QUMS WORKING PAPER SERIES
(Finance)

http://www.qub.ac.uk/schools/QueensUniversityManagementSchool/Research/ResearchandResearchers/WorkingPapers/ByGroup/

The Economics of Fertility Timing: An Euler Equation Approach.

David Canning, Declan French and Michael Moore.

Harvard School of Public Health, Queen's University Belfast, Warwick Business School

Working Paper FIN 14-01

QUEEN'S UNIVERSITY MANAGEMENT SCHOOL
WORKING PAPER SERIES
Queen's University Belfast
185 Stranmillis Road
Belfast BT9 5EE

October 2014

The Economics of Fertility Timing:

An Euler Equation Approach

David Canning,[1] Declan French[2] and Michael Moore[3]

Abstract

We develop a dynamic model of fertility, female labour supply and consumption to explain birth timing, particularly why more educated women delay fertility longer. We express the birth timing decision in an Euler equation framework by treating the probability of fertility in each period as a continuous choice variable, with actual fertility a random outcome given this probability. Within this framework, it is easy to see the effects of economic forces on fertility timing decisions. Using US data we show that women with higher levels of education delay fertility because they can accrue greater benefits from work experience.

Keywords: births, female labour supply, optimisation.
JEL Classification: J13, J31

1. Harvard School of Public Health
2. Queen's University Management School, Belfast
3. Warwick Business School

Acknowledgements

We would like to acknowledge funding for this research from The William and Flora Hewlett Foundation, Grant #2011-6455.

1 Introduction

Our aim is to create a fully optimising dynamic model of fertility, female labour supply and consumption decisions. Using an Euler equation approach we derive three first-order conditions that must be satisfied, one for each decision. These conditions have simple intuitive explanations and allow us to see how the exogenous variables in the model affect decisions and how these decisions interact to determine outcomes. We then estimate the parameters of the model using the method of moments applied to data from the United States. A key to our approach is to think of fertility choice as a continuous variable. Rather than decide simply whether or not to have a child at each period, a women decides on a level of fertility effort (or contraceptive effort if their desired fertility is below the natural rate), which determines the probability of having a child.

Women who want to have a child can raise their fertility effort by increasing the frequency of sexual intercourse and the timing of intercourse or fertility treatments, which will raise the probability of becoming pregnant; they can also adopt health-related behaviours to help avoid foetal death. On the other hand, women who do not want to become pregnant can lower their fertility by avoiding sexual intercourse, undertaking long periods of post-partum lactation, using and maintaining contraceptive methods or through abortion. Contraception comes with costs, financial as well as in terms of side effects. It is also not failsafe. For example, 8% of women using oral contraceptives have an unintended pregnancy per year; this compares to a 0.3% rate for women with perfect adherence in trials (Trussell, Lalla et al., 2009) . The costs of perfect adherence, in money, time, and concentration, are high and women may prefer to risk conception than pay them.

We think of these conception and contraception actions as a continuum that allows women to regulate their probability of fertility away from its natural level – but at a cost. We take actual fertility to be a discrete outcome based on this probability. This allows us to derive and estimate a simple first-order optimal condition, an Euler equation, for fertility (or contraceptive) effort. Our approach is similar to that used by Newman (1988) who examines the timing of fertility in response to child mortality with the choice of continuous level contraceptive control to lower fertility from its natural level. We differ in allowing two-sided control of fertility around its natural level and also by combining the model with endogenous female labour supply and consumption decisions.

We estimate the model using data on fertility, labour supply and consumption decisions observed in the US National Longitudinal Survey data on young women (1968–2003), which

gives us panel data on a cohort of women over their entire reproductive lives. Our fertility Euler equation shows that the expected change in fertility between one period and the next depends on the expected change in labour supply between the next period and the following one, with the fall in expected labour supply associated with a rise in expected fertility effort. It also depends on the level of consumption, adjusted for household composition, with a sign that depends on whether children and consumption are substitutes or complements in the utility function. Estimates from our consumption Euler equation imply that children and consumption are complements, so that high consumption tends to raise current fertility relative to future fertility. Our parameter estimates imply that, as found by Ahn (1995) , households have a direct welfare benefit from children – ignoring the time costs, an additional child would give higher utility.

Estimates from our labour supply Euler equation imply that more highly educated women, who have higher prospective wages, tend initially to have high labour supply; however, for them this falls faster over time than it does for less educated women, as the incentive to gain experience is higher for the more highly educated.

Our explanation for the delay in fertility of highly educated women is therefore very simple. Women have an incentive to work more early in their career than later since early work has experience gains that affect future earnings while later work does not affect earlier earnings. Highly educated women have higher expected future wages and labour supply, and hence earnings, so that the gains from experience, which are multiplicative (rather than additive) on earnings, are larger for these highly educated women. Highly educated women therefore work more than less educated women when young and delay fertility to later in life, when their expected labour supply is lower, and this labour supply effect dominates their incentive for early fertility, as they are richer and have higher consumption.

The structure of the paper is as follows. In the next section we discuss the relationship between our contribution and the existing literature. In section 3 we develop the theoretical model and in section 4 parameterise it in an empirically tractable form. Section 5 describes the data set and the results from estimating the model are reported in section 6. We conclude in the final section.

2 Relation to Previous Literature

The theory of fertility choices often uses a static model to examine the question of why the total fertility rate varies across socio-economic groups and over time in the United States. Various arguments have been put forward for these relationships, including income effects,

time costs of children and wage effects, and heterogeneous preferences (Jones, Schoonbroodt et al., 2011). A full model of fertility includes timing as well as the total number of births. However, in a model of fertility timing the fact that fertility, female labour supply and household consumption are jointly determined each period, and decisions are forward-looking, makes the model very complex to analyse. Arroyo and Zhang (1997) survey the theoretical and empirical approaches to the timing of fertility decisions.

Our approach contrasts with the dominant approach in the literature, which is to think of fertility as a discrete decision, with two states, and model the discrete dynamic optimisation problem. With a finite time horizon these discrete choice models can be solved by backward induction. Wolpin (1984) uses this approach to analyse optimal fertility in response to child mortality over time. Francesconi (2002) allows for joint decisions on both fertility and female labour supply in a discrete dynamic model. Sheran (2007) has a model with discrete fertility, labour supply, schooling and marriage and similarly solves it by backward induction.

Assuming fertility decisions are discrete makes solving the models in these papers very complex. With a finite horizon the model can be solved by backward induction over all possible time paths for a particular set of utility parameters and random shocks. The number of such paths generated by all possible combinations of possible choices, and random shocks, at each point in time is generally very large and often the number of states and possible shocks are severely limited to ease estimation. Given optimal choice paths for each set of variables and random shocks, parameters are chosen to maximise the likelihood of the observed choices given by the data. The complexity of the model means this is often implemented by the simulation of outcomes for a set of possible parameters and then choosing between these. It is difficult to interpret the forces at work in fertility decisions in this approach. The utility function gives rise to an optimisation problem, and we can use backward induction to solve for the parameters that best fit the data, but we have little insight into the nature of the forces at work in fertility decisions.

An alternative to this approach is to take a reduced form model in which we use the fact that the optimal fertility and other decisions must be functions of the information set at the time of decision making to model decisions as functions of all variables in the model, and their lags, and to estimate a simplified version of this reduced form. Moffitt (1984) takes this approach to estimate a dynamic model of fertility and female labour supply, allowing for potential wages to vary with work experience. Bloemen and Kalwij (2001) estimate a reduced form model of female labour supply and fertility for the Netherlands and find that women

with higher education are more likely to be employed and to delay fertility. Del Boca and Sauer (2009) specify a dynamic model of female labour supply and fertility and estimate it using simple approximate decision rules that are a possible reduced form. Eckstein and Lifshitz (2011) take a hybrid approach, analysing a fully optimising dynamic model of female labour supply but taking fertility each period to be a simple function of age, schooling, the number of previous births and employment status.

Our Euler equation approach is similar to a reduced form model in that we derive first-order conditions between observables that should be satisfied. However, our approach allows us to characterise from theory which variables should be in each estimating equation and the functional form of the relationship, which helps our understanding the relationships, and also enables comparative statics to be undertaken. Our estimating equations are much simpler than those in the usual reduced form approach since we can exclude all variables that do not appear in the Euler equation. In addition, we find it is the expected value of future variables rather than lags that matters for current decisions, and we include these future expectations by instrumenting future variables with the current information set, including lags, rather than adding lags in an ad-hoc fashion.

The difficulty posed by jointly estimating fertility and female labour supply means that a common approach in explaining the time pattern of female labour supply is to avoid the issue, and to treat the timing of fertility as exogenous, as do Eckstein and Wolpin (1989) . Olivetti (2006) takes a similar approach and explains rising participation by married women in terms of a rise in the returns to experience for women. Attanasio et al. (2008) estimate an optimising dynamic model of female labour supply and inter-temporal consumption but treat fertility as exogenous. An advantage of our approach is that it has a very simple set of first-order conditions (Euler equations) for each endogenous decision variable. This makes the decisions in the model easy to interpret and also allows the joint estimation of dynamic fertility, labour supply and consumption decisions.

