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

Orazio Attanasio,*Hamish Low†and Virginia Sanchez-Marcos‡

July 2003

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 among these three cohorts. We then use 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 cost (relative to earnings), return to experience and the rate of depreciation of human capital when out of the labour market. 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 shifts in the cost of children relative to life time earnings are the most likely explanation.

*University College London, IFS and NBER
†University of Cambridge and IFS
‡University of Cantabria

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 compares cohorts of women born in the 1930s (such as Elizabeth Dole), 1940s (Hillary Clinton) and 1950s (Oprah Winfrey), two main features emerge when considering labour supply in its various dimensions. First, comparing the Elizabeth Dole cohort to the Hillary Clinton one, we can see a substantial shift of the age profile of labour supply. However, the shape of the profile does not change much. In particular, in both profiles we observe a low participation 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 propose a life cycle model of saving and labour supply that could account for these dramatic changes. We explore whether 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 unobservable factors needed to explain the observed patterns.

The main change in labour supply behavior in the data is on the extensive margin. We consider a number of possible determinants of these changes in participation. First, wages may have increased relative to the fixed cost of participation. For example, the costs of child-care may have fallen. This would lead to greater participation at all ages and especially among mothers of infants. Second, on-the-job learning or the return to experience may have increased. As argued by Olivetti (2001), this increases the opportunity cost of reduced labor supply. Third, the depreciation of skills that occurs if an individual is not participating may have increased. Finally, we look at other possible explanations, such as a delay in the arrival of the first child and an increase in uncertainty over husband's income. Our structural model of life-cycle behavior attempts to evaluate these alternative explanations.

Obviously, wages are likely to be an important determinant of female labour supply. However, by looking at the dynamics of wages alone, it is difficult to disentangle the return to experience, the depreciation rate of human capital and the extent of 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. Furthermore, a simple analysis that relates

wages to labour supply, neglects general equilibrium effects that also imply an effect running from labour supply to wages. The main purpose of this paper is to build a realistic life-cycle environment in which we can explicitly model the participation choice. We can then calibrate the model to fit the behaviour of a given cohort and experiment with changes in the basic determinants of labour supply to determine which are more likely to yield the profiles of other cohorts.

In our life cycle model households face uncertainty about the wages of the husband and the wife; maternity is exogenously given and children impose some monetary fixed cost when mothers decide to work. Decisions are taken at an annual frequency. The model takes into account returns to experience as a result of participation and depreciation of human capital when labor market interruptions are made. Households are able to save and borrow and women choose whether or not to work. This makes our model different from Keane and Wolpin (1989) and van Der Klaauw (1996), who estimate structural models of females' employment decision in the first case and females' employment and marital status decisions in the second case imposing that consumption coincides with income. Without the saving choice, the only way to intertemporally substitute consumption 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 means of intertemporal substitution and so ignoring savings overstates the importance of the labor supply choice in life-cycle smoothing.

We calibrate our model by matching observed participation profiles to simulated participation and observed wage profiles to the simulated wages of those who choose to work. Wage profiles in both the data and in the simulations are subject to selection; that is we only observe the wages of the women who choose to participate. Our selection model enables us to identify the depreciation effect separately from the return to experience. We use observed profiles from the cohort born at the start of the '40s for our calibration. We then explore the role of different factors in shaping changes of the life-cycle wage profile and participation profile.

Pencavel (1998) and Coleman and Pencavel (1993) report similar paths for participation to the paths we report. The facts on employment are not in dispute. More controversial is understanding the data on wage profiles, on depreciation of human capital and on the underlying question of why participation has changed. Mincer and Pollachek (1974) and Mincer and Olfek (1982) discuss the extent of human capital depreciation under different assumptions on the permanence of depreciation. We report some statistics on depreciation but

without a structural model of participation it is hard to identify the depreciation rate. Olivetti (2000) suggests that changes in wage profiles across cohorts reflect a change in the return to experience. The evidence we present is somewhat weaker: first, the cohort effect which leads to an increase in the return to experience can plausibly be interpreted as a year effect with wages in the 1980s growing faster than in previous periods. Second, 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 (2001) uses a four period model and the estimates of the returns to experience in Olivetti (2000) to show the effect that increase in the returns to experience has on hours worked by women. Greenwood and Seshadri (2002) measure the impact of technological progress on the increase in women's participation. The contribution of the current paper is primarily to use a realistic life-cycle model of saving and participation to compare alternative explanations.

The paper is organized as follows. In Section 2 we describe the data that motives the paper, in particular women employment behavior. In particular, we focus on 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 with the literature. The model is a partial equilibrium model and in section 4 we report simulations for individuals with different parameter specifications. We describe first the parameters we use for uncertainty, returns to experience, depreciation, child costs, maternity age and preferences. Second, we describe our baseline simulations, showing participation statistics and life-cycle profiles. In section 5 we carry out comparative statistics for several parameters values. In Section 6, we discuss the implications of our simulations for the changes in participation described in Section 2 and conclude.

2 Data

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 17 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 women employment in the US over the last decades is the change of married women's behavior. The main issue is whether the trend towards marrying later might affect or bias our results. 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 Facts to explain: Employment

We start our analysis by looking at the life cycle profiles of hours worked. In Figure **??**, 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, the large increase 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, by age 37 is minimal. The life cycle profile starts much higher for the third cohort.

In Figure **??**, 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.

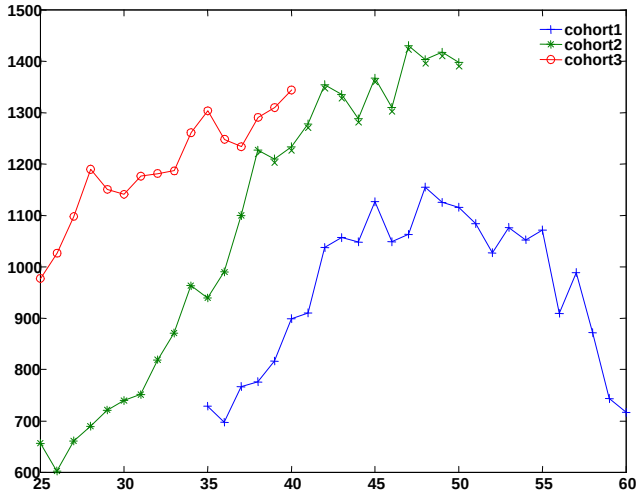

Figure 1: Average Hours Worked

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 former 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 6 and 7, 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 65% of mothers with a child less than 3 are working, while the same figure for cohort 2 is only 45%. This evidence is consistent with some facts reported

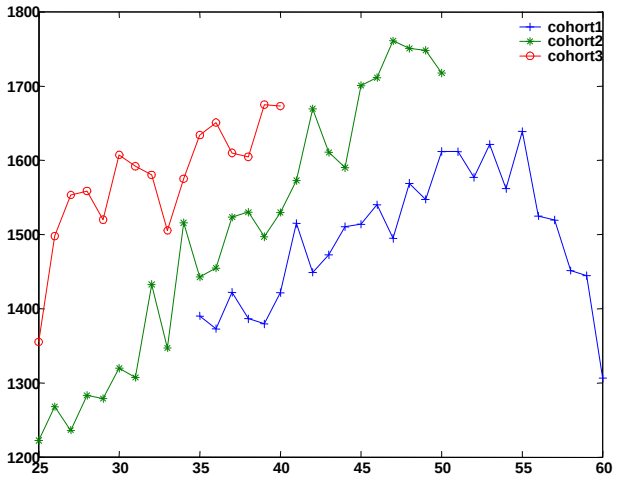

Figure 2: Average Hours Worked, Employed Women Only

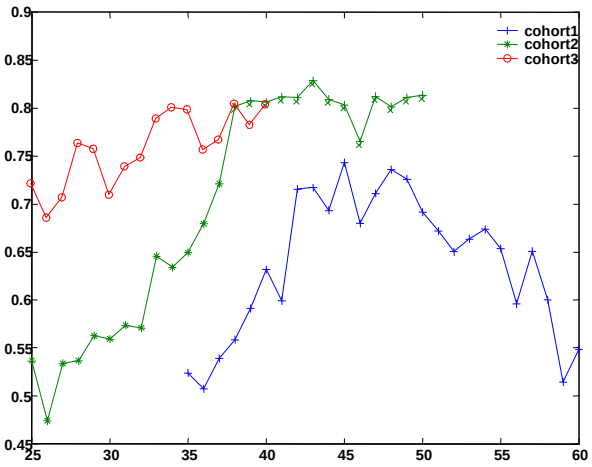

Figure 3: Employment Rate

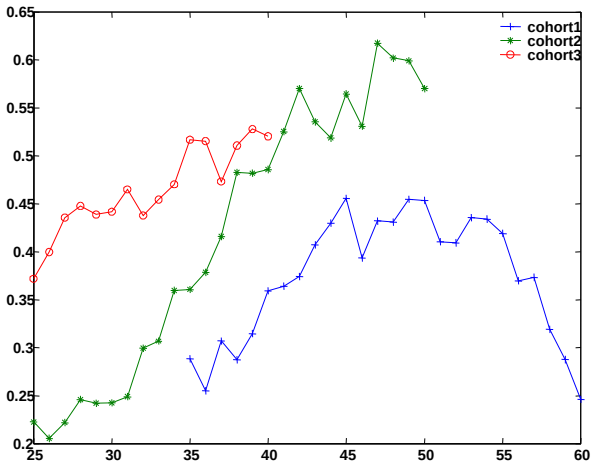

Figure 4: Full-Time Employment Rate

by the US Census Bureau, reproduced in figure 5. They consider women that 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 5 shows that the first two cohorts were unlikely to have 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. We therefore look at employment rates by education. We divide married women into two groups, those with more than high school education and those who are at most high school graduates. Figures 8 and 9 show employment and full-time employment rates for women with low educa-

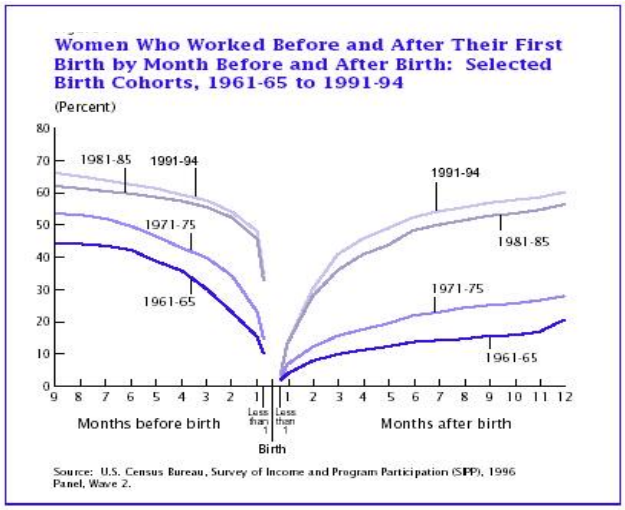

Figure 5: Speed of Return After Childbirth

tion. Figures 10 and 11 show the rates for women with high education. We observe that the increase of women's employment and full time employment rates is important for both educational groups and in both cases the most significant change is the flattening of the age - participation profile between ages 25 and 35.

