged under 40.
5. The wage loss observed for women on re-entry, for those women exiting the labor market before 32 and coming back before 40. The statistic is the exponential depreciation rate calculated as:

$$D_{t,k} = \frac{1}{k} \ln \left(\frac{w_t}{w_{t-k}} \right)$$

where k is the duration of an exit that began in period $t - k$.⁹

6. The wage growth of wages for women who have worked 90% of their lifetime at each age for 2 groups of women: those who are younger than 35 and those who are older than 35. For wage growth we match, for the two groups of women, the parameter θ_1 in the following regression

$$\log y_t^f = \theta_0 + \theta_1 t + u_t$$

⁸In the simulations, the participation rate of mothers is computed considering the participation of mothers within 5 years of the birth of their first child because the second child arrives when the first child is 2.

⁹We follow women who have employment interruptions and who are observed before and after the interruption and we control for the duration of non participation. This calculation is similar to the one performed by Mincer and Polachek (1974) and Mincer and Olfek (1982).

It should be stressed that in computing all the statistics that involve female labour market outcomes, we implicitly control for the selection process that is operative in our model, that is both in the data and in the simulations, the women on which these statistics are computed are a selected sample.

Parameters		Value
Utility Cost of Participation	ψ_1	0.038
Utility Cost of Participation	ψ_2	0.0020
Childcare Cost	p	12791
Depreciation Rate	δ	0.074
Offered Wage Gender Gap	y_0^f/y_0^m	0.64
Returns to human capital	η_0	0.0266
Returns to human capital	η_1	-0.00038
<hr/>		
Targets in the Data	Cohort 2	Model
Participation	0.71	0.72
Participation Mothers Children Aged under 3	0.47	0.47
Participation Non Mothers	0.83	0.83
Median Annualised Wage Loss for Non-Participants	0.05	0.02
Observed Wage Gender Gap	0.69	0.69
FemaleWage Growth (if younger than 35)	0.020	0.023
FemaleWage Growth (if older than 35)	0.009	0.006

Table 3: Calibration Process of Model Parameters

We calibrate the price of child care to be 12791. This implies that the child care cost associated with a new born is 68% of mean earnings for a 30 year old woman who worked continuously prior to childbirth. Child Care costs in our model are clearly substantial and this accords with other direct evidence on childcare costs, such as Schulman (2000), who reports data on child care costs in the 50 US states. In 1999, according to the Child Care Bureau of the US Department of Health and Human Services, Child Care costs varied between \$300 and \$700 per month per child. In certain areas of the country these costs could be even higher, especially for families with infants, whose cost is substantially higher than that of children aged more than 3. In the same year, median earnings of female workers was, according to the Census Bureau, \$1922 per month. Therefore for a woman earning average wages with 2 children and facing a cost of \$500 per child, child care cost was around 53% of her earnings.

The ψ_2 parameter reflects the direct utility cost of participating and it takes

a value equal to 0.0020 in our baseline. We could alternatively have a fixed cost of working in the budget constraint that is not child related. The ψ_1 parameter reflects the reduction in the utility of consumption caused by participation. Since $\gamma > 1$ and $\psi_1 > 0$, the marginal utility of consumption is greater when participating than when not participating, in other words, consumption and participation are complements in utility, as in van der Klaauw. Parameter $\psi_1 = 0.038$ in our baseline corresponds to a utility cost of participation equivalent to 7.3% of consumption.

The ratio of initial offer wages for men and women is set to 0.64 and is lower than the observed gender gap. This implies that offered wages lie below accepted wages for women. These selection effects are an important part of our model and are discussed further below.

The depreciation parameter, δ , is equal to 0.074. This is high relative to Mincer and Olfek (1982) who estimate $\delta = 0.02$. However, the implied annualised wage loss from non-participation in our model is 0.02. This lower annualised wage loss arises because the impact of depreciation on human capital is constrained by the assumption that human capital cannot fall below its starting value. Selection goes also in the direction of decreasing the observed annualised wage depreciation.

The parameters η_0 and η_1 measure the return to experience for those participating, giving a return to one year of participation of about 2%, which is similar to Eckstein and Wolpin (1988). Olivetti (2006) estimates the return to one extra year of full-time work at between 3 and 5%, based on the return to extra hours work rather than the return to participation. One implication of the presence of a return to experience is that individuals may choose to work early on despite current earnings being less than the fixed cost, particularly if they are able to finance this investment through borrowing.

4.2 Baseline life cycle profiles

Given the parameters in Table 3 we can simulate the model and generate life cycle profiles for the variables of interest. We can also calculate statistics that have not been used in the calibration to compare to the data of cohort 2. We focus our discussion on the labour supply behaviour of women who differ in their fertility experience. In particular, we consider childless women, young mothers (age 24 at childbirth) and older mothers (age 29 at childbirth).

Table 4 reports statistics for participation for the three groups of women in the model and compares these to the data. Figure 10 shows corresponding participation profiles. The simulated model is able to match not only the participation statistics that were used in the calibration but also other salient statistics

Statistics	Model	Data
Participation Mothers with Children Aged under 3	0.47	0.47
Participation Young Mothers with Children Aged under 3	0.30	0.42
Participation Old Mothers with Children Aged under 3	0.59	0.53
Participation Mothers with Children Aged 4-18	0.71	0.70
Median Duration in Years	7	7

In the data, a fraction of women return to part-time work in between the birth of the first and second child. Reported median duration is calculated for return to full-time work

Table 4: Other Statistics in the Baseline Simulations

describing female labour market behaviour: first, as in the data, young mothers in the model participate less than older mothers; second, we match almost exactly the participation rates of mothers of older children. Further, median duration out of full-time employment is 7 years in the model and 7 years in the data.¹⁰




Figure 10: Simulated Participation Profiles by Age at Childbirth

¹⁰We report the statistic from the data for duration until return to full-time employment. In the data, as reported in Table 1, duration until return to any employment is faster, often because women return to work temporarily between the birth of their children. The effects of these short return spells are discussed in Mincer and Olfek (1982).