The model closest in spirit to ours is that of Happel et al. (1984), who assume there is only one birth per woman, which has a fixed cost in terms of labour supply forgone, and then examine the optimal timing of this birth with consumption smoothing. We generalise to optimal timing over all births, with optimal labour supply responses and consumption smoothing.

Our model has the advantage of being very simple to understand and estimate. It is not fully realistic. We treat the schooling decision and changes in household size other than through fertility (for example, marriage) as exogenous; a full model would make these

decisions endogenous. Waldfogel (1998) argues that there is a pay penalty for women with children in the United States which Wilde et al. (2010) suggest comes about because children lower the return to experience for women. We do not model a direct effect of having children on pay, which would give an extra incentive to delay childbearing. We allow for heterogeneity among women in their productivity and wages through a fixed effect, but do not allow heterogeneity in preferences.

3 Theoretical Framework

We assume that the woman is the central decision-maker for fertility, her labour supply and household consumption. We define the dynamic maximisation problem facing the woman at each time t as:

$$\underset{c_t, l_t, f_t}{Max}\left[U\left(c_t, f_t, l_t, g_t, n_t, y_t, t\right) + E_t\left(\sum_{s=t+1}^{T} \beta^s U\left(c_s, f_s, l_s, g_s, n_s, y_s, t\right) \middle| c_t, f_t, l_t \right) \right] \quad (1)$$

Her utility in each period t depends on three choice variables: c_t is family consumption, f_t is her fertility effort , l_t is her labour supply. There are two state variables: n_t is the number of children she has and y_t the number of children aged under two years[1] (in our empirical application the period of measurement is two years), both of which depend on realised fertility. Utility also depends on g_t, the number of adults in the household. For simplicity we treat this as an exogenous random variable rather than a state variable. We also allow utility to vary with time t. The number children and young children in her family are known at time t when she makes her current choices, but future values of these are considered as random variables that evolve over time given her choices. In her decision-making she takes into account the effect of her current decisions on expected future utility, discounted at the rate β, and assuming that future choices are made in an optimal fashion in the same way as at time t, given the information available at that future time. The instantaneous utility function $U\left(c_t, f_t, l_t, g_t, n_t, y_t, t\right)$ is presumed to be concave in consumption, fertility effort and leisure.

In addition to the state variables in the utility function, we have two important economic state variables at each time t. The first is household wealth, given by w_t and the second is the woman's work experience given by e_t. These state variables do not enter the

[1] We use a separate state variable for young children as childcare is more intensive in mothers' time when the child is very young (see Table I in Smith et al., 2001).

utility function but will affect the budget constraint. The wage the women can earn in period t, given by p_t depends on experience e_t.

In each period t, the woman chooses consumption, fertility effort and labour supply given the current state variables to maximise the sum of current and expected, discounted, future utility. Actual fertility F_t is a discrete outcome; the actual number of births in the period that depends on fertility effort. We think of fertility effort as a continuous choice variable. A more detailed approach would be a dynamic model of a range of choices that affect fertility, such as the selection of contraceptive method (see Montgomery, 1989) .

Her choices, plus random shocks, determine the next period's state variables. The sequence of events in each time period is summarised in Figure 1. A feature of our model is that we think of the woman as the decision-making unit independent of the other adults in the household. An alternative approach would be to think of fertility as a joint decision of a woman and her partner in which the partner's preferences would also matter through a bargaining process. In this case assortative mating issues would become important. As in Bongaarts (1978), our approach views the absence of marriage as a method of fertility control. In the era prior to the widespread availability of contraception and abortion the major method of regulating fertility was to delay the age of marriage and sexual debut. In this world the decision to marry is essentially a decision to increase fertility effort. In the United States the link between marriage and fertility has weakened considerably with contraception lowering fertility within marriage and high rates of fertility outside marriage (Pagnini and Rindfuss (1993) . A very different alternative theory from ours about why highly educated women have later fertility, is that they are more selective, and so take longer to find partners (Caucutt et al., 2002) .

The Bellman equation for this problem is:

$$V(n_t, y_t, w_t, e_t) = U(c_t, f_t, l_t, g_t, n_t, y_t, t) + \beta E_t \{V(n_{t+1}, y_{t+1}, w_{t+1}, e_{t+1})\} \qquad (2)$$

where the value function V is the sum of current and future expected utility associated with the current state variables assuming all future decisions are optimal, which can be defined recursively. The woman maximises her lifetime utility subject to the equations of motion of the state variables given by:

$$w_{t+1} = r_{t+1}(w_t + l_t p_t(e_t) - c_t) \qquad (3)$$

$$e_{t+1} = e_t + l_t \tag{4}$$

$$n_{t+1} = n_t + F_t \tag{5}$$

$$y_{t+1} = F_t \tag{6}$$

Equation (3) gives the evolution of wealth. The stock of wealth at time t+1 is wealth at time t plus wage income, less consumption, multiplied by the rate of return. r_t is the gross real rate of interest and $p_t(e_t)$ real wage rate at time t. Note that wages will depend on work experience. Households are allowed have negative wealth, i.e. they can borrow in order to smooth income intertemporally.

Equation (4) gives the evolution of work experience: experience increases by the amount of labour supply in the current period. The number of children the woman has in period t+1adds the realised fertility in the previous period given by F_t to the previous number of children. The number of young children in period t+1, is simply F_t. For simplicity we do not allow for child mortality, which is very low in our sample.

The future interest rate r_{t+1} is taken to be a exogenous random variable that is not perfectly known at time t. Similarly the wage rate p_{t+1} (for simplicity we make the dependency of the wage on experience implicit) of the woman at time t will have a random component. The actual fertility of the woman at time t given by F_t is also random but we impose the condition that $E(F_t) = f_t$ so that we can think of the woman choosing her expected fertility. f_t would be the probability of fertility if all births were singletons.

From the first-order and envelope conditions (see Appendix for details), we derive the following Euler equations:

$$\textit{Consumption}: \qquad U_{ct-1} = \beta E_{t-1}\left(U_{ct} r_t\right) \tag{7}$$

$$\textit{Labour}: \quad U_{lt-1} + p_{t-1}U_{ct-1} = \beta E_{t-1}\left[\left(U_{lt} + p_t U_{ct}\right) - l_t \frac{\delta p_t}{\delta e_t} U_{ct}\right] \tag{8}$$

$$\textit{Fertility}: \quad U_{ft-1} + \beta E_{t-1}U_{nt} + \beta E_{t-1}U_{yt} = \beta E_{t-1}U_{ft} + \beta^2 E_{t-1}\left(U_{yt+1}\right) \tag{9}$$

The equations are derived from the point of view of a woman making a decision at time t-1 and all future variables involve expectations based on information available at time , due to

the random elements in future interest rates and wages and in actual fertility outcomes. We use the expectation operator E_{t-1} where the subscript denotes the timing of the information set available. The terms in the Euler equation are all marginal utilities where the first subscript on the utility function denotes the variable with repect to which we are differentiating and the second is the time period. The Euler equations (7), (8) and (9) can be thought of as implications of the fact that reallocating consumption, labour supply or fertility from one period to the next cannot raise expected utility to the maximum.

Equation (7) is the consumption Euler equation. The left-hand side is the marginal utility of an extra dollar of current consumption. The right-hand side is the gain in expected utility if the women saves an extra dollar and consumes it in the next period, adding any interest to it but discounting this future consumption. It gives the usual result that consumption is smoothed over time so that the expected marginal utility of an extra dollar of consumption and saving is equalised. This equalisation is exact if the rate of return is deterministic and $\beta r_t = 1$ so that the rate of return on savings exactly offsets discounting of future consumption.

Equation (8) is the labour supply Euler. If a woman works an hour more this period, at time t-1, and spends the income generated, she gets the marginal utility of labour which we take as negative, but gains the wage times the marginal utility of consumption, which is the left-hand side of equation (8). On the other hand, if she works an hour more next period she gets the future marginal utility of labour plus the future wage times the marginal utility of consumption, which is part of the right-hand side of equation (8). Under both plans work experience will be the same in two periods' time and going forward. However, working in period t rather than period t-1 means the woman loses the experience effect from work in t-1 on her wage in period t, which is the final term in the right-hand side of equation (8). Due to the experience effect, women will typically work more hours early in their working lives than later.