Figures 6 and 7 show 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 and to check how these differences present themselves within education groups, in Table 1 we compute, for cohorts 2 and 3, the percentage of women that exit from the labour market for each age between 26 and 35. The table shows that a much smaller fraction of cohort 3 women exit the labour market between ages 26 and 35. These differences are relevant both for the high and the low education group.

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 the first two panels of Table 2 we report the mean and median duration of time out of the labour market for women who return, both in the whole sample and in the two education groups. In the right panel of the table, we report median duration for all women who exit, including those who do not return. For women whose duration is censored, in this computation it is assumed that they never

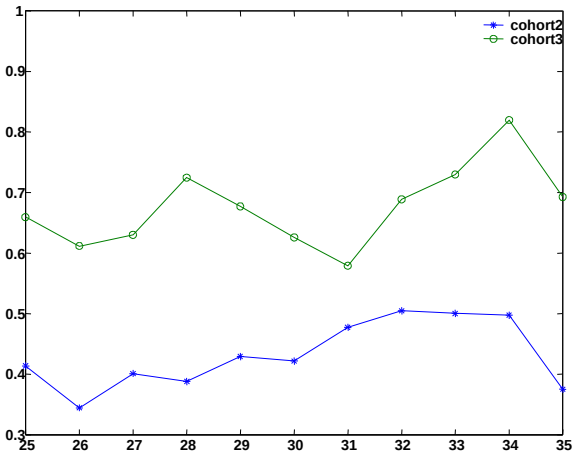

Figure 6: Employment Rates of Mothers of Children less than 3 yrs

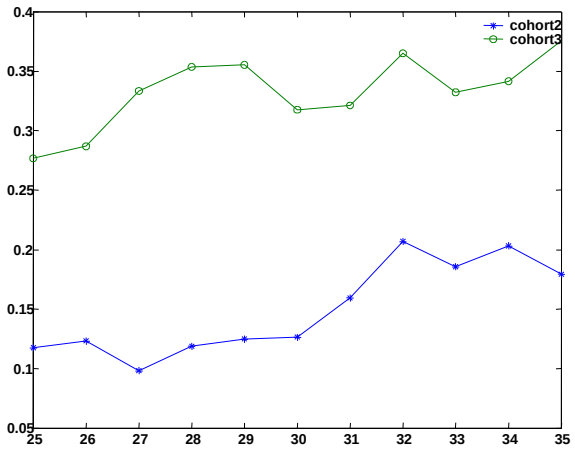

Figure 7: Full-Time Employment Rates, Mothers of children less than 3 yrs

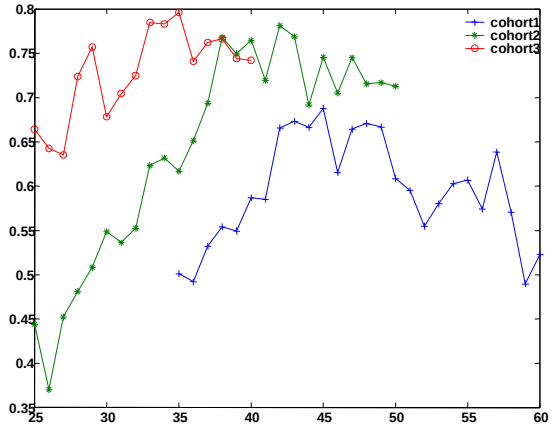

Figure 8: Employment Rate, Low Education

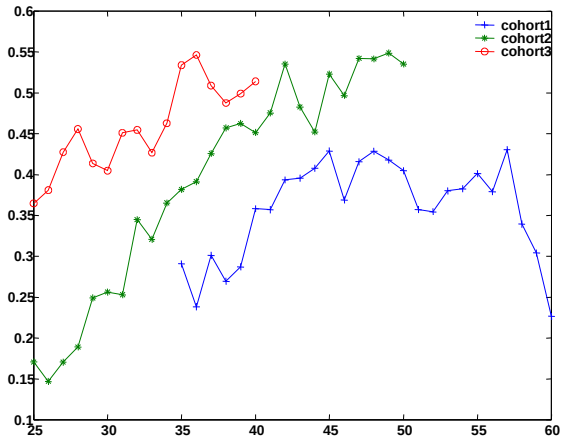

Figure 9: Full-Time Employment Rate, Low Education

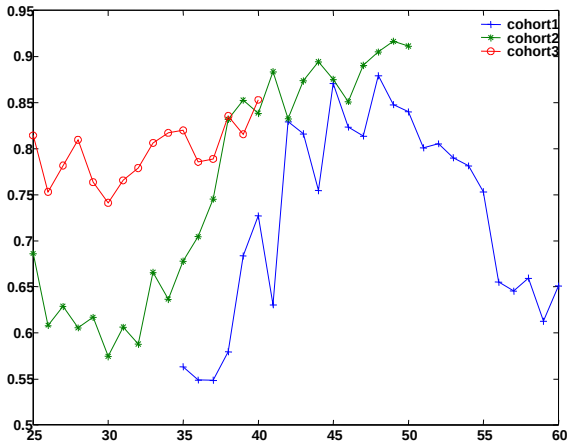

Figure 10: Employment Rate, High Education

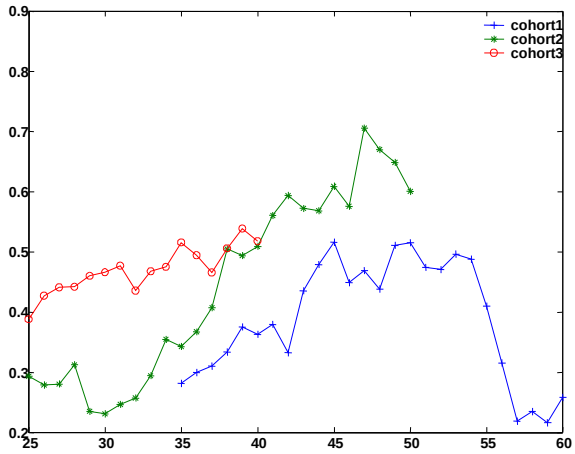

Figure 11: Full-Time Employment Rate, High Education

Table 1: Fraction of Women Exiting at Each Age

	Cohort 2			*Cohort 3*		
Age at exit	*Total*	*Low Edu*	*High Edu*	*Total*	*Low Edu*	*High Edu*
26	27	32	22	16	16	15
27	15	18	14	12	12	9
28	15	15	14	7	7	7
29	16	15	16	10	10	12
30	20	18	20	14	14	10
31	14	17	11	10	10	8
32	15	15	16	10	10	8
33	7	8	6	7	7	9
34	12	13	12	7	7	7
35	10	10	10	8	8	7

Fraction of women working at age t who exit at age t.

re-enter and the duration is set at 65 minus the date of exit. As cohort 1 is only observed from age 35 onwards, the interesting comparison is between cohorts 2 and 3. Average duration is considerably shorter for women in the younger cohort. Further, median duration is also shorter for cohort 3 women.

Table 2: Duration of Non-Participation

	Women who reneter						*All Women*		
	Mean			Median			Median		
	Total	*Low*	*High*	*Total*	*Low*	*High*	*Total*	*Low*	*High*
Cohort 1	2.62	3.05	1.76	1	2	1	4	4	3
Cohort 2	3.42	2.84	3.74	2	2	2	4	3	3
Cohort 3	2.14	2.04	2.17	1	1	1	3	5	3

When duration is censored, age of reentry is set at 65. Censoring is larger for the younger cohort and so using mean values without conditioning on reentry is uninformative.

2.2 Wages

The price of human capital is obviously 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 in the next section, 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 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 is 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 12 plots for each of the three cohorts, the median female hourly wage against age. Figures 13 and 14 re-do the exercise for the two education groups we have considered above. Cohort 3 appears to face a much steeper wage profile in the early part of the life-cycle than cohort 2. This may be interpreted as an increase in the return to experience. However, it is difficult to separate out the cohort and year effects. In particular, all three profiles of wages exhibit sizeable increases in the second half of our sample period. During the first half of the sample, which covers the 1970s and for which we only observe cohort 1 and 2, real wages were basically flat. This is particularly apparent for the high educated women. For the high educated women belonging to cohort 1 median real wages were actually declining during the first part of the sample. By the time cohort 3 comes in, and starts enjoying relatively fast growth of real wages, the wages of cohort 1 and 2 also start increasing. Of course the three cohorts experience these changes in real wages at different ages. To stress how difficult is to interpret these patterns in Figure 15 we plot median wages for the three cohorts against time, rather than age. This suggests we should attribute all or

part of the increase in wages to aggregate factors that move the wages of women belonging to all cohorts.