Quantifying Selection Figure 11 shows life-cycle profiles of log earnings averaged over women who are participating, by fertility experience. The figure shows that, following childbirth, only the more productive mothers remain in the labour force. This explains the hump in the earnings profile which occurs at different ages for mothers with different maternity ages. The figure also shows how child care costs compare to realised earnings. It is difficult to compare these profiles to the data directly because, given our sample sizes, we cannot restrict to women who have exactly the same maternity experience as women in our model. This picture, however, illustrates the potential importance of the selection bias induced by the participation decision following childbirth and has implications for estimates of the return to experience and of wage dispersion that are based on observed wage data.




Figure 11: Simulated Earnings Profiles by Age at Childbirth

There are two forces underlying the selection effects. First, only women with a certain productivity enter the labour market and are observed. Second, those who enter increase their wages through the return to experience, so that selection impacts on the observed wage dynamics as well as on the observed wage level. We use our model to quantify the extent to which selection into the labour market matters for the life-cycle profiles of observed wages compared to offered wages and to disentangle the two effects. In Figure 12, we show the average offer wage (for all women) by age and compare it to the average wage for working women and to the average wage for women who have worked at least 90% of the time. Observed wages lie above offered wages reflecting the selection

effect on the level of the wage. Observed wages for those who have worked at least 90% of the time lie above the wages for all women: this confounds the selection effect and the return to experience.




Figure 12: The Effect of Selection on Wages

To disentangle the pure selection effects from the indirect effect induced by the process of human capital accumulation, we simulate the effect of a 10% fall in childcare costs. This induces an increase in participation relative to the baseline simulations. Note that the exogenous change in child care cost does not affect directly the offer wage. However, new participants will start accumulating human capital and this will affect their offer wage in subsequent periods. To assess the magnitude of these effects we start by plotting in Figure 13 three lines: (i) the average wages observed for all women participating with the lower childcare costs, (ii) the average wage profile observed with the lower child care costs for those women that were participating even in the baseline (who we will call ‘original participants’) and (iii) the average observed wages for the women who are induced to participate by the lower childcare costs (who will call ‘new participants’). The labour supply behaviour of the ‘original’ participants is not affected by the change in child care cost and neither are their observed wages. The overall average wage, however, is lower because less productive women now enter the labour market. This is made clear by the lower level of the third line. The difference between the first and second line is not large because the new participants account for a relatively small fraction of the total. The difference in average wages for the ‘new’ participants those for the ‘original’ participants reflects both the selection and the new process of human capital accumulation. A slightly surprising feature of the average wage for the ‘new’ participants is




Figure 13: The Dynamics of Selection

the fact that it declines considerably after age 28-30.




Figure 14: Offered wage for new participants

To disentangle the selection effects from the human capital effects and to better understand the decline in the wage of the marginal participants, we plot in Figure 14 the offered wage of the ‘new’ participants ‘before’ and ‘after’ the decrease in child care cost, separately for young and older mothers. The difference between the two lines in each figure is induced by the changed return to experience. The solid line corresponds to the average offered wage in the baseline, while the dotted line corresponds to the average offered wage with lower child care costs. Wages are lower for older mothers because, among older mothers,

fewer mothers exit and so average productivity of those who exit is lower. For both young and old mothers, the wage level is the same at the point of childbirth because work histories are almost identical at that point. After childbirth, offered wages are higher in the environment with lower child care costs when the group is participating and earning the return to experience. Indeed this difference is due only to the return to experience. The decline in the offered wage occurs for two reasons: first, there is wage depreciation in the baseline when the group is not participating; second, there is a composition change in the group because women have different durations of time out of the labour force: less productive workers return to work more slowly. Given the parameters of the model this effect is more visible for younger mothers.

The pictures we have discussed here show the importance and complexity through which selection affects observed wages in our model. The last experiment is particularly interesting as it shows that average wages can be affected quite substantially simply by a change in the composition of working women that is induced by a change in a parameter unrelated to wages or productivity.

Properties of Preferences The preference specification we use is not standard and it is useful to characterise the implications of these preferences for how responsive individuals are to changes in their environment. We do not impose the restriction of balanced growth preferences precisely so that we do not predetermine how individuals respond to these changes in, for example, wages. We look at the implication of our preferences by computing the elasticity of participation to child care costs and comparing this elasticity to other estimates in the literature.

Changes in child care costs will be one of the central hypotheses we propose to explain the change in female labour supply in the next section. To compute the elasticity of participation to child care cost implied by our model we simulate for many individuals the effect of a change in these costs and compute participation rates before and after the change. Given the nature of our model, the child care cost elasticity will be dependent on the age of the child and the age of the mother at childbirth. We compute the elasticity for mothers of children age 0-6 for comparability with the literature and we distinguish between the responsiveness of young mothers and that of older mothers. We compute the elasticity for young mothers to be -0.98 and that for older mothers to be -0.82. How do these values compare with values reported in the literature? Blau (2003) discusses carefully alternative estimates of this elasticity. Estimates vary substantially depending on the methodology, with some estimates greater than 1 and others at 0. For married women with children less than 6, Anderson and Levine (2000), for example, estimate an elasticity of -0.46. Closer to our num-

bers are the results in Connelly and Kimmel (2000) who estimate an elasticity of -0.71.

5 Explaining changes in female participation

In this section, we describe the impact that various changes in the economic environment have on female participation over the life-cycle. Our main aim is to establish what are the most likely explanations for the change in the shape of the life cycle profiles of cohort 3 relative to cohort 2. In particular, we will focus on changes in child-care costs, return to experience and the level of the wage. For each of these variables, we first provide some evidence on the observed changes in these variables between cohorts 2 and 3 and then simulate the behaviour of a cohort that faces these changes relative to the baseline scenario we presented in the previous section. We then compare the resulting labour supply behaviour to that of cohort 3. The statistics reported in this section differ slightly from those used in the calibration: first, we report statistics only for women age 40 or under because we do not have more recent data for cohort 3. Second, we report more disaggregated statistics on participation showing how participation differs by the age of child and by age at childbirth and we compare these statistics to the data. Finally, we do not report changes in the wage loss following time out of the labour force. These numbers are very noisy in the data and our model does not do a good job of matching the observed fall even in the baseline.