Our main object of interest is equation (9), the fertility Euler equation. A woman can reduce her fertility effort in this period and increase it in the next so as to keep lifetime expected fertility and long-term outcomes the same. This means that at the optimum the woman has to balance the short-term costs and benefits of moving fertility effort between adjacent periods. The left-hand side of equation (9) is the benefit of current fertility effort in period t-1. This is the direct effect of fertility effort on current utility plus the expected utility of having a young child in the next period, so that both the total number of children and the

number of young children increase at time t. Alternatively, she can delay fertility effort to period t. In this case she gets the direct utility effect of extra fertility effort in period t plus the expected utility of a young child in period t+1. Note that under both plans the woman has an extra child from period t+1, so this effect is cancelled. In principle having a child a period earlier will mean that this child will also leave home a period earlier. however, this effect is in the distant future and we assume it is negligible due to discounting.

4 Empirical Implementation

The Euler equations (7), (8), and (9) are the first-order conditions for an optimum. In order to operationalise them empirically we need to make an assumption of the explicit form of the utility function and how work experience affects wages.

We use a utility function of the form:

$$U\left(c_t, f_t, l_t, g_t, n_t, y_t, t\right) = e^{\rho g_t} e^{\pi n_t} \log\left(1+c_t\right) - \frac{\omega_l}{2}\left(l_t + \phi y_t\right)^2 + \alpha n_t - \frac{\omega_f}{2}\left(f_t - \lambda_t\right)^2 \quad (10)$$

The utility function depends on the same variables set out in equation (1). The first term is the effect of household consumption c_t on utility, which we assume depends on the number of adults g_t and the number of children n_t in the household. The reason for this choice for the form for the utility of consumption will become clearer when we see implied consumption Euler equation; ρ and π are parameters that measure the effect of an extra adult and an extra child on optimal consumption growth.

The second term in the utility function is the disutility due to working and the time costs of children. We assume working and childcare reduce utility and that the time cost of each young child is equivalent to ϕ hours of work each per year. The parameter ω_l is the disutility weight on labour and childcare.

The third term is the effect on utility of having a child in the household: the direct welfare effect of children. The final term is the cost of fertility effort. We assume that there is a natural level of fecundity, the expected fertility a woman would have without any control on her part, that is, varying by her age and given by the parameters λ_t. For simplicity in our theory we assume the woman is born at time zero and take the time variable to measure age. In our empirical work our cohort of women has slightly different birth years and we use age rather than time dummies in the utility function. Woman can deviate from their natural

fecundity rate but at a cost; ω_f >0 is the utility cost of deviations from the normal pattern of expected fertility. Not having children incurs cost due to contraceptive effort, abortion, abstinence or delay in sexual activity (Bongaarts, 1978), while raising fertility above the normal rate may also have costs.

We assume that the wage p_t at time t is given by

$$\log p_t = \log p_t^* + \gamma e_t - \frac{\delta}{2} e_t^{\,2} \tag{11}$$

where γ, δ are parameters and p_t^* is an exogenous wage effect that we model in our empirical work as depending on a woman's education, a time trend and random shocks. The second and third terms capture the effect of experience, e_t, on wages. We expect $\gamma > 0, \delta < 0$ so that wages increase with experience but at a decreasing rate as experience accumulates. For simplicity, we presume that there is no depreciation in human capital due to absences from the workforce (unlike Mincer and Polachek, 1974). Substitution of equations (10) and (11) into the Euler equations (7), (8) and (9) gives the following system of explicit Euler equations, where we log linearise the consumption Euler equation (see Attanasio and Low, 2004) and the error terms $\varepsilon_{ct}, \varepsilon_{lt}, \varepsilon_{ft}$ are mean zero and orthogonal to the information set at time t-1 (see Appendix for details).

$$\log\left(\frac{1+c_t}{1+c_{t-1}}\right) = \log\beta + \log r_t + \rho\left(g_t - g_{t-1}\right) + \pi\left(n_t - n_{t-1}\right) + \varepsilon_{ct} \tag{12}$$

$$\begin{aligned} &-\omega_l\left(l_{t-1} + \phi y_{t-1}\right) + p_{t-1}\frac{e^{\rho g_{t-1}} e^{\pi n_{t-1}}}{1+c_{t-1}} \\ &= \beta\left(-\omega_l\left(l_t + \phi y_t\right) + p_t\frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t}\left(1 - l_t\left(\gamma - \delta e_t\right)\right)\right) + \varepsilon_{lt} \end{aligned} \tag{13}$$

$$\begin{aligned} &-\omega_f\left(f_{t-1} - \lambda_{t-1}\right) + \beta\alpha + \beta\pi e^{\rho g_t} e^{\pi n_t}\log\left(1+c_t\right) - \beta\omega_l\phi\left(l_t + \phi y_t\right) \\ &= -\beta\omega_f\left(f_t - \lambda_t\right) - \beta^2\omega_l\phi\left(l_{t+1} + \phi y_{t+1}\right) + \varepsilon_{ft} \end{aligned} \tag{14}$$

Equation (12) indicates that the expected household consumption in the next period relative to this period depends on the expected interest rate and discount rate but also on expected changes in the number of adults and children in the household. If ρ and π are positive then households will want to move consumption into periods when there are more household members, which is consistent with diminishing marginal utility of consumption per capita. ρ and π times 100 measures the expected percentage increase in household consumption

with an extra adult and child respectively. We expect that $\rho > \pi$ if the consumption needs of children are lower than those of adults. ρ and π will tend to be smaller the greater the non-rivalries in the consumption of household goods.

The labour Euler equation is quite complex but it is easier to relate to the conventional labour economics literature if we take the case of a single woman living alone who is just starting out on her working life: in effect $n_t = n_{t-1} = y_t = y_{t-1} = 0$ and $g_t = g_{t-1} = 1$. We assume we have $\beta r_t = 1$ so her optimal consumption is steady over time. In this case her optimal time path of labour supply given by equation (13) simplifies to

$$l_t - l_{t-1} = \frac{1}{\omega_l}\frac{e^{\rho}}{1+c_t}\left[(p_t - p_{t-1}) - p_t l_t(\gamma - \delta e_t)\right] + \frac{1}{\omega_l}\varepsilon_{lt} \tag{15}$$

The first term in the square brackets on the right-hand side of equation (15) indicates that women will tend to shift their labour supply into periods where they expect high wages. The second term is the experience effect. For women just starting work, experience will be low and so $\gamma > \delta e_t$. The negative sign on the experience effect implies women will want to work more when they are young and have a declining labour supply over time to benefit from experience. The size of this effect depends on the level of wages and is larger for women earning high wages. This is because while the experience effect is linear in log wages by equation (11), it is multiplicative in the level of wages.

In order to understand the evolution of fertility over time we can rewrite equation (14) taking $\beta = 1$ for simplicity as

$$\begin{aligned} f_t - f_{t-1} &= \lambda_t - \lambda_{t-1} \\ &+ \frac{1}{\omega_f}\left[\omega_l\phi\left(l_t + \phi y_t\right) - \omega_l\phi\left(l_{t+1} + \phi y_{t+1}\right) - \alpha - \pi e^{\rho g_t} e^{\pi n_t}\log\left(1+c_t\right) + \varepsilon_{ft}\right] \end{aligned} \tag{16}$$

Equation (16) indicates that fertility is likely to follow natural fecundity, with deviations due to economic incentives. The size of these deviations depends inversely on the cost of fertility effort; if deviations from natural fertility are very costly the effect of economic incentives on fertility will be small, while if the cost of fertility control is low these deviations will be large. When labour supply is expected to fall over time, expected fertility effort will rise over time; woman want to have the time costs of children when they are working less. Recall, however, that our labour supply Euler implies higher-waged women have a faster decline in labour supply over time since they have greater payoffs to experience – hence our model predicts that higher-waged women will have lower initial fertility and faster rising fertility over time

than lower-waged women. The term α indicates the more women like having children, the earlier they will be fertile to enjoy this flow of utility. The final term indicates that if $\pi > 0$ and children and consumption are complements, women with higher consumption will tend to have their children earlier.

We cannot estimate equation (14) directly since we do not have a measure of fertility effort but using the fact that $E(F_t) = f_t$ we have the moment condition on actual fertility given by

$$\begin{aligned} &-\omega_f \left(F_{t-1} - \lambda_{t-1}\right) + \beta\alpha + \beta\pi e^{\rho g_t} e^{\pi n_t} \log\left(1 + c_t\right) - \beta\omega_l \phi\left(l_t + \phi y_t\right) \\ &= -\beta\omega_f \left(F_t - \lambda_t\right) - \beta^2 \omega_l \phi\left(l_{t+1} + \phi y_{t+1}\right) + u_{ft} \end{aligned} \tag{17}$$

Where $u_{ft} = \varepsilon_{ft} - \omega_f \left(F_{t-1} - f_{t-1}\right) + \beta\omega_f \left(F_t - f_t\right)$ which is mean zero and orthogonal to the information available when decisions are made at time t-1 since actual fertility is determined after decisions are made at time t-1 (see Figure 1).