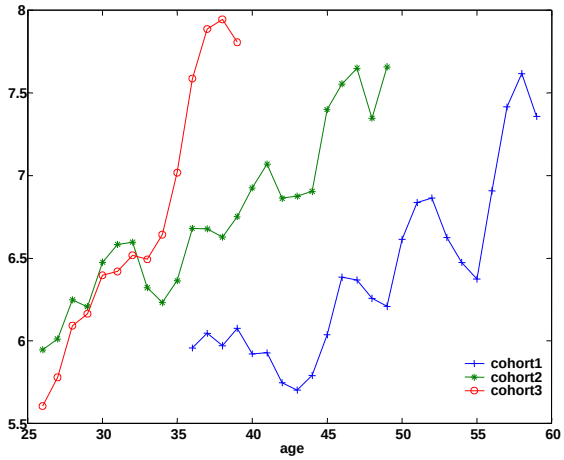

Figure 12: Median Hourly Wage by Age

In addition to the problems relating to the impossibility of disentangling age, time and cohort effects, the life cycle profiles in Figures 12, 13 and 14 are only partially informative about the return to experience because at each age we average the wages of women with different levels of labour market 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 that have been observed since age 25 and have been not had more than one year off work. These two profiles, plotted in Figure **??**, 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 Figures 12, 13 and 14 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 that we know 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

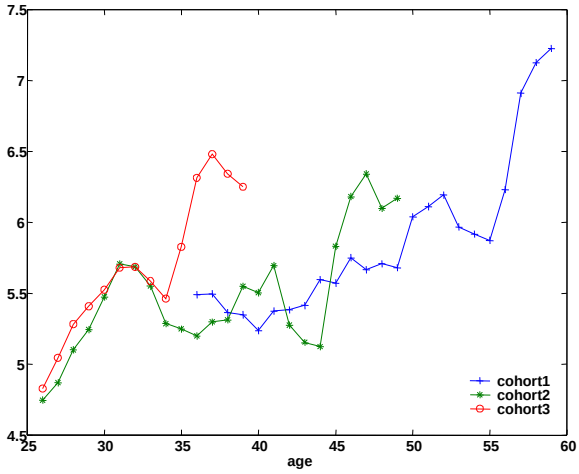

Figure 13: Median Hourly Wage, Low Education

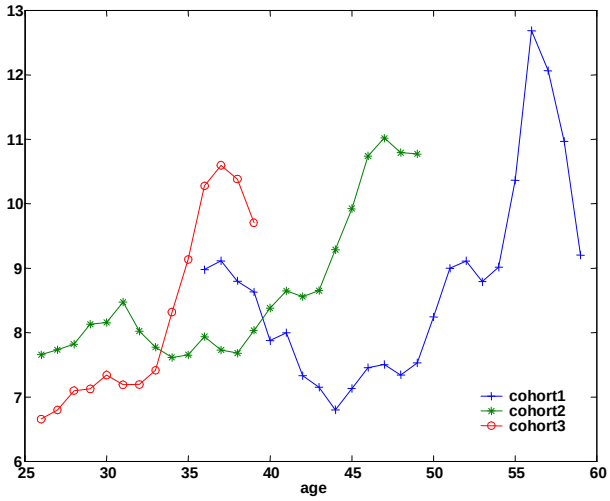

Figure 14: Median Hourly Wage, High Education

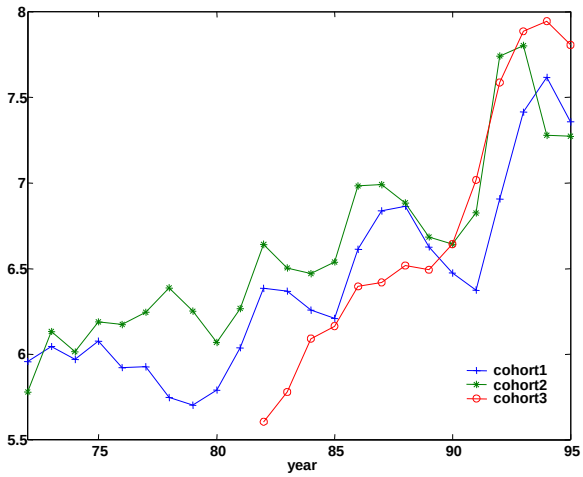

Figure 15: Median Hourly Wage by Year

to experience between cohort 2 and 3. However, increases in the return to experience for cohort 3 could be hidden by the shorter duration of exits from the labour market for this cohort (as well as by changes in the selection process).

An important determinant of the participation decision is possibly the depreciation of human capital when out of the labour force. A clean estimate of a depreciation model is remarkably hard because of the selection problems involved: one only observes the wage of women who decide to go out and then to come back. There will be women who might have not gone out because of the depreciation and, more importantly, women who have not come back because of the extent to which their human capital had depreciated. Without a structural model (and the assumptions connected to it) it is not possible to solve these selection issues. Nonetheless, we look at what happens to the wages of women who have exited the labour market and decided to reenter after an interruption. Our exercise is similar to the one performed by Mincer and Polachek (1974) and Mincer and Olfek (1982). In particular, we follow women that have employment interruptions and who are observed before and after the interruption. In Table 3, we report the average level of wages before and after the interruption for these two cohorts and for different education levels, as well as mean and median depreciation.

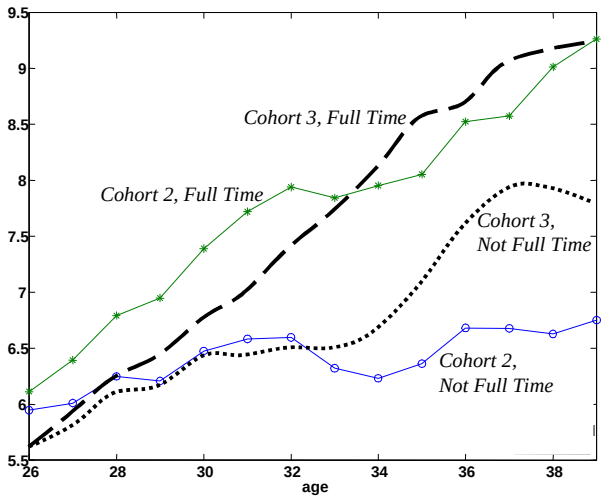

Figure 16: Returns to Experience for Cohort 2 and 3

Table 3: Wage Depreciation

	Cohort 2			*Cohort 3*		
	Total	*Low Edu*	*High Edu*	*Total*	*Low Edu*	*High Edu*
Wage before	6.13	4.42	7.32	6.21	4.45	7.82
Wage after	5.45	4.40	6.26	5.86	5.24	6.67
Mean dep	1.59	1.48	1.69	1.62	1.54	1.70
Median dep	1.24	0.93	1.24	1.01	0.92	1.22
No. obs	167	85	75	153	63	68

Depreciation is the ratio of the wage before exit to the wage after return. The sample is restricted to the women who were married both before and after the interruption. For the education split, we consider only women that had the same education level before and after the interruption.

Mean and median depreciation do not change substantially between the two cohorts. However, the fact that women in cohort 3 stay out for a shorter amount of time does imply a faster depreciation rate. Median depreciation is much smaller than mean depreciation. In fact, for low educated women the median change in wages after an interruption is positive rather than negative.

We are of course aware of the limitations of these measures. First, there is a selection bias that can give misleading returns to experience and penalty of employment interruptions. Second, concerning the measure of the penalty for employment interruptions, we are comparing cohorts that are observed at different ages and it could be that the cost of employment interruptions in terms of future wages is higher at the beginning of the life cycle. However, the average age at labor market interruptions is 30 years for both cohorts.

2.3 Fertility and child care arrangements

Fertility behaviour is obviously important for labour supply decisions, either as a determinant or as a joint decision variable. The cost of children can be an important indirect determinant of labour supply decisions. In Figure **??**, we plot the average number of children aged less than 17 present in the household against their mother's age. From the figure the massive decline in fertility is apparent: at age 35 a woman of cohort 1 had on average 3 children living with her, while at the same age a woman belonging to cohort 1, had less than 2.

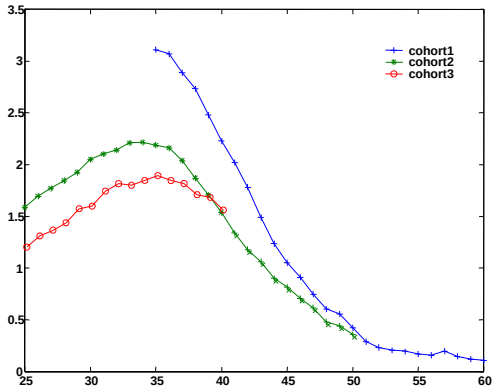

Figure 17: Average Number of Children Younger than 17, per Household

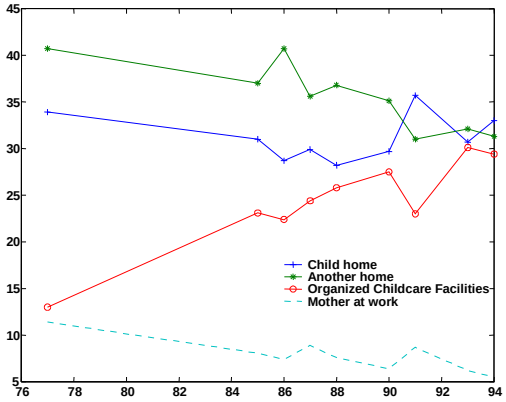

Figure 18: Child Care Arrangements

As in other developed countries, the reduction in the number of children happened at the same time as a substantial 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 that have their first child in any given age interval. The numbers reported in Table 4 are revealing: 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).

Table 4: Distribution of Maternity Age

Maternity Age	*Percentage of Women*		
	Cohort 1	*Cohort 2*	*Cohort 3*
$18-22$	31	30	26
$23-26$	40	33	33
$27\leq$	29	37	41

Child Care costs can be substantial. 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. Concerning these costs, Ferrero and Iza (2002) argue that the skill-biased technological change over the last decades, that 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. Ahn and Mira (2002) argue that there was a decline in the price of child care relative to female wages based on Blau (1992).[1] The availability of organized child care facilities is also a key issue in explaining women's employment behavior, especially for mothers of young children. We have shown that the increase of mothers' employment might be behind a large part of the increase of women employment over time. It is therefore worthwhile looking at the availability and cost of child care arrangements. Figure 18 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.[2] Historically child care subsidies have not been important. The Child Care and Development Fund is the major source of Federal child care assistance for low and moderate income families and yet in 1998 only 15% of eligible children actually received help through the program.