We do not consider differences across cohorts in the number of children or in the age at childbirth as strong candidates for explaining the change in labour supply. The main reason for this decision is that, although there are large changes in the level and timing of fertility between cohorts 1 and 2 (the Clinton cohort has fewer children and later in the life cycle than the Dole cohort), the changes in this dimension between cohorts 2 and 3 are small.¹¹

5.1 Child-care costs

In section 2, we stressed that the main difference in the labour supply of women in cohort 2 and 3 is caused by the increase of mothers' employment. It is

¹¹The reduction in the number of children happened at the same time as a delay in the birth of the first child. Using the PSID 1993 additional fertility module, we can calculate, for each of our three cohorts, the proportion of women who have their first child in any given age interval. In cohorts 1 and 2, 30% of women had their first child between 18 and 22; this percentage falls to 26% for cohort 3. On the other hand, while only 29% of cohort 1 women had their first child when aged over 26, that percentage goes up to 37% for cohort 2 and to 41% for cohort 3. These proportions are similar to the data from the National Vital Statistics System published by the US Department of Health and Human Services (DHHS, 1989).

therefore natural to look at the cost and availability of child care arrangements and to consider how changes in these costs might impact on participation.

5.1.1 Measuring the exogenous change across cohorts

In section 4, we discussed the high level of childcare costs that we calibrated for cohort 2. To understand the role of child care costs in explaining the change in participation across cohorts, we need to know how much these costs have changed in relation to the potential earnings of women. Unfortunately, it is not easy to find consistent data on the child care costs for the period we study. We use two sources of information to provide evidence on the dynamics of child care costs relative to female earnings. First, we consider detailed information on child care expenditure that can be obtained from the BLS Consumer Expenditure Survey, and second, we consider indirect evidence on the unit cost of child care.

The CEX collects detailed information on consumption expenditure by US households. For our exercise, in which we compare the women born in the 1940s to women born in the 1950s, we need data from the 1970s, which are not available: data is only available on a consistent basis since 1980.¹² However, we use the CEX to give an indication of the trend in child care costs. We consider households observed between 1980-1984 and between 1998-1999, where the reference person or the spouse is a working woman with children between 0 and 7. We construct the ratio of child cost expenditure to female earnings and run a Tobit regression of this ratio on the log of the annual number of hours worked, the number of children between 0 and 3 and the number of children between 4 and 7 years of age. In addition to these variables, we introduce a dummy for the late 1990s year. In Table 5 we report the results of two regressions: the first for mothers with at least one child aged 0 to 3, the second for mothers with at least one child aged 4 to 7. Notice that the two samples are partly overlapping. In both columns we see that the coefficient on the late 1990s is significantly negative, indicating a decline in the ratio of child care costs to women earnings, even after controlling for differences in the number of hours worked. The decline is particularly large for younger children. A similar exercise can be executed in the PSID, although the number of observations is small. If we compare the ratio of child care cost to women earnings in 1975 and 1988 we find that, conditional on having positive costs, the ratio is considerably

¹²The 1972-1973 CEX used very different definitions of most commodities, including the components of child care. The numbers between surveys, therefore, are not comparable. The SIPP is an alternative data source, which has been used by several authors, including Anderson and Levine (2000) (who use waves from the 1990s) and Ribar (1992) who uses the 1984 wave. However, as with the CEX, we do not have information on childcare costs for cohort 2.

higher in 1975 than 1988, even after controlling for the number of hours and number of children.¹³

Table 5: Evidence for Changes in Childcare Costs

<i>Tobit regression, Dep. Var.: Childcare cost / woman's earnings</i>		
	<i># children age 0-3 > 0</i>	<i># children age 4-7 > 0</i>
Log hours	-0.029 (0.006)	- 0.023 (0.003)
# of children 0-3	0.057 (0.017)	0.076 (0.006)
# of children 4-7	0.035 (0.011)	0.048 (0.009)
Year>1997	-0.236 (0.036)	-0.050 (0.012)
constant	-0.163 (0.048)	-0.123 (0.028)
Number obs	9201	10438
Pseudo R^2	0.0098	0.0375

CEX data.

Direct evidence on the cost of child care, especially historically, is difficult to obtain. Some insights on the evolution of these costs can be obtained by considering the wages of child care workers. According to O'Neill and O'Connell (2001) employees in child care establishments have fared very poorly over the past 20 years. Figure 15 reproduces data from the US Census Bureau on payroll costs for employees working in the child care sector.

The average payroll dollars per employee in 1997 was \$11,076, which is \$548 less than in 1977 (in 1997 dollars). This fall of about 5% comprises an initial fall of about 15% that occurred between 1977 and 1982, followed by a partial recovery. Further, between 1982 and 1997, in the economy-wide recovery of wages after the high inflation period at the end of the 1970s, wages for all female workers increased by 79% (from \$13,366 to \$16,849) compared with the 11% increase (from \$9,690 to \$11,076) for child care employees. Given the importance of labour cost in total child care costs, this evidence would imply a relative decline of the latter. Women in cohort 2 would have incurred child care costs in the 1970s, while women in cohort 3 would have incurred costs in the

¹³Interestingly, the fraction of working women reporting a positive amount of child care costs is higher in 1988 than in 1975.

1980s when the level of costs had fallen by as much as 15% and the cost relative to female earnings had fallen even further. The fiscal treatment of Child Care Costs also changed over the period. In 1954 a deduction for employment related care expenses was established. The deduction became a credit in 1976 and in 1981 the limits were \$2400 for one child and \$4800 for two or more.¹⁴

Consistently with this evidence, Ahn and Mira (2002) argue that there was a decline in the price of child care relative to female wages based on Blau (1992). A possible explanation for the decline in child care cost is discussed in Ferrero and Iza (2002) who argue that recent skill-biased technological change, which implies an increase of the skill premia and a relative decrease in the market price of child-caring with respect to female mean wages, could contribute to explaining the increase in women’s employment rate.