5 Data

The data are taken from the US National Longitudinal Survey data on young women 1968–2003, which tracks 5,159 women aged 14–24 in 1968. The information collected relevant to this study covers the respondent's schooling, family income and assets as well as their family and household composition and fertility history.

Surveys were conducted in each of the first five years of the study. The Bureau of Labor Statistics then adopted a 2–2–1 year cycle until 1988 after which surveys were biennial. We constructed a variable s, the number of years since the last survey, to account for this irregularity. The discount rate β between periods is therefore replaced in our estimation by β^s. Similarly, the real rate of interest between periods is measured cumulatively over the gap between surveys.

Calculation of the number of adults and children in the household is based on household record questions and includes the respondent herself, all blood relatives, in-laws and adopted/step-/foster children but excludes non-family members living in the household. When surveys were annual, fertility F_t is an indicator variable that takes on the value 1 if the number of the respondent's biological or adopted/step-/foster children is larger in period t+1 than in period t and is otherwise 0. If surveys were more than one year apart and the number of the respondent's children was larger in the later survey, we take the fertility rate F_t = 1/s

where s is the number of years between surveys. For s small we have $E(F_t) = f_t$ the probability of fertility per year.

Total family consumption, was determined by total family income less changes in total net family assets per year. It was expressed in real terms using the Bureau of Labor Statistics Consumer Price Index. Negative consumption data were set to zero (4.6% of all consumption data). Annual labour hours, l_t, were taken from responses to the numbers of hours worked in the week prior to the survey, which were then annualised. Those who were not working or were unable to work were recoded to zero and excessively large responses were truncated to 50 hours (some responses imply working 24 hours a day). Work experience, e_t, was expressed in hours by adding work experience at the previous survey to the hours of work experience since last interview.[2] Hourly rates of pay were measured in real terms using the consumer price index. Descriptive statistics for these variables are given in Table 1. For estimation purposes when women are not working we impute their wages as shown in the next section. We count women in full-time education as working full time, which means school time is assumed to have the same effect on fertility as work time, although school time does not add to work experience.

Figure 2 shows fertility rates by age for women with a high level of education (more than 12 years of schooling, corresponding to some college education) and a low level (12 years of less of schooling, corresponding to high school or below). Women with lower levels of education have high initial and then rising fertility between ages 18 and 22, after which fertility declines steadily. Women with higher levels of education have low initial fertility at age 18 and then rapidly rising fertility rates up to age around 26, when fertility declines in line with women with low education. Completed fertility, the integral of the area under the age-specific fertility rates, is higher for women with lower levels of education.

Figure 3 shows the hours of work of each education group. After age 22 the work time of highly educated women is higher than that of women with lower levels of education. However, before age 22 the former have lower working time due to being in college. Figure 4 combines work and college time. Now we see that the combined work and college time of highly educated women starts off higher and falls faster than for women with low levels of education.

[2] Weeks worked since last survey by labour hours in survey week.

6 Estimation

Our estimation requires wage rates, which we do not observe if a woman is not working. In order to estimate the effect of experience on wages and impute a wage rate for women when they are not working, we estimate the following Mincer equation for wages of a woman i at time t

$$\log p_{it} = \xi * edu_{it} + \gamma * e_{it} - \frac{\delta}{2} * e_{it}^{\ 2} + \tau_t + c_i + \varepsilon_{it} \tag{18}$$

The first term captures the dependence of log wages on education, edu, while the second and third terms capture the concave dependence of log wages on experience. τ_t captures any increase in labour productivity over time, c_i accounts for individual fixed effects in wages and ε_{it} is a random shock. Results from estimating this equation for our sample are reported in Table 2. Column 1 of Table 2 reports the effect of education and experience measured in years on log wages. Each year of education is estimated to raise wages by about 11.6% while the first year of experience raises wages by 4.4%, with subsequent years of experience raising wages by less, due to the negative coefficient on experience squared. The effect of experience on wages has a turning point after around 37 years of work. In our estimation of the Euler equations we measure labour supply and experience in hours and column 2 in Table 2 estimates the relationship in these units. The estimates in column 2 are exactly consistent with those in column 1 and are used for our estimates of the parameters γ and δ.

For women who report working and have observed wages at some stage in their lives we use the estimates of equation (18), reported in Table 2, to impute wages when they are not working.

A difficulty with equation (17) is that the natural fertility rates λ_t are not identified if $\omega_f = 0$, which causes estimation problems.[3] We therefore actually estimate

$$\begin{aligned} &-\left(F_{t-1} - \lambda_{t-1}\right) + \chi_f\left(\beta\alpha + \beta\pi e^{\rho g_t} e^{\pi n_t} \log\left(1 + c_t\right)\right) - \chi_f \beta\omega_l\phi\left(l_t + \phi y_t\right) \\ &= -\beta\left(F_t - \lambda_t\right) - \chi_f \beta^2 \omega_l \phi\left(l_{t+1} + \phi y_{t+1}\right) + u_{ft} \end{aligned} \tag{19}$$

where $\chi_f = 1/\omega_f$.

The Euler equations hold in expectation given the information available when decisions are made at time t-1. This means that, given the true parameter values, the errors from these equations (the amount by which they differ from exact balance) should be

[3] The estimation procedure is iterative so we search over the parameter space. This transformation prevents a singularity from arising.

orthogonal to any variable in the information or choice set at time t-1. This gives us a method of moment estimator where any variable in the information set is a potential instrument and the moments are the products of the Euler equations with these instruments. The instruments we use in our method of moments are all the state and choice variables (other than fertility) measured at time t-1 and at time t-2. Since actual fertility in period t-1, F_{t-1}, is not observed until after fertility effort f_{t-1} is chosen, we include as instruments the fertility outcomes F_{t-2}, F_{t-3}. With this large instrument set the model is over-identified and we minimise the average deviation of the moment conditions from zero. All our moment conditions use the same instrument set, which represents the information available at time t-1. An advantage of the method of moments estimator is that it is consistent under quite general conditions and does not require distributional assumptions about the form of the error terms, as in maximum likelihood.

In principle, the three Euler equations (12), (13) and (19) could be estimated jointly to determine the parameter estimates. However, this proved difficult in practice because of convergence problems. We therefore proceeded in steps. We first estimated the consumption Euler equation (12). The labour Euler equation (13) was then estimated conditional on the parameter estimates for β, ρ, π found in the consumption Euler. Finally, the fertility Euler equation (19) was estimated conditional on estimates of β, ρ, π from the consumption Euler and the estimates of ω_l, ϕ from the fertility Euler equations. This sequential approach corresponds to imposing a particular fixed weighting matrix on the system estimate, and provides estimates that are consistent and asymptotically normal, but may not be as efficient as using the optimal weighting matrix (Hansen, 1982) . The results are reported in Table 3.

From the consumption Euler in Table 3 we see that the discount rate is estimated to be 0.984, which is within the normal range for this parameter. The effect of adding an extra adult to the household is to raise consumption by about 20%, while adding a child raises consumption by about 6%. Households want to save when they do not have children and redirect consumption to periods when they do have them. All the parameters in the Euler consumption equation are very precisely determined.

Estimates from the labour Euler equation give us a figure for the disutility of working parameter ω_l of 0.142 x 10^{-6} that is significant at the 5% level. This is best interpreted in terms of consumption units. For a single woman with no children spending $25,000 a year, and working full time during the year, working one hour less has the same effect on utility as

$5.81 of consumption (the equivalent variation). This money figure for the disutility of working is, as we would expect, similar to the hourly wage rate in the sample as shown in Table 1. The results from the labour Euler equation suggest that a young child is equivalent to 681 hours of work per year, or around two hours of work per day. This point estimate is consistent with Craig and Bittman (2008), who estimate that a young (less than two years old) first child increases the unpaid work of a woman by about 8.1 hours per day, but reduces her paid work by only 2.7 hours a day, so her total working time goes up by 5.4 hours per day on average with a young child. Time spent in childcare does not seem to be a perfect substitute for paid work; rather women reduce their leisure time by far more than their time in paid work when they have to undertake extra childcare as the result of a young child. While our point estimate for the hours of labour supply lost to childcare is reasonable it is not statistically significant. Turning to the fertility Euler equation, our estimate of the direct utility α of a child to a woman is 0.236 and is highly statistically significant. Again, taking the benchmark of a single woman working full time and spending $25,000 a year, a child gives about the same direct utility as around $5,320 of extra spending (not including the effect through making consumption more valuable). The time costs of a child, at 681 hours a year, have a utility equivalent to around $4,232 of spending. The inverse of the weight on deviations of fertility from its natural level is estimated to be 0.212. This corresponds to a weight of 2.72, which means that, for our benchmark woman, the cost of changing the fertility rate by 0.1 is around $477 while changing the fertility rate by 0.2 is $1,868 a year. The non-linearity in the utility function means that large adjustments in fertility become increasingly expensive. Our results for adjustments of this magnitude are in line with the cost of contraception in the United States (Trussell et al., 2009), although it is lower than the high costs of raising fertility above the natural level (Collins, (2002) . Our approach treats the costs of lowering and raising fertility symmetrically and is an average of these costs. While women on average want to have lower fertility than the natural level, if the costs of raising fertility are very high a woman may not lower fertility initially, even though it is cheap to do so, because of the high costs of raising fertility above the natural level later in her reproductive life.