The introduction of maternity leaves in the US is very recent. It took place in 1993 when the Family and Medical Leave Act was approved. Even then, the approved period of maternity leave was 12 weeks, the lowest of the OECD and does not mandate paid benefits during the period. However, the existence of maternity leave can help women to reconcile employment and maternity. The availability of maternity periods allows women time out of the labour force around childbirth without compelled them to quit their jobs. Indeed, as shown in figure 5, maternity leave is associated with women returning to work after child birth more quickly.[3]

[1] Blau (1992) reports that real wages of child care workers, as well as other workers, were flat from 1976 to 1986.

[2] 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 where $2400 for one child and $4800 for two or more. In 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.

[3] A further issue that could affect married women's employment decision is the risk of

3 Model

In this section we describe the model we use to explain the changes in female labour supply behaviour. 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 husbands always work and receive earnings that are determined by a stochastic process that we introduce below. All households have two adults and receive one child at an exogenous and deterministic time. We consider heterogeneity in the age of arrival, and we consider the impact of exogenous changes in the age of arrival. 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 variable, 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,

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

where β is the discount factor and E_t the expectations operator conditional on information available in period t. We use a utility function of the form

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

As we typically use values of $\gamma > 1$, we constrain $\psi_1 > 0$ so that participation reduces the utility of consumption. Consumption is equivalised by the factor

divorce. When married women face uncertainty on 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 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.

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

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

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 child. Female earnings are given by y_t, 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. As we discuss below, this difference turns out to be substantial.

We model the fixed cost associated with children following Hotz and Miller (1988), who specify the functional form for the time cost of children as

$$G\left(a_t\right) = \theta\phi^{a_t - 1} \tag{3}$$

and estimate parameters θ and ϕ to match the time cost associated with child care for children of different ages. The price of this time cost is then given by p, giving an expression for the fixed cost of work

$$F\left(a_t\right) = pG\left(a_t\right), \tag{4}$$

Female earnings, y_t, are subject to transitory shocks, ε_t, and permanent shocks, v_t, and are given by

$$\ln y_t = \ln y_0 + h_t - \delta_T I\left(P_{t-1} = 0\right) + v_t + \varepsilon_t \quad \text{where } \varepsilon_t \sim N\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \tag{5}$$

$$v_t = v_{t-1} + \xi_t \quad \text{where } \xi_t \sim N\left(-\frac{\sigma_\xi^2}{2}, \sigma_\xi^2\right) \tag{6}$$

where h_t is the level of human capital at the start of the period. If individuals have not participated in the previous period, they have to pay a premium on

[4] According 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 one child which arrives at a predetermined age and leaves at age 18.

their wage for one period for reentry, given by δ_T. We think of δ_T as the temporary depreciation in skills associated with a period out of the labour force, as estimated by van der Klaauw (1996).

Human capital evolves with employment decisions in the following way

$$h_t = h_{t-1} + (\eta_0 + \eta_1 t) I(P_{t-1} = 1) - \delta_P h_{t-1} I(P_{t-1} = 0)$$

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

We think of δ_P as the permanent depreciation in human capital associated with non-participation, as discussed and estimated by Mincer and Polachek (1974) and Mincer and Olfek (1982). Notice that the two forms of depreciation of human capital (temporary and permanent) have very different effects on participation choices. The return to experience that individuals receive from participating depends on their age, as in Olivetti (2000), with the increase in human capital decreasing with age if $\eta_1 < 0$. Notice that 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).

The process of human capital accumulation is both crucial to our model and differentiates it from other models in the literature, such as those in Eckstein and Wolpin (1989), van der Klauw (1996) and Olivetti (2001). We discuss its implications and compare it to the alternative available models in the calibration section where we specify the parameters that we use in the simulations.

Since men always work, male earnings are given by

$$\ln y_t^m = \ln y_0^m + h_t^m + v_t^m + \varepsilon_t^m \quad \text{where } \varepsilon_t^m \sim N\left(-\frac{\sigma_{\varepsilon,m}^2}{2}, \sigma_{\varepsilon,m}^2\right) \tag{7}$$

$$v_t^m = v_{t-1}^m + \xi_t^m \quad \text{where } \xi_t^m \sim N\left(-\frac{\sigma_{\xi,m}^2}{2}, \sigma_{\xi,m}^2\right) \tag{8}$$

$$h_t^m = h_{t-1}^m + (\eta_0 + \eta_1 t) \tag{9}$$

We assume that the return to experience for men is the same as the return to experience for women in our baseline. However, when considering changes to returns to experience, we hold the return for men constant.

In each period, if the women chooses to participate, the value function is

given by

$$V_t^1\left(A_t, v_t, P_{t-1}, h_t\right) =$$
$$\max_{c_t}\left\{u\left(c_t, P_t=1\right)+\ \beta E_t\left[\max\left\{\begin{array}{l} V_{t+1}^0\left(A_{t+1}, v_{t+1}, P_t=1, h_{t+1}\right) \\ V_{t+1}^1\left(A_{t+1}, v_{t+1}, P_t=1, h_{t+1}\right)\end{array}\right\}\right]\right\} \tag{10}$$

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

$$V_t^0\left(A_t, \nu_t, P_{t-1}, h_t\right) =$$
$$\max_{c_t}\left\{u\left(c_t, P_t=0\right)+\ \beta E_t\left[\max\left\{\begin{array}{l} V_{t+1}^0\left(A_{t+1}, \nu_{t+1}, P_t=0, h_{t+1}\right) \\ V_{t+1}^1\left(A_{t+1}, \nu_{t+1}, P_t=0, h_{t+1}\right)\end{array}\right\}\right]\right\} \tag{11}$$

The decision of whether or not to participate in period t is determined by comparing $V_t^0\left(A_t, \nu_t, P_{t-1}, h_t\right)$ and $V_t^1\left(A_t, \nu_t, P_{t-1}, h_t\right)$. The participation choice and the consumption choice in t determines the endogenous state variables (assets, human capital and the reentry decision) 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 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. This has a number of implications: first, the utility cost of non-participation is lower in our model because consumption can be smoothed over time through saving. The presence of husband earnings means consumption is not as variable as the wife's wage income, but, without savings, consumption will still be highly sensitive to the participation choice because of the budget constraint. Since we observe periods of non-participation, a model without saving that tries to match the data on non-participation requires the estimated direct utility benefit of non-participation to be large, the estimated negative impact on future wages of non-participation to be small and estimated child-care costs to be large. Without this, individuals will not be willing to accept the consumption loss in a given period and the variability in consumption across periods implied by non-participation. On the other hand, if individuals are able to save or borrow, the income loss in a particular period associated with non-participation in that period can be spread over the life-cycle.

In effect, if participation does not enter the utility function directly, then the choice of participation affects life-time income, but has only a second-order effect on consumption in any one period. Further, the participation decision will be independent of husband earnings because it is a simple earnings maximisation choice. If there were no saving, then even if participation does not enter the utility function directly, the participation decision will depend on husband earnings. Husband earnings provide a cushion to consumption so choices about participation do not lead to such serious effects for consumption as would arise in the absence of husband earnings.

The effect of ignoring saving and borrowing on estimates of child-care costs goes two ways: on the one hand, if individuals are able to borrow, then they may be willing to participate despite their current wage not covering the child-care costs, in order to reap the benefit of the return to experience in the future. Since they can borrow against this future return, they choose to participate. This implies models without saving will underestimate child-care costs. In a model with certainty, this effect seems to dominate. On the other hand, in the data we observe individuals not participating despite the loss of per period consumption implied by a model with no saving, assuming that the fixed cost is less than potential within-period earnings. In order to generate this observed non-participation we would need estimated fixed costs of work that are higher than in a model when individuals can save and borrow.

4 Baseline Parameters and Simulations

This section is organized as follows. First we describe the parameters we use in our baseline model. We then show the life-cycle profiles 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 the baseline parameters for female participation.

4.1 Baseline parameters

In Table 5, we report the parameters we use in our baseline simulations. The first column contains those parameters that have been measured directly in the data or come from other studies, while the second column contains those that have been calibrated to statistics in the data.

The parameters used to match the data are the utility cost of working (given

Table 5: Baseline Parameters

Exogenous Parameters		*Calibrated Parameters*	
σ^2_ξ	0.031	δ_P	0.02
$\sigma^2_{\xi,m}$	0.031	p	18.0
γ	1.5	ψ_1	0.1
β	$\frac{1}{1.02}$	ψ_2	0.001
k	24	δ_T	0.0
R	1.015	η_0	0.065
		η_1	-0.00108

by 2 parameters), the depreciation rate, returns to experience and the price of child-care. As discussed below, the returns to experience parameters are set using observed wage profiles. Table 6 reports the statistics that the remaining parameters are calibrated to, together with the baseline model values for those statistics. They refer to Cohort 2, that is women aged between 23-27 in 1971. The statistics we match are the average participation across the life-cycle, average participation rates by mothers with children less than 3, the ratio of participation in the year before child birth to participation in the year after, and finally, the median duration of exit from employment. We are comparing the median duration of exit, rather than the mean, to avoid assigning values to individuals who we never observe returning to work. In comparing average participation rates of mothers of children younger than 3, we include all mothers age less than 36 in order to make our sample size large enough. In the table, we also show the statistics for cohort 3. The aim of section **??** on comparative statics is to show what changes may induce the observed changes in these statistics.

Uncertainty Values for the variances of permanent shocks are taken from Meghir and Pistaferri (2002). In our baseline parameterisation, we only allow for permanent shocks. For both adults, the variance of the permanent shocks, σ^2_ξ, are assumed to be constant across the life-time and we assume that shocks to the earnings of the husband and wife are uncorrelated. In the comparative statics, we consider the impact on female participation of changing uncertainty about husband earnings. The omission of temporary shocks for women earnings was only done to simplify the computations. We do not expect the introduction of this type of shocks to affect our results in a substantial fashion.