Another indication of the ‘cost’ of child care is its ‘availability’. Figure 16 uses data from the US Census Bureau to show child care arrangements used by mothers over time. The figure shows the large increase in the use of organized child care facilities, which seems to indicate their increased availability. An interesting fact is that until 1986/87 child care workers were not considered as a separate occupation in The Occupational Outlook Handbook that the Bureau of Labour Statistics publishes every other year.

While this evidence clearly indicates a decline in child care costs, it is difficult to quantify the size of such a decline. Moreover, it is likely that there were changes also in the quality of child care that would constitute an additional problem in quantifying the change in its cost. Blau (2001, 2003) has stressed that part of the problem with child care is the prevalence of low quality childcare, and the fall in childcare cost is an indication of this fall in quality. We do not pursue the issue of a decline in the quality of childcare in our experiments, but we try different hypotheses about the size of the decline in the price.

5.1.2 Experiment

In Table 6 we show how participation reacts to the price of child-care in our model. The first column of the Table reports the change in the statistics of interest between cohort 2 and cohort 3. The remaining columns report the simulated change relative to the baseline when we decrease the cost of child care, p , by 5%, 10%, 15% and 20%. Given the evidence above on the evolution of child care workers pay relative to other workers, a decrease in 15% is not unreasonable. We see that decreasing child care costs by 15% increases total

¹⁴In 2002 these figures were raised to \$3000 for one child and \$6000 for two. These increases do not even make up for inflation. Family income determines the percentage of child care expenses that can be claimed for credit (between 20 and 35%). The credit is not refundable which keeps the lowest-income families from benefiting from it.

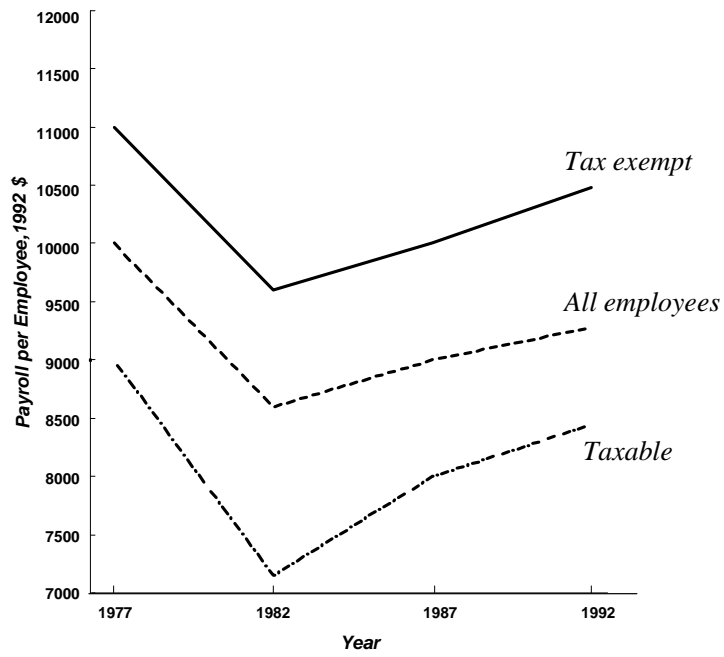


Figure 15: Annual Payroll for Employees in Child Care Establishments




Figure 16: Child Care Arrangements

participation by 6% points, which compares with an increase of 12% points in the data. When we look at the different groups of women in our simulations, we find that a 15% decrease in child care cost is able to generate changes in each of these groups that are similar to what we observe in the data: for young mothers, we observe a 22% points increase in the data and an 18% points increase in the simulations; for older mothers, we observe a 17% points increase in the data and an 11% points increase in the simulations. Our simulations match closely the increased participation rates of mothers aged 4 to 18 which increases by 8% points in the data and 6% in the simulations. With a fall in child care costs of 20%, the match on participation does even better, but this fall is larger than we have evidence for. Furthermore, on the negative side, we should stress that the experiment is not successful in capturing changes in the observed wage growth¹⁵ or the gender gap.

Child Care Cost (p)	Δ Cohort	-5%	-10%	-15%	-20%
Δ Participation All Women	0.12	0.02	0.04	0.06	0.07
Δ Part. Mothers Child 0-3	0.20	0.05	0.09	0.14	0.18
Δ Part. Mothers (Age < 29) Child 0-3	0.22	0.06	0.11	0.18	0.23
Δ Part. Mothers (Age \geq 29) Child 0-3	0.17	0.05	0.07	0.11	0.14
Δ Part. Women with no Child < 18	0.06	0.00	0.00	0.00	0.00
Δ Part. Mothers Child 4-18	0.08	0.02	0.04	0.06	0.08
Δ Wage Gender Gap	0.07	0.00	0.00	0.00	-0.01
Δ Female Wage Growth (Age \leq 35)	0.021	0.00	-0.003	-0.004	-0.006

Table 6: Decreasing Child Care Costs

5.2 Returns to experience

Olivetti (2006) presents evidence on the change in the return to experience between the 1970s and the 1990s. She shows that estimated changes in the return to experience can explain changes in hours worked across this time frame. In this section, we explore whether a change in the return to experience might explain the observed changes in participation.