Figure 6 shows our estimated age-specific natural fertility rates from our estimated fertility Euler equation and their confidence intervals. These estimated natural fertility rates are similar to the natural fertility rates found in pre-industrial societies (Knodel, 1978) where there was no evidence of fertility control; in particular we find a fairly linear decline in natural fertility from around age 20 to age 50. A 25-year-old woman has a natural fertility

rate of around 0.36. Our cost of fertility control estimates suggest that reducing this rate to zero would have a utility cost equivalent of losing $5,500. Since the average out-of-pocket cost of abortion in the United States is lower than this – Henshaw and Finer, (2003) estimate it to be just under $500 in real terms over our time period – our results suggest that there may be large direct utility costs of abortion over and above any monetary costs. This is also consistent with the high use of fertility control methods that are more expensive than abortion in money terms.

7 Conclusion

Our approach gives us estimates of the parameters of a simple utility function that we use to explain the timing of fertility. In contrast to the 'black box' optimisation using backward induction our Euler equation approach gives us a simple intuition for why highly educated women delay their fertility longer than women with lower levels of education. Highly educated women have large gains from work experience causing their optimal labour supply to be higher initially and to fall faster as they age. This means that they want to move childcare into later periods when they are working less. The parameters of the model we estimate are reasonable and support this view. One of the major benefits of our approach is that we can interpret fertility and labour effects in terms of consumption units.

The model could be developed in several ways. A key issue is the cost of fertility control. This may be better modelled as having asymmetric costs depending on whether women want to raise or lower the level. We might also raise costs substantially as women get near the boundaries on zero and certain fertility. More generally there are issues with the precise utility function we have used for our estimation. Different ways of formulating the utility of consumption, and disutility of work and childcare, may give different results. Marriage might be included as an additional choice variable but is complicated by the fact that it is a two-sided decision. We leave these issues to future research.

Appendix

Derivation of Euler equations

We begin by deriving the Euler equation for the general case. Bellman's equation gives the value of the current state variables (n_t, y_t, w_t, e_t) as :

$$V(n_t, y_t, w_t, e_t) = U(c_t, f_t, l_t, g_t, n_t, y_t, t) + \beta E_t \{V(n_{t+1}, y_{t+1}, w_{t+1}, e_{t+1})\} \qquad \text{(A1)}$$

Subject to the equations of motion:

$$w_{t+1} = r_{t+1}(w_t + l_t p_t(e_t) - c_t) \qquad \text{(A2)}$$

$$n_{t+1} = n_t + F_t \qquad \text{(A3)}$$

$$y_{t+1} = F_t \qquad \text{(A4)}$$

$$e_{t+1} = e_t + l_t \qquad \text{(A5)}$$

where variables may be stochastic. The first-order conditions in the choice variables are:

$$\begin{aligned} &U_{ct} - \beta E_t(V_{wt+1} r_{t+1}) = 0 && \text{(a)} \\ &U_{ft} + \beta E_t(V_{nt+1}) + \beta E_t(V_{yt+1}) = 0 && \text{(b)} \\ &U_{lt} + p_t \beta E_t(V_{wt+1} r_{t+1}) + \beta E_t(V_{et+1}) = 0 && \text{(c)} \end{aligned} \qquad \text{(A6)}$$

Where we use the fact that $E_t(F_t) = f_t$. The envelope conditions are:

$$\begin{aligned} &V_{nt} = U_{nt} + \beta E_t(V_{nt+1}) && \text{(a)} \\ &V_{wt} = \beta E_t(V_{wt+1} r_{t+1}) && \text{(b)} \\ &V_{et} = \beta E_t(V_{et+1}) + l_t \frac{\delta p_t}{\delta e_t} \beta E_t(V_{wt+1} r_{t+1}) && \text{(c)} \\ &V_{yt} = U_{yt} && \text{(d)} \end{aligned} \qquad \text{(A7)}$$

The consumption Euler equation is obtained by substituting A7(b) into A6(a), leading, multiplying by βr_{t+1}, taking expectations at time t-1 and substituting back into A6(a) to obtain:

$$U_{ct-1} = \beta E_{t-1}(U_{ct} r_t) \qquad \text{(A8)}$$

This is equation (7) in main text.

The labour supply Euler equation is obtained from A6(a), A6(c), and A7(a). Substituting A6(a) into A6(c), we obtain $U_{lt} + p_t U_{ct} + \beta E_t(V_{et+1}) = 0$. Using this, substitute for $\beta E_t(V_{et+1})$

in A7(c), use A6(a) and lead to obtain $V_{et+1} = -U_{lt+1} - p_{t+1}U_{ct+1} + l_{t+1}\frac{\delta p_{t+1}}{\delta e_{t+1}}U_{ct+1}$. Taking expectations at time t and using $U_{lt} + p_t U_{ct} + \beta E_t\left(V_{et+1}\right) = 0$, we obtain $U_{lt} + p_t U_{ct} + \beta E_t\left[-U_{lt+1} - p_{t+1}U_{ct+1} + l_{t+1}\frac{\delta p_{t+1}}{\delta e_{t+1}}U_{ct+1}\right] = 0$. Simplifying, and leading by one period we have our labour supply Euler equation:

$$U_{lt-1} + p_{t-1}U_{ct-1} = \beta E_{t-1}\left[U_{lt} + p_t U_{ct} - l_t\frac{\delta p_t}{\delta e_t}U_{ct}\right] \qquad \text{(A9)}$$

This is equation (8) in main text.

The fertility Euler equation is obtained from A6(b), A7(a) and A7(d) . Substitute A7(d) into A6(b) and substitute the modified A6(b) into A7(a) to obtain $V_{nt} = U_{nt} - U_{ft} - \beta E_t\left(U_{yt+1}\right)$. In the derivation of the labour supply Euler equation above, we found $\beta E_t\left(V_{et+1}\right) = -\left(U_{lt} + \tilde{p}_t U_{ct}\right)$. Substitute this. Then multiply by β, lead and take expectations from time t. Substitute the resulting expression and A7(d) into A6(b) to obtain

$$U_{ft-1} + \beta E_{t-1}U_{nt} + \beta E_{t-1}U_{yt} = \beta E_{t-1}U_{ft} + \beta^2 E_{t-1}\left(U_{yt+1}\right) \qquad \text{(A10)}$$

This is equation (9) in main text.

We now turn to the derivations of the explicit Euler equations we estimate. From the Mincer equation for wages we have:

$$\log p_t = \log p^* + \gamma e_t - \frac{\delta}{2}e_t^{\,2} \Rightarrow \frac{\delta p_t}{\delta e_t} = \left(\gamma - \delta e_t\right)p_t \qquad \text{(A11)}$$

From the utility function we can derive the marginal utilities

$$U\left(c_t, f_t, l_t, n_t, y_t\right) = e^{\rho g_t}e^{\pi n_t}\log\left(1+c_t\right) - \frac{\omega_l}{2}\left(l_t + \phi y_t\right)^2 + \alpha n_t - \frac{\omega_f}{2}\left(f_t - \lambda_t\right)^2$$

$\Rightarrow$

$$\begin{aligned}
U_{ct} &= \frac{e^{\rho g_t}e^{\pi n_t}}{1+c_t} \\
U_{lt} &= -\omega_l\left(l_t + \phi y_t\right) \\
U_{ft} &= -\omega_f\left(f_t - \lambda_t\right) \\
U_{yt} &= -\omega_l\phi\left(l_t + \phi y_t\right) \\
U_{nt} &= \pi e^{\rho g_t}e^{\pi n_t}\log\left(1+c_t\right) + \alpha
\end{aligned} \qquad \text{(A13)}$$

We now derive our explicit Euler equations.