Table 6: Baseline Model and Main Statistics

	Baseline	*Data, Cohort 2*	*Data, Cohort 3*
Participation	0.68	0.70	0.76
Part Kid	0.44	0. 44	0.67
Ratio B/A	2.03	1.80	1.22
Median duration of exit	4	4	3

The rowr *Participation* reports the average participation rates of women across all ages. The rorw *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The row *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

Interest Rate and Discount Rate. We assume a discount rate equal to 0.02, which is slightly higher than the interest rate, fixed at 0.015.

Risk aversion. 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).

Average age of maternity. In the baseline we assume that one child is born when the mother is 24, which is approximately the average first maternity age for women of Cohort 2. We experiment below with delays in maternity age.

Returns to Experience and Depreciation As we mentioned in Section 3, there are three issues to address in the accumulation of human capital: first, how fast does human capital accumulate when working; second how fast does human capital depreciate when not working; and third, how permanent is the depreciation.

We set on-the-job accumulation so that the earnings growth experienced by a worker who participates in every period matches the PSID data, controlling for cohort effects. For a worker who participates from age 22 to 62, earnings grow by two and a half times, with earnings growth being fastest when young.

We report the equation for the accumulation of human capital in Section 3, for convenience.[5]

[5] This accumulation process differs from the approach taken in Eckstein and Wolpin (1988) and van der Klaauw (1996) where years of experience and years of experience squared are included directly in the earnings equation. Both papers are counting years of experience as human capital, but having the return to human capital diminishing. This can be thought of

$$h_t = h_{t-1} + (\eta_0 + \eta_1 t)\, I\left(P_{t-1} = 1\right) - \delta_P h_{t-1} I\left(P_{t-1} = 0\right)$$

$$\ln y_t = \ln y_0 + h_t - \delta_T I\left(P_{t-1} = 0\right) + \nu_t + \varepsilon_t \tag{12}$$

In table 7, we report the extent of the returns to experience by varying η_0 and showing the effect that being out of the labour force has on lifetime earnings assuming zero depreciation: $\delta_p = 0$ and $\delta_T = 0$. We consider 3 years out - this value is the median length of exit, conditional on reentering employment, for cohort 2. The importance of returns to experience is best illustrated by the difference between the cost of exit in the two columns.

Table 7: Earnings Cost of Labour Market Interuptions

	3 years at 24	*3 years at 25* $\eta_0 = 0$
$\eta_0 = 0.03$	0.13	0.09
$\eta_0 = 0.065$	0.18	0.09
$\eta_0 = 0.1$	0.23	0.09
Eckstein and Wolpin	0.11	0.09
der Klaauw	0.13	0.10

No uncertainty over wages and depreciation rates set to zero.

These calculations assume that wages are known with certainty. The effect of uncertainty and, therefore, ex-post heterogeneity, is to introduce selection into the exit decision. In particular, individuals who exit tend to be those with (ex-post) lower wages. For these women, earnings foregone will be lower than shown above suggesting that these are over-estimates of the monetary cost paid by those who choose to exit. Olivetti (2000) estimates the return to experience, but looking at the effect of varying hours of work, rather than participation. This means we cannot use her estimates for looking at the effect of non-participation on human capital: if we took her estimates as being relevant for the participation choice, then a spell of 3 years out of the labour force at age 25 would cause a cost of 54% of lifetime earnings.

Returns to experience in our baseline are greater than those estimated in Eckstein and Wolpin (1988) and van der Klaauw (1996). As discussed above,

as equivalent to having an accumulation process for human capital in which an extra year of experience has a diminishing effect of total human capital but with the price of human capital being constant.

the estimates of the cost of being out of the labour force in those papers are likely to be underestimates of the true costs: they are determined by the fraction of individuals who do not participate and, given that consumption equals income in those models, it is hard in a dynamic programming model to induce individuals to choose not to participate unless the cost of doing so is small. In section 5 we show the effect of taking these lower estimates of the returns for participation in our model: in the presence of saving, there is too much exit from the labour force if there is no depreciation and the returns to experience are low.

In our baseline, we fix the permanent depreciation rate at $\delta_P = 0.02$, and the transitory depreciation rate at $\delta_T = 0.0$. In Mincer and Polachek (1974) there is only permanent depreciation, δ_p, which they estimate to be around 1.5% per year. There is no controlling for selection issues in that paper but their estimates are the only ones available for the permanent depreciation. Van der Klaauw (1996) and Keane and Wolpin (1997) assume there is only transitory depreciation, δ_T, and in Eckstein and Wolpin (1988) there is no depreciation. Mincer and Olfek (1982) estimate both δ_p and δ_T and find δ_T to be important. Mincer and Olfek (1982) also discuss whether the depreciation rate is not constant. In particular, they discuss the possibility that depreciation rates increase with duration of exit. We discuss the implications of such depreciation in section 6.

In table 8, we show the effect on the cost of time out of the labour force for different values of permanent depreciation, ignoring uncertainty. The striking point about this table is how fast the cost of exit increases with permanent depreciation whether or not there are returns to experience. The table does not report differences in earnings loss associated with increases in transitory depreciation. This is because such increases make almost no difference to the earnings loss associated with exit. This is simply because we are considering life-time earnings and in a life-cycle context, the value of δ_T is the one-off cost of returning to work.

Mincer and Polachek (1974) explore the issue of returning to work in between the arrival of children and show that investment in human capital is less by those who work at this stage. In the current paper, we abstract from this issue by assuming that each family has only one child.[6] This assumption also avoids issues of returns to scale in child-care.

[6] We could equivalently assume that families have multiple children but at the same age.

Table 8: Human Capital Depreciation: Earnings Cost of Labor Market Interruptions

	3 years at 24	*3 years at 25* $\eta_0 = 0$
$\delta_P = 0.0$	0.18	0.09
$\delta_P = 0.01$	0.21	0.12
$\delta_P = 0.02$	0.24	0.15
$\delta_P = 0.04$	0.30	0.20
$\delta_P = 0.06$	0.35	0.25
$\delta_P = 0.10$	0.43	0.33

No uncertainty over wages. In column 2, $\eta_0 = 0.065$.

Child Costs Equation (4) determines how the fixed cost of work varies with children. We take the parameters of the function G from Hotz and Miller (1988) who estimate $\theta = 660$ and $\phi = 0.89$ using 1970s data from the PSID. The child cost price p is set to capture both the direct cost associated with the child-care and additional loss associated with that child care[7]. As shown in Table 5 a value of this parameter equal to 18 is needed. Given this, the ratio of child-care expenses for a newborn to average annual earnings of the husband at age 30 is 46%.. In Table **??** we calculate the net present value of the income path from taking a spell out relative to not doing so. If there were no utility cost of participation and no uncertainty, this opportunity cost of non-participation would determine the size of the fixed cost necessary to induce that spell of non-participation. Individuals may choose to work early on despite current earnings being less than the fixed cost if there is a significant return to experience and they are able to finance this investment through borrowing.

Preferences The ψ_2 parameter reflects the direct utility cost of participating and it takes a value equal to 0.1 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. Parameter ψ_1 takes

[7] Including some kind of utility cost.

value 0.001 in our baseline.

4.2 Baseline life cycle profiles

Given the parameters in Table 5 we can simulate the model and generate life cycle profiles for some of the variables of interest. First, Figure 19 shows profiles of participation for simulated data and for actual data This Figure shows that matching to the statistics in Table 6, the model is able to reproduce the salient features of the life cycle profiles of Cohort 2 labour force participation. In particular, the model generates the increase in participation rates with age when children are young. Second, Figure 21 shows the average earnings profile under the assumption of full participation versus the average earnings profile for participating women. We can see that the selection effect means that average earnings of those actually participating are higher than the potential earnings of all women. Figure 21 also show the profile of fixed costs. Finally, Figure 23 shows asset accumulation. In our model, individuals save prior to the arrival of children, then run down savings and borrow while children are young, particularly if not participating. As they get older, debt is paid off, and then individuals save to cover periods of non-participation at the end of the life-time (although there is no exogenous retirement).

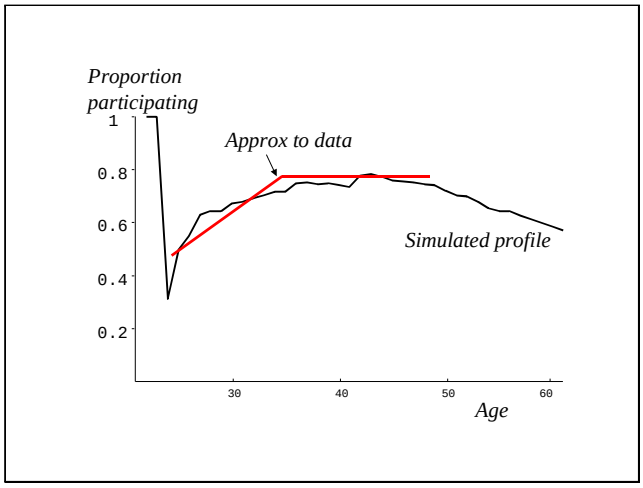

Figure 19: Simulated and Actual Participation Rates over the Life-cycle

Figure 20:

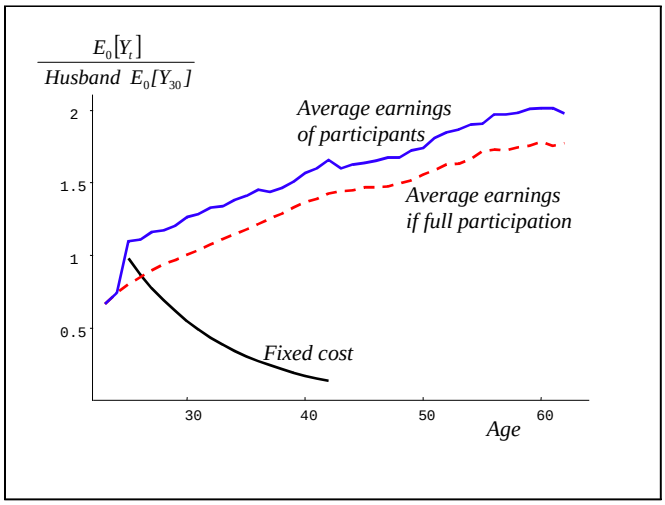

Figure 21: Simulated Earnings, Potential Earnings and Childcare Costs

Figure 22:

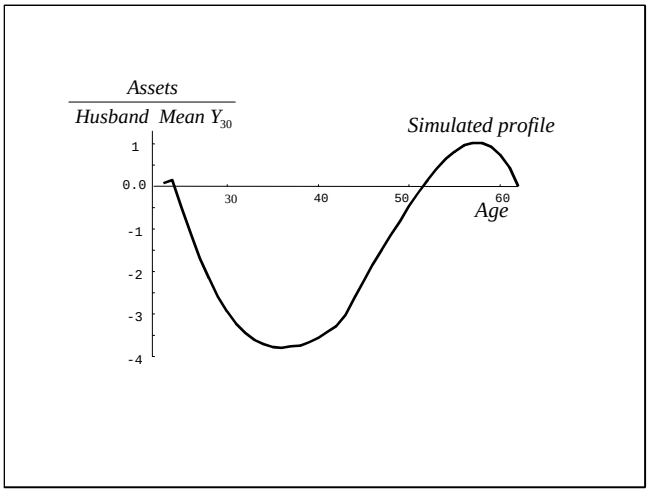

Figure 23: Simulated Asset Accumulation

5 Explaining changes in female participation

In this section we describe the impact that changes in the economic environment have on female participation over the life-cycle.[8] 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 and return to experience. Our strategy is first to establish how much participation changes with given changes in the variable of interest. We then discuss what changes constitute plausible explanations for the observed changes in labour supply.