5.2.1 Measuring the exogenous change across cohorts

The evidence in section 2 suggests that changes in wage profiles across the two cohorts are driven partly by year effects common to both cohorts, and partly by

¹⁵We do not report wage growth for older women because cohort 3 is only observed until age 40.

cohort specific increased returns which may be caused by an increased return to experience. In order to have additional information on the returns to experience we plot, for each cohort, two life cycle wage profiles. At each age we compare the wage of all women with the wage of those women who have been observed since age 25 and have not had more than one year off work. These two profiles, plotted in Figure 17, are observed from age 26 to age 39 and give an idea of the return to experience faced by these cohorts. Computing the ‘return to experience’ as the difference between these two profiles has an important advantage relative to the profiles in Figure 8 and some drawbacks. The advantage is that, to a certain extent, year effects are common to the two profiles and might be ‘differenced out’. However, we do not take into account the duration of spells out of the labour force and we know these to be shorter for cohort 3. Moreover, we ignore selection (and possible changes in it). Perhaps surprisingly, in these figures we do not find very strong differences in this particular measure of return to experience between cohort 2 and 3. On the other hand, we again see the pattern that wage growth for all women in cohort 3 has been faster than for cohort 2. Since the profiles in Figure 17 are obviously affected by the interaction of the selection in and out of employment, changes in productivity and depreciation, they cannot be used formally to disentangle the various factors at play. As we know that more women participate in cohort 3, selection will play a greater role in affecting the difference between the two profiles shown in the figure. The fact that this difference is similar between the two cohorts is not inconsistent with the hypothesis that cohort 3 enjoys a higher return to experience.

5.2.2 Experiment

As suggested by Olivetti (2006), an increase in the returns to experience may affect labour supply. She estimates an increase of the returns to experience of about 25%. In Table 7 we report the same participation statistics reported in Table 6 for different returns to experience. As with Table 6, in the first column we report the change in the statistics between cohorts, while in columns 3 to 5, we report the statistics for simulations where we increase the parameters of the human capital accumulation function. In particular, we increase η_0 so that the implied increase in the marginal returns to experience averaged along the life cycle is approximately 10%, 20% and 40%.¹⁶

An increase in the return to experience by 10% has only small effects on the participation of mothers or any women. When the increase is 20%, the participation rates of both young and older mothers increases by 7% points

¹⁶Note that the marginal increase in the return to experience implied by the human capital function parameter depends on age.




Figure 17: Returns to Experience for Cohort 2 and 3

This compares to increases of 22% points and 17% points between cohort 2 and 3 for the two groups. Even when we increase the return to experience by 40%, we still do not get an increase that is as large as that observed in the data. One explanation of this may be the fact that an increase in the return to experience not only tilts the wage age profile but also increases the amount of resources available to our households. Moreover, a large increase in the return to experience allows women who take time out of employment to raise children to make up for the lost years more quickly.

Turning to the changes in earning profiles, we see that the increase in the return to experience does not explain the changes in wage growth, although it performs better than the previous experiment. For the gender gap, we do not get any substantial changes.

5.3 Wage gender gap

5.3.1 Measuring changes in the relative wages of men and women

There is a substantial amount of evidence that, over the period we are considering, there has been a reduction in the observed wage gender gap. In other words, a reduction in the difference between average observed wages for working men relative to women, controlling for observed skills. The evidence on the wage gender gap is vast and we could not do justice to it here. Instead, we use the PSID to calculate the change in the observed gender wage gap across the

Returns to Experience	Δ Cohort	+10%	+20%	+40%
Participation All Women	0.12	0.02	0.04	0.07
Part. Mothers Child 0-3	0.20	0.04	0.07	0.14
Part. Mothers (Age < 29) Child 0-3	0.22	0.05	0.07	0.15
Part. Mothers (Age \geq 29) Child 0-3	0.17	0.04	0.07	0.13
Part. Women with no Child < 18	0.06	0.00	0.01	0.01
Part. Mothers Child 4-18	0.08	0.02	0.04	0.07
Wage Gender Gap	0.07	0.00	0.01	0.01
Female Wage Growth (Age \leq 35)	0.021	0.00	0.00	-0.001

Table 7: Increasing the Returns to Experience

cohorts. The ratio of observed female to male wages has increased from 0.69 to 0.76 over the two cohorts for women under 40.

This change in the observed gender wage gap is likely to be a result of a reduction in the offered wage gender gap. However, exogenous changes in the offered wage are filtered into changes in observed wages through the selection mechanism that responds to specific incentives. Further, observed wages for women relative to men will be affected by additional human capital accumulation earned by women through any increased participation. In what follows, we experiment with different changes in the offered wage gender gap to show the effects on participation and wage dynamics.

5.3.2 Experiment

There are several ways in which we can introduce a change in the gender wage gap in our model. In particular, we can have an exogenous change in the level of female wages relative to male earnings or we could have an increase in the rate of growth for female wages (but not for male earnings or wages) that over time would induce a reduction in the wage gender gap. In Table 8, we consider the two mechanisms for changing the gender gap in turn.

As with the other tables, in Table 8 we first report the change in the statistics across the two cohorts. We then report the results we get by increasing the wage level by 5%, 10% and 15%. Finally, in the last two columns, we consider introducing exogenous, deterministic wage growth equal to 1% per year and 2% per year for women.

The increase in the female wage level does increase substantially the labour supply of mothers. Indeed an increase in the wage of around 15% generates an increase in participation of both young and older mothers similar to what we observe in the data. The similarity of this result with that obtained with

a change in child care cost is intuitive: what matters for participation is the ratio between wages and child care costs. However, an increase in the level of wages generates an increase in the participation of mothers of older children (aged 4-18) of 11% points compared to the increase of 8% points in the data. Further, the required increase in the gender gap in the offered wage gives rise to an improvement in the gender gap in observed wages that is greater than observed in the data, while the increase in the wage level does not improve our matching of wage growth. A lower increase of 11% in the wage level for females matches exactly the observed changes in the wage gender gap, however this is not enough to explain the observed changes in the participation of mothers, particularly young mothers.