Consumption Euler: $U_{ct-1} = \beta E_{t-1}\left(U_{ct} r_t\right)$

Substituting in the derivatives for the utility function gives us

$$\frac{e^{\rho g_{t-1}} e^{\pi n_{t-1}}}{1+c_{t-1}} = \beta E_{t-1}\left(r_t \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right) \tag{A14}$$

Now define the shock to the marginal utility of consumption at time t as

$$v_{ct} = \frac{\left(r_t \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right)}{E_{t-1}\left(r_t \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right)} - 1 \tag{A15}$$

Clearly $E_{t-1} v_{ct} = 0$. Now we can write the Euler Equation as

$$\frac{e^{\rho g_{t-1}} e^{\pi n_{t-1}}}{1+c_{t-1}} = \beta \left(r_t \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right) \frac{1}{1+v_{ct}} \tag{A16}$$

and taking $\varepsilon_{ct} = -\log(1+v_{ct})$ we have

$$\log\left(\frac{1+c_t}{1+c_{t-1}} \right) = \log\beta + \log r_t + \rho\left(g_t - g_{t-1}\right) + \pi\left(n_t - n_{t-1}\right) + \varepsilon_{ct} \tag{A17}$$

Now provided v_{ct} is small we have $\varepsilon_{ct} = -\log(1+v_{ct}) \approx -v_{ct}$ and $E_{t-1}\varepsilon_{ct} \approx 0$.

This approach to log linearising the utility Euler equation depends on the shocks that affect the marginal utility of consumption being small.

We now take the labour Euler Equation

$$\textit{Labour}: \quad U_{lt-1} + p_{t-1} U_{ct-1} = \beta E_{t-1}\left[\left(U_{lt} + p_t U_{ct}\right) - l_t \frac{\delta p_t}{\delta e_t} U_{ct} \right]$$

Substituting in the derivatives from our explicit functional forms we have

$$\begin{aligned} &-\omega_l\left(l_{t-1} + \phi y_{t-1}\right) + p_{t-1} \frac{e^{\rho g_{t-1}} e^{\pi n_{t-1}}}{1+c_{t-1}} \\ &= \beta E_{t-1}\left[\left(-\omega_l\left(l_t + \phi y_t\right) + p_t \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right) - l_t p_t\left(\gamma - \delta e_t\right) \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right] \end{aligned} \tag{A18}$$

Now noting that the actual future outcome is the expected outcome plus a shock we have

$$\begin{aligned} &-\omega_l\left(l_{t-1} + \phi y_{t-1}\right) + p_{t-1} \frac{e^{\rho g_{t-1}} e^{\pi n_{t-1}}}{1+c_{t-1}} \\ &= \beta\left[-\omega_l\left(l_t + \phi y_t\right) + p_t \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} - l_t p_t\left(\gamma - \delta e_t\right) \frac{e^{\rho g_t} e^{\pi n_t}}{1+c_t} \right] + \varepsilon_{lt} \end{aligned} \tag{A19}$$

Where $E_{t-1}\varepsilon_{lt} = 0$.

Finally we consider the fertility Euler equation.

$$Fertility: \quad U_{ft-1} + \beta E_{t-1}U_{nt} + \beta E_{t-1}U_{yt} = \beta E_{t-1}U_{ft} + \beta^2 E_{t-1}\left(U_{yt+1}\right)$$

Again substituting in the derivatives from our explicit function forms we have

$$\begin{aligned} &-\omega_f\left(f_{t-1} - \lambda_{t-1}\right) + E_{t-1}\left[\beta\alpha + \beta\pi e^{\rho g_t} e^{\pi n_t}\log\left(1+c_t\right) - \beta\omega_l\phi\left(l_t + \phi y_t\right)\right] \\ &= E_{t-1}\left[-\beta\omega_f\left(f_t - \lambda_t\right) - \beta^2\omega_l\phi\left(l_{t+1} + \phi y_{t+1}\right)\right] \end{aligned} \tag{A20}$$

And again noting that the future outcomes are the expected outcomes plus a shock we have

$$\begin{aligned} &-\omega_f\left(f_{t-1} - \lambda_{t-1}\right) + \beta\alpha + \beta\pi e^{\rho g_t} e^{\pi n_t}\log\left(1+c_t\right) - \beta\omega_l\phi\left(l_t + \phi y_t\right) \\ &= -\beta\omega_f\left(f_t - \lambda_t\right) - \beta^2\omega_l\phi\left(l_{t+1} + \phi y_{t+1}\right) + \varepsilon_{ft} \end{aligned} \tag{A21}$$

Where $E_{t-1}\varepsilon_{ft} = 0$.

Table 1 Descriptive statistics (mean/standard deviation)

	c_t	r_t	g_t	n_t	l_t	p_t	f_t	Observations
1968	22,755	1.02	2.7	2.2	1,569	4.18	0.10	3,447
	16,292	-	1.2	2.2	885	1.75	0.29	
1969	22,825	1.02	2.7	2.0	1,516	4.40	0.12	3,330
	15,146	-	1.2	2.1	903	1.77	0.32	
1970	23,152	1.02	2.6	1.7	1,399	4.63	0.13	3,451
	14,701	-	1.2	1.9	963	1.82	0.34	
1971	23,434	1.02	2.6	1.6	1,281	4.82	0.14	3,244
	22,178	-	1.2	1.7	989	1.81	0.34	
1972	23,115	1.02	2.5	1.4	1,222	4.96	0.11	3,234
	20,049	-	1.2	1.6	1,005	1.87	0.32	
1973	20,207	1.02	2.3	1.4	1,215	5.10	0.11	2,453
	13,892	-	1.1	1.5	1,001	1.94	0.21	
1975	19,479	0.99	2.0	1.4	1,167	5.52	0.10	2,432
	13,000	-	0.8	1.4	1,004	2.16	0.20	
1977	20,720	1.01	1.9	1.5	1,172	5.79	0.12	2,469
	12,475	-	0.8	1.3	991	2.34	0.32	
1978	22,054	1.01	1.9	1.6	1,155	5.75	0.08	2,185
	16,250	-	0.8	1.4	996	2.40	0.19	
1980	21,713	1.03	2.0	1.8	1,151	5.80	0.07	1,815
	13,806	-	0.7	1.2	985	2.44	0.18	
1982	21,160	1.19	1.9	1.8	1,266	6.14	0.07	2,231
	12,656	-	0.8	1.3	981	2.69	0.26	
1983	23,885	1.08	2.0	1.8	1,270	6.14	0.04	1,595
	17,413	-	0.8	1.3	985	2.68	0.14	
1985	25,554	1.14	2.0	1.6	1,416	6.60	0.03	1,742
	17,796	-	0.8	1.2	954	2.99	0.13	
1987	28,107	1.11	2.0	1.6	1,489	6.79	0.03	1,844
	18,994	-	0.8	1.3	933	3.08	0.17	
1988	30,499	1.05	2.2	1.4	1,398	6.96	0.01	1,477
	25,536	-	0.9	1.2	961	3.23	0.07	
1991	32,676	1.16	2.2	1.1	1,411	7.29	0.01	1,283
	25,058	-	0.9	1.2	962	3.44	0.06	
1993	31,446	1.06	1.6	0.8	1,641	7.60	0.01	766
	30,328	-	0.8	1.1	909	3.52	0.06	
1995	42,568	1.10	2.1	0.6	1,351	7.79	0.00	674
	47,118	-	1.1	1.0	1,012	3.73	0.05	
1997	39,570	1.12	2.0	0.4	1,346	7.57	0.03	659
	46,607	-	0.8	0.8	1,016	3.66	0.12	
1999	45,200	1.13	2.4	0.4	1,278	7.50	0.00	611
	61,663	-	1.2	0.8	1,015	3.53	0.03	
2001	54,485	1.10	2.5	0.2	1,179	7.92	0.00	632
	86,723	-	1.3	0.5	1,018	3.76	0.03	

Note : c_t, annual real family consumption (1983 prices) ; r_t , real rate of interest ; g_t, number of adults in family ; n_t, number of children in family ; l_t, woman's annual labour hours ; p_t, hourly real rate of pay (1983 prices) ; f_t, fertility.

Table 2 Mincer wage equation estimates

Parameter	Explanatory variable	Coefficients	
		Experience in years	Experience in hours
ξ	Years of schooling	0. 116 ** (0.002)	0. 116 ** (0.002)
γ	Experience	0.044** (0.001)	0.220 x10^{-4}** (0.504 x10^{-6})
$-\delta/2$	Experience^2	-0.00062** (0.00002)	-0.156 x10^{-09}** (0.603x10^{-11})
	Time trend	-0.0078** (0.0004)	-0.0078** (0.0004)
	Observations Number of women R-squared	53,011 4933 0.641	53,011 4933 0.641

Note: Both regressions include woman fixed effects. Standard errors in brackets. ** significant at 1%.