Child-care Costs One would expect a reduction in the cost of child care to increase participation by mothers. In Table **??** and Figure 24 we show how participation reacts to the price of child-care in our model. In this section, we do not change the way in which the age of the child affects child costs. We discuss in section 6 how changes in the relative costs of child-care over the age of the child may affect participation. In the Table we report, for each level of child cost (and leaving the other baseline parameters unchanged), the average participation over the life cycle, the participation rate of mothers with a child younger than three, the ratio of exit wages to re-entry wage and the median duration of the exit.

A reduction in child-care costs increases substantially participation by mothers of young children and has a small positive effect on overall participation. The reduction in child-care costs can therefore flatten the age - participation profile of young mothers, without having a substantial effect on the level of participation. For instance, moving from the baseline cost of 18 to a cost of 16 (or a bit below) essentially explains the difference in the participation rate of mothers of young children in the second and third cohort. On the other hand, median duration of exits does not decrease in the way we observe in the data. On the contrary, median duration, given our model of participation, increases rather than going down. We return to this issue in the next section.

Returns to experience

As suggested by Olivetti (2001), an increase in the returns to experience may affect labour supply. An increase in the return to experience has two effects: first, it increases the opportunity cost of not participating and second,

[8] For each set of parameter changes, we present the baseline statistics in bold text in the tables and baseline profiles with solid lines in the figures.

Table 9: Childcare Costs

p	*Part*	*Part Kid*	*Ratio B/A*	Median Duration
12.0	0.78	0.89	1.19	18
14. 0	0.78	0.86	1.21	15.5
15.5	0.76	0.79	1.29	18
16.0	0.71	0.53	1.86	5
18.0	**0.68**	**0.44**	**2.03**	**4**
20.0	0.67	0.34	2.79	3.5
22.0	0.64	0.25	4.09	3

The column *Part* reports the average participation rates of women across all ages. The column *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The column *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

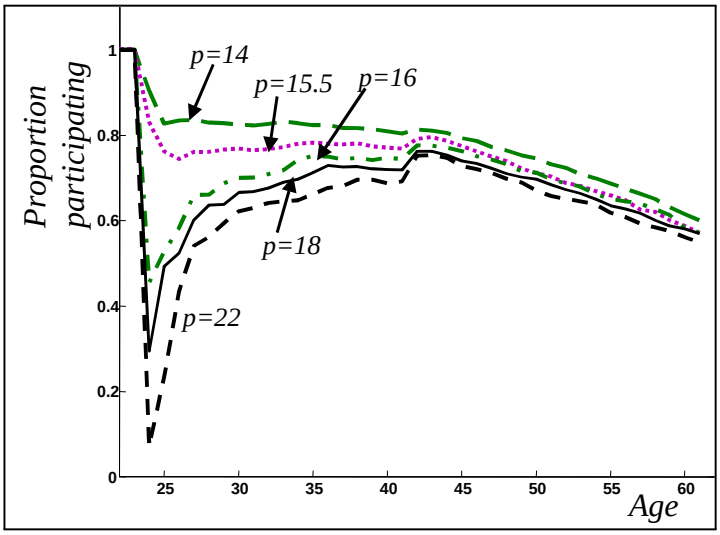

Figure 24: Participation over the Life-cycle, varying childcare costs

it increases potential lifetime wealth. In carrying out comparative statics, we want to isolate these two effects. The increase in lifetime income means that the proportion of lifetime earnings that would be spent on child care are smaller and this makes the working decision during early motherhood less costly.

In Table 10 we report the same participation statistics reported in Table **??** for different returns to experience, without correcting for the change to lifetime wealth. We also report the effect of the increased returns on the opportunity cost of exit. This table shows that an increase in the return to experience can have a large effect on participation, in particular inducing greater participation during early motherhood. Greater returns to experience are however, associated with increases in the median duration of exits. This increase in median durations is partly a composition effect in that some women who would choose to exit in the baseline environment are no longer exiting. Offsetting this composition effect is the impact on behaviour of those who would choose to exit in both the baseline case and the case with greater returns: greater returns to experience should induce shorter durations for those who exit. In the simulations, the composition effect dominates, whereas in the data, the behavioural effect dominates. In Figure 26 we show the corresponding life-cycle profiles for participation. The important question here is how large are the increases in returns necessary to induce the changes in participation. We show the implications of these different scenarios for earnings in Figure 25

In the absence of returns to experience, individuals are more able to intertemporally substitute their labour supply, working when the wage is high or child costs low and exiting when they receive a bad wage shock or face high child-care costs. Hence, there is more exit for mothers with infants.

We now try to disentangle the pure dynamic effect induced by an increase in the return to experience from the wealth effect. There are several ways we can change the return to experience and keep expected lifetime income constant. First, in Table 11, we adjust the initial wage for women, y_0, for each value of the return to experience to hold constant the net present value of life-cycle earnings for a women who participates in every period. The key point about this table is that increases in returns to experience lead to lower participation overall and lower participation by mothers of infants. This suggests that the increases in participation seen in Table 10 are due to the increase in lifetime earnings associated with the increased returns to experience rather than to the changed return to accumulating human capital per se.

Second, in Table 12, we adjust the husband's initial wage in each of the

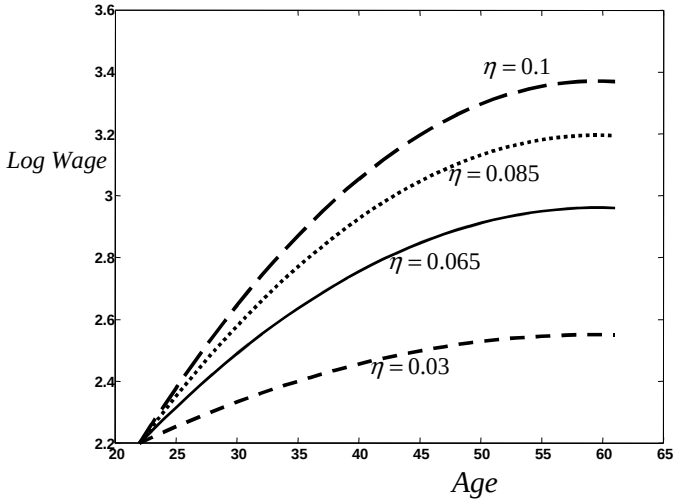

Figure 25: Log Wage by Age, Varying Return to Experience

Table 10: Returns to Experience

η_0	*Opp Cost*	*Part*	*Part Kid*	*Ratio B/A*	Median Duration
0.0	0.17	0.52	0.21	4.93	4
0.03	0.20	0.62	0.40	2.06	4
0.065	**0.24**	**0.68**	**0.44**	**2.03**	**4**
0.080	0.29	0.72	0.52	1.88	8
0.083	0.30	0.75	0.68	1.46	14
0.085	0.30	0.79	0.86	1.16	
0.10	0.33	0.84	0.92	1.11	28

The column *Opp Cost* reports the fraction of lifetime earnings lost for a 3 years spell of non-participation at age 24. The column *Part* reports the average participation rates of women across all ages. The column *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The column *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

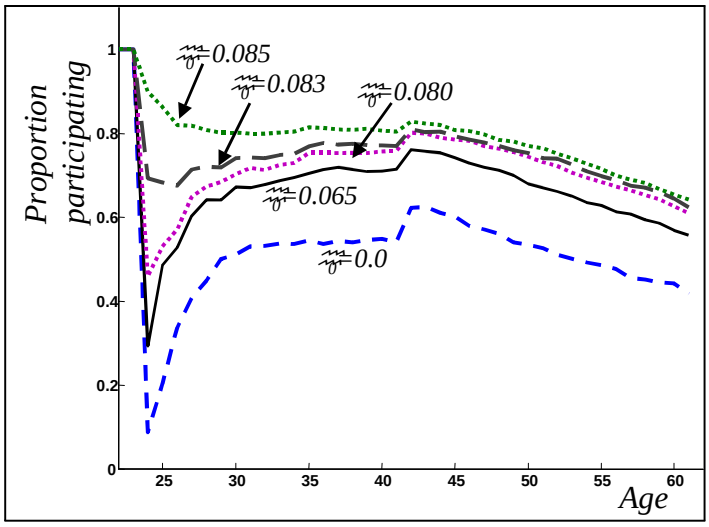

Figure 26: Participation over the Life-cylce, varying returns to experience

Table 11: Returns to Experience, Wage Level Adjusted

η	$\ln y_0$	*Part*	*Partkid*	*Ratio bef/after*	*Median duration*
0.0	1.51	0.70	0.61	1.74	7
0.03	1.28	0.70	0.59	1.50	5
0.065	**1.00**	**0.68**	**0.44**	**2.03**	**4**
0.085	0.83	0.65	0.34	2.82	4
0.10	0.71	0.62	0.18	5.70	3

The wage schedule is shifted up or down to hold constant the net present value of earnings for a woman who participates in every year.