Introducing a deterministic growth component to female wages does lead to better matching of the wage growth statistics. However, the implications for participation do not match the data. In particular, deterministic growth has a larger impact on old mothers, relative to young mothers, whereas in the data the relative change is the other way around.¹⁷

Increasing Wages	Δ Cohort	Δ Gender Gap			Wage Growth	
		+5%	+11%	+15%	+1%	+2%
Δ Participation All Women	0.12	0.04	0.08	0.10	0.07	0.14
Δ Part. Mothers Child 0-3	0.20	0.07	0.15	0.19	0.12	0.23
Δ Part. Mothers (Age < 29) Child 0-3	0.22	0.09	0.19	0.25	0.09	0.19
Δ Part. Mothers (Age \geq 29) Child 0-3	0.17	0.06	0.12	0.16	0.14	0.26
Δ Part. Women with no Child < 18	0.06	0.01	0.02	0.02	0.02	0.04
Δ Part. Mothers Child 4-18	0.08	0.04	0.08	0.11	0.08	0.16
Δ Wage Gender Gap	0.07	0.03	0.07	0.09	0.06	0.13
Δ Female Wage Growth (Age \leq 35)	0.021	-0.002	-0.004	-0.006	0.005	0.009

Table 8: Increasing Female Wages

5.4 Combining explanations

Clearly none of the experiments we have considered so far can fully explain the observed patterns in the data. The hypothesis that appeals to an increase in the return to experience is the least attractive in terms of explaining the changes between cohorts 2 and 3 in the variables we considered. On the other

¹⁷We did consider exogenous growth affecting both male and female earnings. This induced income effects which substantially dampened the incentive for women to participate and we do not report these results further.

hand, both changes in the wage gap and changes in child care costs can explain several aspects of the data.

The main problems we encounter with the explanation that appeals only to changes in child care costs is the inability of such an experiment to explain changes in the observed wage gap and wage growth between the two cohorts. Although, as we discussed above, selection is an important channel through which changes in various variables are reflected in aggregate wages, clearly it is not enough to generate the patterns of wages we observe. On the other hand, as we saw above, plausible changes either in the level or the growth of the wage gender gap cannot explain the labour supply behaviour of mothers. For this reason, our next exercise consists in combining two different explanations: that of an increase in the wage gap and a simultaneous decrease in child care costs. Clearly, as the two effects cumulate, we consider smaller changes in the two variables.

In Table 9 we show the combination of an 11% increase in female wages with a 6% decrease in the child care costs. Including the 6% decrease to child care costs helps to explain observed changes in participation of mothers, whereas it still predicts well the change in the wage gender gap and the participation of mothers of children age 4 to 18. However, the prediction for wage growth worsens. In Table 10 we combine a 1% increase in the exogenous wage growth with a 9% decrease in the child care cost. Again, the combination of two exogenous changes improves the predictions of the simulations in explaining simultaneously the observed wage and labour supply data. In this case, our match of the relative participation of young versus old mothers is less close than when adjusting the wage level, while our match on observed wage growth improves, although we still do not match the size of the change.

With 11% Increase in Female Wage	Δ Cohort	Δ ChildCare Cost	
		0%	-6%
Δ Participation All Women	0.12	0.08	0.10
Δ Part. Mothers Child 0-3	0.20	0.15	0.20
Δ Part. Mothers (Age < 29) Child 0-3	0.22	0.19	0.25
Δ Part. Mothers (Age \geq 29) Child 0-3	0.17	0.12	0.15
Δ Part. Women with no Child < 18	0.06	0.02	0.02
Δ Part. Mothers Child 4-18	0.08	0.08	0.10
Δ Wage Gender Gap	0.07	0.07	0.07
Δ Female Wage Growth (Age \leq 35)	0.021	-0.004	-0.006

Table 9: Gender Gap and Child Care Costs

With 1% Exogenous Wage Growth	Δ Cohort	Δ ChildCare Costs	
		0%	-9%
Δ Participation All Women	0.12	0.07	0.11
Δ Part. Mothers Child 0-3	0.20	0.12	0.20
Δ Part. Mothers (Age < 29) Child 0-3	0.22	0.09	0.20
Δ Part. Mothers (Age \geq 29) Child 0-3	0.17	0.14	0.20
Δ Part. Women with no Child < 18	0.06	0.02	0.02
Δ Part. Mothers Child 4-18	0.08	0.08	0.12
Δ Wage Gender Gap	0.07	0.06	0.06
Δ Female Wage Growth (Age \leq 35)	0.021	0.005	0.002

Table 10: Wage Growth and Child Care Costs

Combining factors in this way enables us to match both participation and wage statistics with more plausible parameter changes. The combination with the wage level change is better for matching the difference between young and old mothers, while the change in the wage growth matches the average behaviour of mothers and goes some way towards matching wage growth.¹⁸ The fact that a combination of changes is necessary to explain differences across the cohorts is highly plausible because, in reality, several of the factors that affect female participation have changed and different explanations are not mutually exclusive.

6 Implications and Conclusions

It is now time to take stock of the simulations performed in the previous section and discuss what can be learned from them in terms of explaining the differences in the behaviour of the Hillary Clinton and Oprah Winfrey cohorts that we discussed in Section 2. We simulate three changes to important determinants of female labour supply. In particular, having matched the behaviour of the Hillary Clinton cohort in our baseline simulations, we show how the life-cycle profile of female labour supply changes when we decrease the cost of child care relative to earnings, when we increase the returns to experience, and when we reduce the wage gender gap (through both a change in the growth and a change in the level of female). We also presented some evidence on the importance of these

¹⁸While it may seem obvious that adding an exogenous growth to wages improves our ability to match the observed increased wage growth, observed wage growth reflects both the exogenous changes to the offer wage distribution and the selection process induced by the model. It is this selection process which can, in principle, substantially dampen the effect of increased wage growth.

factors over the relevant period. While it is hard to measure the magnitude of some of these changes, our simulations serve the purpose of evaluating the relative merits of alternative explanations of the observed changes in female labour supply. Having done that, and in particular, having determined the magnitude of changes in the environment needed to cause the observed shift in labour supply behaviour, we can assess how realistic they are.

First, we observe that a decrease in child care costs of greater than 15% would be needed to explain the increase in participation of mothers of young children from the observed 0.47 rate in the Clinton cohort to the observed 0.67 in the Winfrey cohort. Such a change also explains the differential shifts in the labour supply of younger and older mothers and the change in participation by mothers of older children. We presented evidence showing that changes in childcare costs of a similar magnitude have occurred.