Table 3 Estimated parameters from Euler equations

	Parameter	Description	Estimate
Consumption Euler	β	Discount rate	0.984** (0.003)
	ρ	Effect of an adult on household consumption	0.202** (0.013)
	π	Effect of a child on household consumption	0.064** (0.008)
Labour Euler	ϕ	Annual hours of childcare per child	681 (829)
	ω_l	Weight on labour supply in utility	0.142×10^{-6}* ($.070 \times 10^{-6}$)
Fertility Euler	$\chi_f = 1/\omega_f$	Inverse of weight on fertility on utility	0.212 ** (0.012)
	α	Weight on children in utility function	0.236** (0.014)
	λ_t	Natural age-specific fertility rates	See figure 6
		Observations (consumption) Observations (labour) Observations (fertility)	21,850 21,850 21,698

Note: This table gives results from the consumption Euler equation (12), the labour Euler equation (13) and fertility Euler equation (17). In all estimates the values of variables other than fertility measured in period t or later are all instrumented with lags measured in periods t-1 and t-2. Observed fertility at time t-1and onwards is also instrumented (with fertility at t-2 and t-3) since it is not known when decisions at t-1 are being made. The consumption Euler is estimated first. The parameter values from the consumption Euler are fixed in the estimation of the labour supply Euler and both these sets of parameters are held fixed when estimating the fertility Euler.

** significant at 1% level. * significant at 5% level.

Figure 1 Sequence of events in each time period

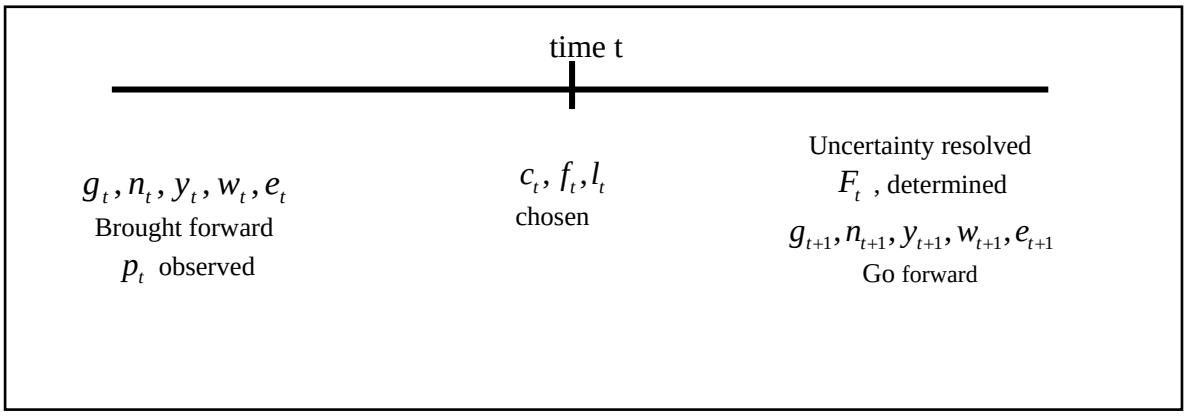

Figure 2 Age-specific fertility rates for women with low and high levels of education

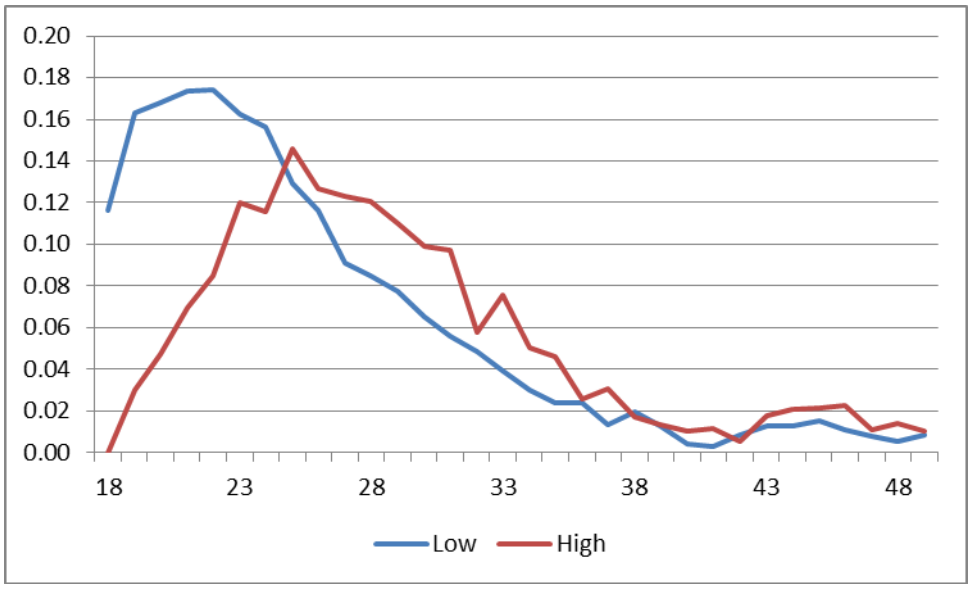

Note: High = women with more than 12 years of education.

Figure 3 Average labour supply by age for women with low and high levels of education

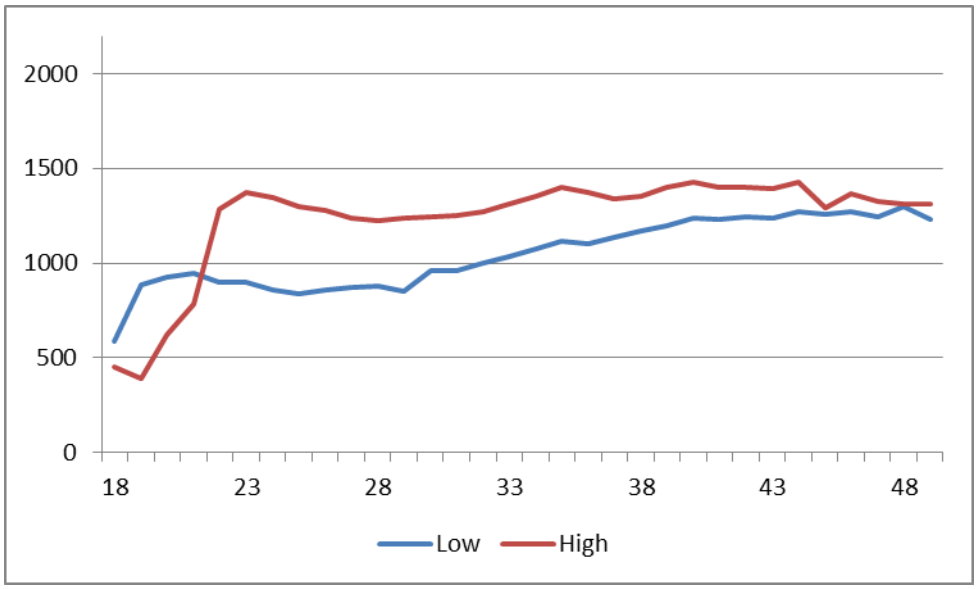

Note: Figures are hours per year. High = women with more than 12 years of education.

Figure 4 Average labour supply by age for women with low and high levels of education, including hours in education

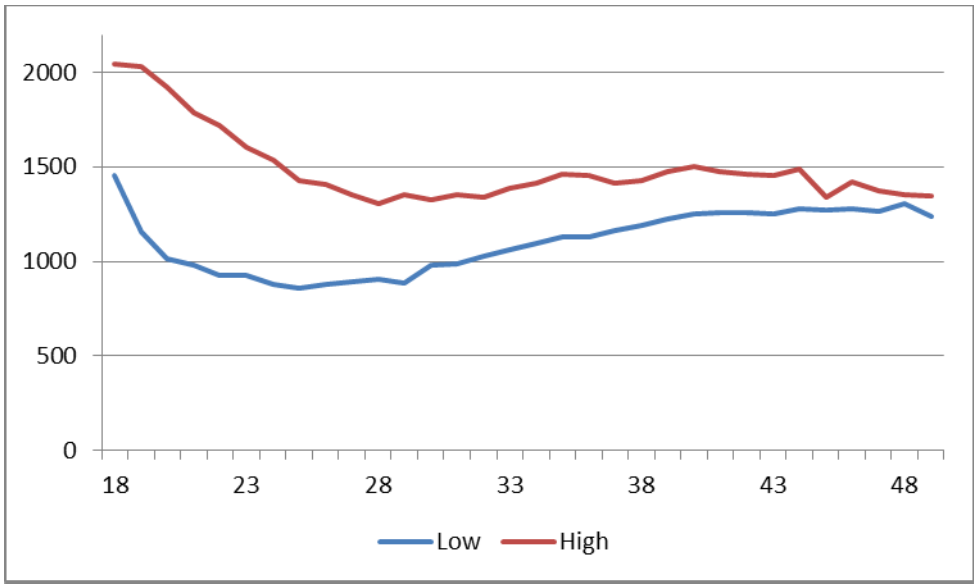

Note: Figures are hours per year. High = women with more than 12 years of education. It is assumed that full-time education is equivalent to an annual labour supply of 2,000 hours (40 hours/week).