Table 12: Returns to Experience, Husband's Wage Adjusted

η	$\ln y_0^m$	*Part*	*Partkid*	*Ratio bef/after*	*Median duration*
0.0	1.34	0.42	0.16	6.37	4
0.03	1.22	0.56	0.39	2.18	2
0.065	**1.00**	**0.68**	**0.44**	**2.03**	**4**
0.085	0.80	0.83	0.87	1.16	11
0.10	0.58	0.96	1.00	1.0	no exits

The husband's wage schedule is shifted up or down to hold constant the net present value of household income for a household where the woman participates in every year.

returns to experience scenarios to hold constant the net present value of total family income for a household where the women works full-time. When this adjustment is made, an increase in the return to experience for women implies lower relative earnings for men because lower male earnings are necessary to hold household lifetime earnings fixed. Female participation increases faster as returns to experience increase with this adjustment than in the scenario with no adjustment. This is primarily because the lower (relative) male wage means the marginal utility from the extra consumption associated with an extra hour of work is greater.

Depreciation Table 13 shows that increases in depreciation have little effect on the participation statistics. Depreciation only impacts on participation when depreciation rates are very high. For values of depreciation between 0.0 and 0.06 we do not observe much difference relative to the baseline simulation. However, the result that depreciation has very little impact on participation is somewhat dependent on the ratio of the wage to child-care costs: if the ratio is higher and so child-care costs less important, then increases in depreciation can have a big impact. From these considerations we conclude, therefore, that the impact of depreciation is only second-order and is relevant only when the decision about participation is marginal.

In our model, the rate of depreciation of human capital is independent of the duration of the spell out of the market. There is some evidence, however, that this assumption is too restrictive in that human capital depreciation seems increasing in duration. Jacobson, LaLonde and Sullivan (1993), for instance,

Table 13: Human Capital Depreciation

δ_P	*Part*	*Part Kid*	*Ratio B/A*	Median Duration
0.0	0.69	0.43	2.07	4
0.01	0.68	0.44	2.02	4
0.02	**0.68**	**0.44**	**2.03**	**4**
0.04	0.66	0.46	1.95	6
0.06	0.63	0.43	2.11	8
0.15	0.56	0.38	2.56	18

The column *Part* reports the average participation rates of women across all ages. The column *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The column *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

found that long-term displaced workers experience large and enduring earnings losses.[9] This pattern of depreciation could contribute to explain the reduction in the median duration of exits of females.

Uncertainty Changing the degree of uncertainty about husband earnings has an effect on individual behaviour through changing the precautionary motive to save and to participate. In table 14 and figure 27, we consider the effect of increasing the variance of the permanent shock to husband earnings. The precautionary effect means an increase in uncertainty leads to increased participation by mothers of young children. Increases in uncertainty lead to a flattening of the age - participation profile associated with young infants, and does not lead to large increases in the overall level of participation. However, the increase in uncertainty necessary to obtain this effect is implausibly large (Meghir and Pistaferri, 2003).

The degree of the 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. This is important because

[9] Ljungqvist and Sargent (1998) argue that this, together with some features of the Welfare State, can contribute to explain the persistently high unemployment in some European countries.

Table 14: Varying Earnings Uncertainty for the Husband

$\sigma^2_{\xi,m}$	*Part*	*Part Kid*	*Ratio B/A*	Median Duration
0.010	0.71	0.51	1.92	9
0.031	**0.68**	**0.44**	**2.03**	**4**
0.062	0.69	0.52	1.73	4
0.093	0.74	0.82	1.25	18

The column *Part* reports the average participation rates of women across all ages. The column *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The column *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

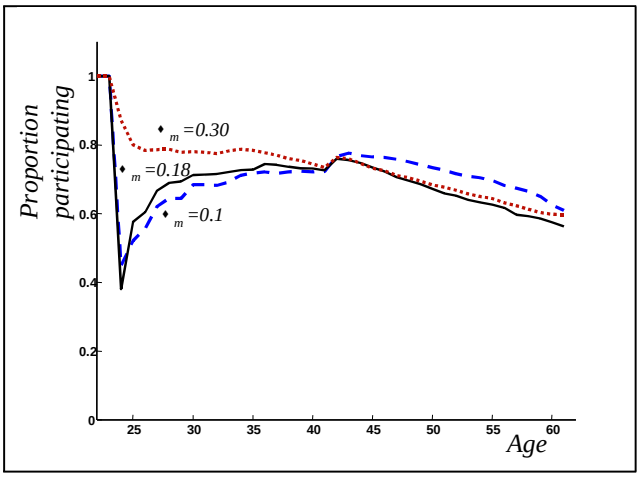

Figure 27: Participation over the Life-cycle, varying income uncertainty

of the assumption of the absence of persistence of shocks in Eckstein and Wolpin (1989) and van der Klaauw (1996).

Maternity Age In our baseline, we set the age of mothers at childbirth at 25. As discussed in section 2, age of mothers at childbirth has increased markedly in recent cohorts and so in Table 15 and in Figure 28 we consider how changes in the age of mothers at childbirth affects participation.

Table 15: Varying Age at Childbirth

Age of childbirth	*Part*	*Part Kid*	*Ratio B/A*	Median Duration
23	0.65	0.33	2.25	3.5
24	**0.68**	**0.44**	**2.03**	**4**
27	0.71	0.57	1.75	13
32	0.71	0.56	1.55	12
37	0.71	0.55	1.56	14
No kids	0.82			

The column *Part* reports the average participation rates of women across all ages. The column *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The column *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

Mothers who have children later, tend to participate more while their children are infants, but this effect is not monotonic. Further, overall participation rates remain fairly constant as age at childbirth increases and the ratio of participation before to participation after motherhood decreases. This means that an exogenous change of the maternity age could explain the fewer exits observed in recent cohorts. This is itself partly due to the returns to experience in the model: once women have accumulated a large amount of human capital, the incentive to leave the labor market are smaller because of the higher opportunity cost it has. The delay of marriage due to the availability of the pill, argued by Goldin (2002), could imply an increase of women employment by itself.

However, median duration increases with age of childbirth. This is again due to a composition effect: as age of childbirth increases, the opportunity cost of exit increases. Women with the highest opportunity cost who previously exited for only a short period, no longer exit. For those that remain, their length of exit may be lower than previously but the composition effect means that median

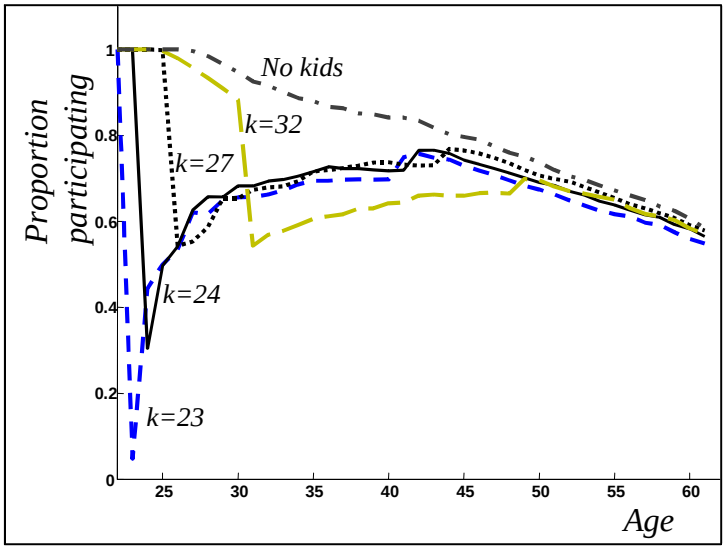

Figure 28: Participation over the life-cycle, varying age of childbirth

duration of exit increases.

6 Implications and Conclusions

It is now time to take stock on the simulations performed in the previous section and discuss what can be learned from them in terms of explaining the difference in the behaviour of the Hillary Clinton and Oprah Winfrey cohorts. we discussed in Section 2. We simulate three main changes to important determinants of female labour supply. While we do not make an explicit attempt to match the size of these changes in actual data (partly because it is very hard to identify and measure these phenomena), our simulations serve the purpose of evaluating the relative merits of different possible explanations of the observed changes in female labour supply. In particular, having matched the behaviour of the Hillary Clinton cohort in our baseline simulations, we check 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 third, when we increase the depreciation rate of human capital when out of the labour market. In addition, we also simulate changes in husband earnings uncertainty and maternity age.

First, we observe that a relative large decrease in child care costs would be needed to explain the increase in participation of mothers of young children from the observed 0.44 rate in Clinton cohort to the observed 0.67 in the Winfrey cohort. Further, reducing child care costs increases median duration of time non-participating. This is the opposite of what is observed in the data.

Second, while a sizeable increase in the returns to experience seems to be able to explain the increase in participation of mothers of young children, we should stress some interpretation issues. The increase in participation seems to be more related to the wealth effect implied by the increase in the return to experience than by the increased cost of being out of the labour force. As with child costs, the increase in participation leads to an increase in mean and median duration of exit for those who do move out of the labour market, and this, once again, is in contrast with what we observe. Notice that the increase in participation by mothers of young children does not translate into large increases in average participation over the life-cycle because part of the increase in participation when children are young reflects an intertemporal substitution effect: women work during motherhood, but then exploit their accumulated returns to experience and reduce their participation later in life. An increase in the return to experience changes both the opportunity cost of exit, but also the level of life-time earnings. Once we (counterfactually) control for the increase in life-time earnings, we do not observe the same reallocation of participation over the life-cycle. Indeed, the effect of increased returns to experience goes in the opposite direction, leading to a reduction in participation when children are very young. Finally, notice that the increase in return to experience necessary to obtain the observed shift in participation early in the life cycle is substantial. In Section 2 we cast some doubt on the size of the change in return to experience.