Second, an increase in the return to experience that is similar to that estimated in the data by Olivetti (2006) is not able to increase the participation rate of mothers by the amount we observed in the data, particularly for younger mothers. More substantial increases in the return to experience increase labour supply somewhat but additional increases imply wealth effects that dampen the improved incentives to work.

Third, an increase in the level of wages does increase substantially the labour supply of women. This is consistent, within the model, with the large effects seen for a decrease in the cost of child care: it is the ratio of wages to child labour costs that is relevant. However, the increase in the level of wages that is necessary to match the change in participation implies an increase in the labour supply of mothers of older children that are larger than observed in the data and it implies an improvement in the observed gender gap that is greater than that observed. Changing the wage gender gap through exogenous growth in wages is less successful at matching participation because such exogenous growth induces a greater change in participation for older mothers than for young mothers, which is the reverse of what is observed in the data.

This evidence leads us to conclude that no single change can explain both changes in participation and changes in wages. Instead, we show that a combination of a decrease in child care costs alongside a reduction in the gender wage gap through aggregate wage growth can match both the changes in participation and key aspects of observed wages. This does not mean, however, that our model can predict all aspects of the data. In particular, we cannot explain the *size* of the increase in the rate of growth of wages.

We have not explicitly investigated the effects of other factors that are, in all likelihood, important determinants of female labour supply. Those that immediately come to mind are changes in fertility (both in number of children

and their timing), changes in home production technology, changes in divorce and marriage rates and changes in education. The reason for these omissions is that these factors have not changed much between cohort 2 and cohort 3. However, they are probably very important, instead, in explaining the shift from cohort 1 to cohort 2, as argued recently by Greenwood et al. (2005) for home production technology, and by Albanesi and Olivetti (2006) for changes in home production.

The results of this paper are very suggestive and open further lines of enquiry. Two important issues that we have ignored are the incentives to accumulate human capital through education or training and the possibility of working part time. Both these issues again, can be very important in the early part of the life cycle, where most of the action happens both in the data and in our simulations. It would therefore be worthwhile incorporating these issues in more realistic incarnations of the model.

References

- [1] Albanesi, S. and Olivetti, C. (2006) “Home Production, Market Production, and the Gender Wage Gap: Incentives and Expectations” CEPR Discussion Paper, 4984
- [2] Altug, S. and Miller, R.A. (1998), “The Effect of Work Experience in Female Wages and Labor Supply”, *Review of Economic Studies*, 65(1): 45-85.
- [3] Ahn, N. and Mira, P. (2002) “A Note on the Changing Relationship Between Fertility and Female Employment Rates in Developed Countries”. *Journal of Population Economics*.
- [4] Anderson, P. and P. Levine (2000): “Child Care and Mother’s Employment Decisions” in *Finding Jobs: Work and Welfare Reform*, David Card and Rebecca Blank eds., New York: Russell Sage
- [5] Attanasio, O.P. and G. Weber (1995): “Is Consumption Growth Consistent with Intertemporal Optimization?”, *Journal of Political Economy*, 103 (6): 1121-57.
- [6] Attanasio, O.P., Low, H.W. and V. Sanchez-Marcos (2005) “Female labour supply as insurance against idiosyncratic risk” *Journal of the European Economic Association Papers and Proceedings*, 3 (2-3): 755- 764
- [7] Blau, D. (1992). “The Child Care Labor Market”, *Journal of Human Resources*, 9-39.
- [8] Blau, D. (2001) *The Child Care Problem: An Economic Analysis*. Russell Sage Foundation
- [9] Blau, D. (2003) “Child care subsidy programs” in: *Means-Tested Transfer Programs*, ed: R. Moffitt, NBER
- [10] Browning, M. and Meghir, C. (1991) “The effects of male and female labour supply on commodity demands”, *Econometrica*, 59 (4): 925-951
- [11] Carroll C., and A. Samwick (1997), “The Nature of Precautionary Wealth”, *Journal of Monetary Economics*, 40 (1): 41-71.
- [12] Caucutt, Guner and Knowles (2002), “Why do Women Wait? Matching, Wage Inequality and The Incentives for Fertility Delay”, *Review of Economic Dynamics*, 5: 815-855.
- [13] Coleman, M. and Pencavel, J. (1993) “Trends in market behavior of women since 1940” *Industrial and Labour Relations Review* 46: 653-76

- [14] Conesa, J.C. (1999), “Educational Attainment and Timing of Fertility Decisions”. Unpublished manuscript.
- [15] Department of Health and Human Services (1989) “Vital and Health Statistics: Trends and variations in first births to older women, 1970-1986” Series 21, No 47
- [16] Eckstein, Z. and K. Wolpin (1989) “Dynamic labor force participation of married women and endogenous work experience” *Review of Economic Studies* 56: 375-390
- [17] Ferrero, D. and Iza, A. (2002), “Skill premium effects on fertility and female labour supply”, *Journal of Population Economics*, 17 (1): 1-16.
- [18] Francesconi, M. (2002) “A joint dynamic model of fertility and work of married women”, *Journal of Labor Economics*, 20 (2): 336-380
- [19] French, E. (2005) “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour”, *Review of Economic Studies* 72 (2): 395-427
- Greenwood, J. and Seshadri, A. (2002), “Technological Progress and Economic Transformation” *Economie d'avant garde*, Research Report No. 3, University of Rochester.
- [20] Greenwood, J. Seshadri, A. and M. Yorukoglu (2005) “Engines of liberation”, *Review of Economic Studies*, 72: 109-133
- [21] Hotz, V.J. and R.A. Miller (1988) “An empirical analysis of life-cycle fertility and female labor supply”, *Econometrica* 56 (1): 91-118
- [22] Household Economic Studies (2001) “Maternity Leave and Employment Patterns: 1961-1995” P70-79. US Census Bureau.
- [23] Jones, L. , Manuelli, R. and McGrattan, E. (2003) “Why are married women working so much?”. Federal Reserve Bank of Minneapolis.
- [24] Keane, M.P. and K. Wolpin (1997) “The career decisions of young men”, *Journal of Political Economy* 105(3): 473-522
- [25] van der Klaauw, W. (1996) “Female labour supply and marital status decisions: a life-cycle model”, *Review of Economic Studies* 63(2): 199-235
- [26] van der Klaauw, W. and Wolpin, K. (2005) “Social Security and the Retirement and Savings Behavior of Low Income Households”, PIER Working Paper 05-20