Figure 5 Average log real household consumption by age for women with low and high levels of education

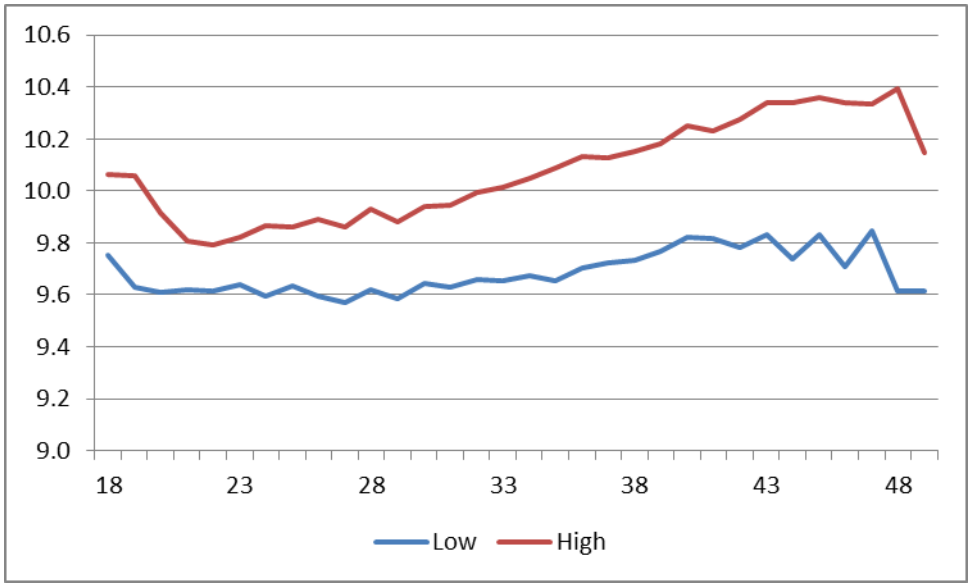

Note: Figures are the natural logarithm of annual household consumption at 1983 prices. High = women with more than 12 years of education.

Figure 6 Estimated natural fertility rate by age

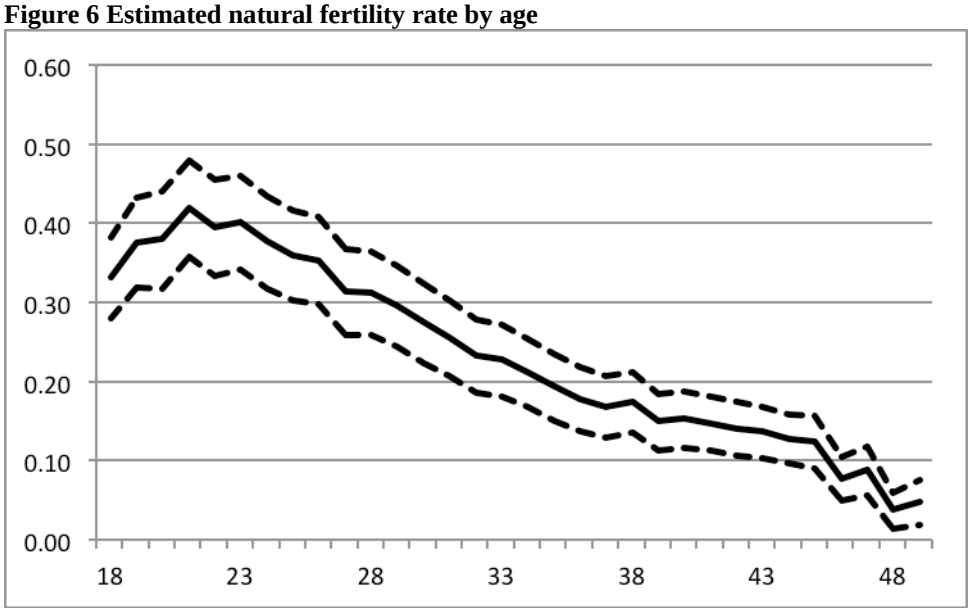

Note: Dotted lines give 95% confidence intervals.

References

Ahn, N. (1995). 'Measuring the value of children by sex and age using a dynamic programming model', *The Review of Economic Studies,* vol. 62(3), pp. 361-379.

Arroyo, C.R. & Zhang, J. (1997). 'Dynamic microeconomic models of fertility choice: A survey', *Journal of Population Economics,* vol. 10(1), pp. 23-65.

Attanasio, O.P. & Low, H. (2004). 'Estimating Euler equations', *Review of Economic Dynamics,* vol. 7(2), pp. 406-435.

Attanasio, O., Low, H. & Sanchez-Marcos, V. (2008). 'Explaining Changes in Female Labor Supply in a Life-Cycle Model', *American Economic Review,* vol. 98(4), pp. 1517-1552.

Bloemen, H. & Kalwij, A.S. (2001). 'Female labor market transitions and the timing of births: a simultaneous analysis of the effects of schooling', *Labour Economics,* vol. 8(5), pp. 593-620.

Bongaarts, J. (1978). 'A Framework for Analyzing the Proximate Determinants of Fertility', *Population and Development Review,* vol. 4(1), pp. pp. 105-132.

Caucutt, E.M., Guner, N. & Knowles, J. (2002). 'Why Do Women Wait? Matching, Wage Inequality, and the Incentives for Fertility Delay', *Review of Economic Dynamics,* vol. 5(4), pp. 815-855.

Collins, J. (2002). 'An international survey of the health economics of IVF and ICSI', *Human Reproduction Update,* vol. 8(3), pp. 265-277.

Craig, L. & Bittman, M. (2008). 'The incremental time costs of children: An analysis of children's impact on adult time use in Australia', *Feminist Economics,* vol. 14(2), pp. 59-88.

Del Boca, D. & Sauer, R.M. (2009). 'Life cycle employment and fertility across institutional environments', *European Economic Review,* vol. 53(3), pp. 274-292.

Eckstein, Z. & Lifshitz, O. (2011). 'Dynamic female labor supply', *Econometrica,* vol. 79(6), pp. 1675-1726.

Eckstein, Z. & Wolpin, K.I. (1989). 'Dynamic labour force participation of married women and endogenous work experience', *The Review of Economic Studies,* vol. 56(3), pp. 375-390.

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

Hansen, L.P. (1982). 'Large sample properties of generalized method of moments estimators', *Econometrica: Journal of the Econometric Society,* vol.50, pp. 1029-1054.

Happel, S.K., Hill, J.K. & Low, S.A. (1984). 'An Economic Analysis of the Timing of Childbirth' , *Population Studies,* vol. 38(2), pp. 299-311.

Henshaw, S.K. & Finer, L.B. (2003). 'The accessibility of abortion services in the United States, 2001', *Perspectives on Sexual and Reproductive Health,* vol. 35(1), pp. 16-24.

Jones, L.E., Schoonbroodt, A. & Tertilt, M. (2011). 'Fertility theories' *in Demography and the Economy,* ed. J.B. Shoven, NBER and University of Chicago Press.

Knodel, J. (1978). 'Natural fertility in pre-industrial Germany', *Population Studies,* vol. 32(3), pp. 481-510.

Moffitt, R. (1984). 'Profiles of fertility, labour supply and wages of married women: a complete life-cycle model', *The Review of Economic Studies,* vol. 51, pp. 263-278.

Montgomery, M. (1989). 'Dynamic behavioural models and contraceptive choice', *Journal of Biosocial Science,* vol. 11, pp. 17-40.

Newman, J.L. (1988). 'A stochastic dynamic model of fertility.', *Research in Population Economics,* vol. 6, pp. 41-68.

Olivetti, C. (2006). 'Changes in women's hours of market work: The role of returns to experience', *Review of Economic Dynamics,* vol. 9(4), pp. 557-587.

Pagnini, D.L. & Rindfuss, R.R. (1993). 'The divorce of marriage and childbearing: Changing attitudes and behavior in the United States', *Population and Development Review,* vol.19(2), pp. 331-347.

Sheran, M. (2007). 'The career and family choices of women: A dynamic analysis of labor force participation, schooling, marriage, and fertility decisions', *Review of Economic Dynamics,* vol. 10(3), pp. 367-399.

Smith, K.E., Downs, B. & O'Connell, M. (2001). *Maternity Leave and Employment Patterns, 1961-1995,* US Department of Commerce, Economics and Statistics Administration, US Census Bureau.

Trussell, J., Lalla, A.M., Doan, Q.V., Reyes, E., Pinto, L. & Gricar, J. (2009). 'Cost effectiveness of contraceptives in the United States', *Contraception,* vol. 79(1), pp. 5-14.

Waldfogel, J. (1998). 'Understanding the 'Family Gap' in Pay for Women with Children', *The Journal of Economic Perspectives,* vol. 12(1), pp. 137-156.

Wilde, E.T., Batchelder, L. & Ellwood, D.T. (2010). 'The Mommy Track Divides: The Impact of Childbearing on Wages of Women of Differing Skill Levels', NBER no.16582.

Wolpin, K.I. (1984). 'An estimable dynamic stochastic model of fertility and child mortality', *The Journal of Political Economy,* vol.92(5), pp. 852-874.