Third, increases in depreciation rates can have a large effect on the opportunity cost of exit. However, somewhat surprisingly, these increases do not seem to have a significant effect on the participation decision. The impact of depreciation appears to be second-order. Only when the decision about participation is marginal do changes in depreciation have a large effect on participation.

From these three experiments, we conclude that reductions in the price of child-care (relative to earnings) can reproduce the observed changes in the participation statistics. Similarly, a large increase in returns to experience can induce the same changes, but changes in depreciation cannot. Behind these results the same determinants operate: for households making decisions in a life-cycle context where they are free to save or borrow the most important determinant

of participation is the size of the fixed cost relative to the increase in lifetime earnings associated with participation. In this sense, both the reduced child-cost and the increased returns to experience have the same effect of reducing the ratio of the cost of participation to the increase in lifetime earnings associated with participation. Since the cost of reentry associated with higher depreciation can be spread over the remaining lifetime, the cost of non-participation is smaller in our model than if the impact was on consumption in the reentry period only. The possibility for intertemporal smoothing of consumption through saving and borrowing is therefore important to our result that changes to intertemporal incentives to work have only a small effect (except through wealth effects). If borrowing were not possible, the importance of intertemporal incentives to work may increase.

We explored two further explanations of the change in participation which do not involve wealth effects in the same way. First, the observed increase in maternity age does increase participation and flatten the age-participation profile in the direction observed in the data. However, the magnitude of the simulated changes is much smaller than that observed in the data. Second, an increase in husband earnings uncertainty might explain an increase in participation of mothers of young children for precautionary reasons. In the simulations, we observe that increases in uncertainty do induce greater participation. There is a reallocation of participation towards the early stages of the life-cycle, leaving the average level of participation fairly constant. There is an increase in the duration of exit of those who still exit, since these are the women who receive very low wage shocks themselves. In the model, we have considered only permanent shocks to earnings and the inclusion of temporary shocks might induce a higher turnover of early participation and may affect the level. While this is conceptually interesting, it should be stressed that the increase in husband uncertainty necessary to flatten the age-participation profile is implausible.

We were unable, in any of our simulations, to capture simultaneously the flattening of the age-participation profile and the decline in median duration of exits. Fewer women exit to have children, but those that do exit, in the simulations, spend longer out of the labour force. This is a composition effect: those that still exit may have reduced the duration of their exits, but the sample of those out of the labour force has changed. Perhaps surprisingly, in our simulations, this composition effect dominates the behavioural change that induces women to spend shorter time out of the labour force. Because of the

prevalence of the composition effect, however, the simulations do not agree with the observed data. We discuss here two possible explanations that deserve further investigation: first, we consider whether the *structure* of child-costs may have changed; second, we consider whether the depreciation rate may be non-constant.

Hotz and Miller (1988) estimate the time cost of child-care on data from the 1970s and we used their estimates in our calibration exercise. In the experiment in section **??**, we varied the price of this child-care holding the time cost constant. To induce a reduction in median duration, we may change the shape of the G function in equation 4. In particular, it is possible to make the function more convex by increasing parameter ϕ in equation 3. This means that child-care costs initially fall more quickly with the age of the child, but then fall more slowly than in the baseline. This leads to an increase in female participation, particularly for mothers of small children and to a reduction in the median duration of exits, as in the data. In this case, we are making the intertemporal

Table 16: Varying Shape of Childcare Costs

ϕ	*Part*	*Part Kid*	*Ratio B/A*	Median Duration
24	**0.68**	**0.44**	**2.03**	**4**
0.78	0.78	0.81	1.23	3
0.76	0.79	0.81	1.22	2

The column *Part* reports the average participation rates of women across all ages. The column *Part Kid* reports participation rates of mothers is for mothers of children aged 3 or younger. The column *Ratio B/A* reports $\frac{\text{Participation Before Childbirth}}{\text{Participation After Childbirth}}$. Median duration is for all women who exit.

incentives to work more pronounced because of the change in child-costs with age. Again, however, we are in effect changing the size of the cost of participation relative to the benefit of participation. Unfortunately, we have no evidence (one way or another) about changes in the shape of the G function.

Mincer and Olfek (1982) discuss the possibility that the depreciation rate may be non-constant. In particular, the rate of depreciation associated with a long spell of non-participation may be disproportionately large, whereas workers who return quickly face only small depreciation rates. While Mincer and Olfek (and in a different context Jacobson, LaLonde and Sullivan (1993)) find some evidence to support this suggestion, it is not clear that there has been any

change in the way that the depreciation rate increases with length of exit, and it is not clear what the net impact would be on participation.

It should also be mentioned that, in the data, the reduction in median duration of spells out of the labour market is particularly apparent for low education women. For high education median duration is roughly constant. The overall effect partly reflects changes in the composition of the younger cohort in terms of education achievement. These aspects are obviously ignored in the simulations.

The results of this paper are very suggestive and open further lines of enquire. We have discussed extensively the importance that savings have in our model. In the current draft, households can borrow up to the present discounted value of the worse income realization. Given the prominent role played by the possibility of intertemporal smoothing, a natural extension to our model would be the consideration of binding borrowing restrictions in the first part of the life cycle. Another two important issues that we have ignored are the incentives to accumulate human capital 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 is therefore worthwhile to incorporate in more realistic incarnations of the model. Finally, as we mentioned in the previous paragraph, it might be worth extending the model to consider depreciation rates that increase with the duration of the exit out of the labour market.

From an empirical point of view, on the one hand it is important to gather additional evidence on the size of the changes in the various determinants of female labour supply, ranging from the cost of child-care, to the return to experience and depreciation. Unfortunately, especially for the various determinants of wages, as we discussed, it is difficult to disentangle various effects. Selection issues and the difficulty of isolating cohort, time and age effects make the separate identification of depreciation and returns to experience particularly hard. Related to this difficulty are also the general equilibrium effects which cause changes in female labour supply to have important effects on wages.

7 References

References

[1] 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.

[2] 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.*

[3] Attanasio, O.P. and G. Weber (1993): "Consumption, Interest Rate and Aggregation, *Review of Economic studies,*

[4] Attanasio, O.P. and G. Weber (1995): "Is Consumption Growth Consistent with Intertemporal Optimization?", *Journal of Political Economy*,

[5] Blau, D. (1992). "The Child Care Labor Market", *Journal of Human Resource*s, 9-39.

[6] Caucutt, Guner and Knowles (2001), "The Timing of Births: A Marriage Market Analysis". Unpublished manuscript.

[7] Coleman, M. and Pencavel, J. (1993) "Trends in market behavior of women since 1940" Industrial and Labour Relations Review 46: 653-76

[8] Conesa, J.C. (1999), "Educational Attainment and Timing of Fertility Decisions". Unpublished manuscript.

[9] 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

[10] Eckstein, Z. and K. Wolpin (1989) "Dynamic labor force participation of married women and endogenous work experience" Review of Economic Studies 56: 375-390

[11] Erosa, Fuster and Restuccia (2001), "A General Equilibrium Analysis of Parental Leave Policies". Unpublished manuscript.

[12] Ferrero, D. and Iza, A. (2002), "Skill premium effects on fertility and female labour supply'. Mimeo.

[13] Goldin, C. (2002) Journal of Political Economy

[14] Greenwood, J. and Seshadri, A. (2002), "Technological Progress and Economic Transformation," Economie davant garde, Research Report No. 3, University of Rochester.

[15] Hotz, V.J. and R.A. Miller (1988) "An empirical analysis of life-cycle fertility and female labor supply" Econometrica 56(1): 91-118

[16] Household Economic Studies (2001) "Maternity Leave and Employment Patterns: 1961-1995" P70-79. US Census Bureau.

[17] Jacobson, LaLonde and Sullivan (1993), "Earnings Losses of Displaced Workers", *American Economy Review* 83.

[18] Jones, L. , Manuelli, R. and McGrattan, E. (2003) "Why are married women working so much?". Federal Reserve Bank of Minneapolis.

[19] Jones, L., Khan, A. and Thomas, J. (2003) "Parental Support and Married Female Labor Supply". Mimeo.

[20] Keane, M.P. and K. Wolpin (1997) "The career decisions of young men" Journal of Political Economy 105(3): 473-522

[21] van der Klaauw, W. (1996) "Female labour supply and marital status decisions: a life-cycle model" Review of Economic Studies 63(2): 199-235

[22] Ljungqvist, L. and Sargent, T. (1998), "The European Unemployment Dilemma", *Journal of Political Economy* 106 (3).

[23] Mincer, J. and S. Polachek (1974) "Family investment in human capital: earnings of women" Journal of Political Economy 82(2): 76-108

[24] Mincer, J. and H. Olfek (1982):"Interrupted Work Careers: Depreciation and Restoration of Human Capital", *Journal of Human Resources,* 17, 3-24.

[25] Olivetti, C. (2000), "Measuring Changes in Returns to Experience: Learning-by-doing and Gender Differentials". Unpublished manuscript.

[26] Olivetti, C. (2001), "Changes in Women's Hours of Market Work: The Effect of Changing Returns to Experience". Unpublished manuscript.

[27] Pencavel, J. (1998) "The Market Work Behavior and Wages of Women: 1975-94." Journal of Human Resources 33(4):771-804.

[28] Sánchez-Marcos, V. (2002) "Women's Employment and Fertility in Spain over the last twenty years". Mimeo.

[29] 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.

8 Appendix

Table 17: **Number of observations at each age by Cohort**

Cohort 1		Cohort 2		Cohort 3	
Age	Obs	Age	Obs	Age	Obs
35	312	25	383	25	385
36	303	26	387	26	386
37	290	27	384	27	409
38	278	28	360	28	418
39	278	29	363	29	417
40	275	30	360	30	414
41	272	31	360	31	404
42	262	32	344	32	401
43	256	33	333	33	395
44	246	34	323	34	384
45	236	35	309	35	383
46	233	36	299	36	389
47	232	37	299	37	367
48	225	38	295	38	358
49	214	39	294	39	386
50	209	40	289	40	276
51	201	41	285		
52	191	42	284		
53	184	43	247		
54	175	44	268		
55	167	45	261		
56	167	46	255		
57	152	47	242		
58	149	48	223		
59	152	49	236		
60	105	50	170		