- [27] Mcgrattan, E. and Rogerson, R. (2004) “Changes in hours worked, 1950–2000”, *Quarterly Review*, Federal Reserve Bank of Minneapolis: 14-33
- [28] Mincer, J. and S. Polachek (1974) “Family investment in human capital: earnings of women”, *Journal of Political Economy* 82(2): 76-108
- [29] Mincer, J. and H. Olfek (1982): “Interrupted Work Careers: Depreciation and Restoration of Human Capital”, *Journal of Human Resources*, 17, 3-24
- [30] Olivetti, C. (2006), “Changes in Women’s Hours of Market Work: The Effect of Changing Returns to Experience”. *Review of Economic Dynamics*, 9 (4): 557-587.
- [31] Pencavel, J. (1998) “The Market Work Behavior and Wages of Women: 1975-94”, *Journal of Human Resources* 33 (4):771-804.
- [32] Rendon, S. (2006) “Job Search And Asset Accumulation Under Borrowing Constraints” *International Economic Review* 47 (1): 233-263
- [33] Ribar, D. (1992): “Child Care and Labour Supply of Married Women: Reduced Form Evidence”, *Journal of Human Resources*, 27: 134-65.
- [34] Sánchez-Marcos, V. (2002) “Women’s Employment and Fertility in Spain over the last twenty years”. Unpublished manuscript.
- [35] Schulman, K. (2000), “The High Cost of Child Care Puts Quality Care Out of Reach for Many Families”, Published by Childrens Defense Fund (CDF).
- [36] Smith, K., B. Downs and M. O’Connell (2001) “Maternity leave and employment patterns: 1961-1995” Current Population Reports, P70-79, U.S. Census Bureau, Washington, D.C.

7 Appendix

7.1 Solution Method.

Households have a finite horizon and so the model is solved numerically by backward recursion from the terminal period. At each age we solve the value function and optimal policy rule, given the current state variables and the solution to the value function in the next period. This approach is standard. The complication in our model arises from the combination of a discrete choice (to participate or not) and a continuous choice (over saving). This combination means that the value function will not necessarily be concave. We briefly describe in this appendix how we deal with this potential non-concavity.

In addition to age, there are four state variables in this problem: the asset stock, the permanent component of earnings of the husband, ν_t^m , the permanent component of earnings of the wife, ν_t , and the experience level of the wife. We discretise both earnings variables and the experience level, leaving the asset stock as the only continuous state variable. Since both permanent components of earnings are non-stationary, we are able to approximate this by a stationary, discrete process only because of the finite horizon of the process. We select the nodes to match the paths of the mean shock and the unconditional variance over the life-cycle. In particular, the unconditional variance of the permanent component must increase linearly with age, with the slope given by the conditional variance of the permanent shock.

Value functions are increasing in assets A_t but they are not necessarily concave, even if we condition on labour market status in t . The non-concavity arises because of changes in labor market status in future periods: the slope of the value function is given by the marginal utility of consumption, but this is not monotonic in the asset stock because consumption can decline as assets increase and expected labour market status in future periods changes. By contrast, in Danforth (1979) employment is an absorbing state and so the conditional value function will be concave. Under certainty, the number of kinks in the conditional value function is given by the number of periods of life remaining. If there is enough uncertainty, then changes in work status in the future will be smoothed out leaving the expected value function concave: whether or not an individual will work in $t + 1$ at a given A_t depends on the realisation of shocks in $t + 1$. Using uncertainty to avoid non-concavities is analogous to the use of lotteries elsewhere in the literature.

The choice of participation status in t is determined by the maximum of the conditional value functions in t . In our solution, we impose and check restrictions on this participation choice. In particular, we use the restriction that the participation decision switches only once as assets increase, conditional on permanent earnings and experience. When this restriction holds, it allows us to interpolate behaviour across the asset grid without losing our ability to determine participation status. We therefore define a reservation asset stock, R_t to separate the value function and the choice of consumption made when participating from the value function and choice of consumption made when not participating.

Solving for the reservation asset stock serves two purposes: one, it makes it easier in the solution to allow for the fixed cost in the budget constraint (rather than having an unconditional policy function with a discontinuity); two, it provides an additional check on our numerical solution: the reservation asset stock should be increasing in the wage rate. A sufficient condition for this

to be unique is that the conditional value functions be concave. This is not true in general, as discussed above, but uniqueness can be achieved by having enough uncertainty to make the conditional expected value function concave. Even when the conditional value functions are not concave, however, we can have a unique reservation asset stock, particularly if individuals are impatient enough: impatience means that individuals prefer periods of non-participation to be earlier in the life-cycle and so avoids indifference about the timing of leisure which can generate non-uniqueness.¹⁹

In solving the maximisation problem at a given point in the state space, we use a simple golden search method. We solve the model and do the calibration assuming this process is appropriate and assuming there is a unique reservation asset stock for each point in the state space. We then check that the results in our baseline case are unaffected when we use a global optimising routine, simulated annealing, and we do not assume a unique reservation asset stock. It is worth stressing that there are parameter values for which the techniques we used do not work. In particular, the assumption of a unique reservation asset stock fails as the variance of shocks gets sufficiently low and if households have discount rates very close to the interest rate.

There are no non-concavities due to borrowing constraints in our model because the only borrowing constraint is generated by the no-bankruptcy condition which is in effect enforced by having infinite marginal utility of consumption at zero consumption.

¹⁹In principle, we could test whether this unique reservation asset stock property held empirically. However, such a test would be difficult given the quality of asset data. Further, a rejection in the data may well be generated by unobserved heterogeneity which would not be a problem in our solution where the solution is carried out separately for each “type” of person.