

Does Fertility Respond to Financial Incentives?

Guy Laroque Bernard Salanié¹

February 2005

¹CREST-INSEE, GRECSTA and CEPR. We benefited from comments by Gary Becker, Didier Blanchet, Olivia Ekert-Jaffé, Philippe Février, Dominique Goux, James Heckman, Mark Killingsworth, Anne Laferrère, Steve Levitt, Annalisa Luporini, Thierry Magnac, Kevin Murphy, Laurent Toulemon, Alain Trannoy, Daniel Verger and James Walker.

Abstract

We attempt here to evaluate the sensitivity of fertility to financial incentives in France. We discuss and implement several estimation strategies; our main focus is on a structural model of female participation and fertility based on a microsimulation model of the tax-benefit system. We estimate this model on individual data from the Labor Force Survey. Our results suggest that financial incentives play a sizable role in determining fertility decisions in France.

JEL codes: J13, J22, H53

Keywords: population, fertility, incentives, benefits.

1 Introduction

It is natural for economists raised on the *New Home Economics* to presume the existence of a link between family transfers and fertility. The standard model of Becker (1960, 1991) implies that the demand for children depends on their cost, which in turn depends on family transfers (see for instance Cigno (1986) for a theoretical study of the impact of taxes on fertility). In fact, the policies towards families implemented in France at the end of the 1930's were in part based on the belief that family benefits increase fertility. Sweden and Québec adopted a similar approach to family benefits. In a number of other countries, this belief is not shared (see Gauthier (1996)), and family benefits are mostly designed as a way to ensure a minimum standard of living to families and children. Germany for instance, while worried by its fertility decline, has not used family benefits in purportedly pronatalist ways. Even in France, decision makers seem to have mixed feelings: while a recent widely quoted administrative report by Thélot and Villac (1998) barely mentions fertility in its analysis of family policies, a 2004 reform of family benefits explicitly mentions fertility concerns as a motivation. Our estimation results indeed suggest that family benefits, by reducing the cost of raising children, can have sizable effects on fertility. We estimate the cost elasticity of the demand for children to be about 0.2. To give a rough order of magnitude, politically feasible benefit reforms can change the cost of children by about 25%; thus according to our estimates, such reforms can move fertility up or down by about 5%.

How can we quantify the effects of financial incentives on fertility? To simplify the discussion, consider a barebones model of fertility with the variable of interest generated according to:

$$F = f(X, E, u) \tag{1}$$

where F is some measure of the fertility of an observational unit, E measures the (current) financial incentives to fertility, X comprises sociological, demographic, and all other predetermined or weakly exogenous determinants of fertility, and u is an error term. Empirical work aims at recovering an estimate of f'_E .

As a very simple example, F could be the occurrence of a birth for a given couple in a given period; if we had some exogenous measure of the variation in net discounted household income ΔR linked to a birth, we could posit a binomial model for F based on the latent variable

$$F^* = X\beta + \alpha\Delta R + u,$$

and our main interest would be in the estimate of α . A major issue, which arises even in this simple example, is in the measure of ΔR ; in practice, the income variation consecutive to a birth, “the price of a birth”, depends heavily on whether the woman stops working or not. Thus the simple model is not good enough: the fertility decision depends on the whole set of incomes conditional on working or not working, on having a child or not having a child. We should include in E the list of values of income for each possible joint choice of fertility and hours (in the language of a demand model, demand depends on the prices of all available goods). The value of current income when working depends on the wage W , so that E is a (hopefully) known function

$$E = R(X^E, W; \pi) \quad (2)$$

where R describes the whole tax-benefit system, whose parameters we denote π for future reference (so that π includes for instance the brackets of the income tax, the value of child credits, and so forth). The identification of f'_E in (1) comes from variation in X^E or W that is not contained in X , or from some functional form restrictions (for instance, the tax-benefit system may exhibit discontinuities, while f may be assumed to be smooth)¹.

If the wage was observed for every individual (and was uncorrelated with u), we could stop there. Unfortunately, wages are only observed for working individuals. Thus as is usual in labor supply models, we need to specify a wage equation

$$W = w(X^W, \varepsilon), \quad (3)$$

which is typically estimated jointly with the choice of hours.

The estimation problem then boils down to estimating f'_E (which immediately gives f'_π since the tax-benefit system R is known), given that we observe the exogenous variables (X, X^E, X^W) , the choice variable F , and W when the individual is working. Several approaches have been used (see for instance the survey of Hotz, Klerman, and Willis (1997)).

A first group of studies uses panel data on countries; then the unit of observation is a given (country, year) (*cy*) pair. Thus the work of Ekert-Jaffé (1986), Blanchet and Ekert-Jaffé (1994) and Gauthier and Hatzius (1997) suggests that the French family benefit system may increase total fertility by 0.1 to 0.2 child per woman. This line of work can be justified by packing our

¹The fertility decision involves predicting the levels of income and family benefits in the future. In model (1), these predictions are implicit, and the derivative f'_E captures the direct effect of a change in the current financial incentives, as well as the indirect effect through the induced change in the expectations of future income and benefits. Typically identification then is more difficult than in a static world. We briefly discuss below structural approaches to a fully dynamic model.

equations together, so that

$$F = f(X, R(X^E, w(X^W, \varepsilon); \pi), u)$$

Assuming, rather heroically, that this functional form can be well approximated by a linear form, we can estimate the regression

$$F_{cy} = X_{cy}\beta + X_{cy}^E\beta^E + X_{cy}^W\beta^W + \pi_{cy}\alpha + \xi_c + \xi_y + \xi_{cy} \quad (4)$$

and α is identified under the reasonable assumption that there is variation in policies over time and space that is not explained entirely by the other explanatory variables.

The first glaring difficulty with equation (4) is that the list of variables in (X, X^W, X^E) should be very long. As we will see later, X should for instance contain not only the age of the woman, but also the current age breakdown of children, and some interaction terms; X^W must contain schooling and experience variables; and X^E must list all variables that enter the computation of taxes and benefits. Instead of using measures of all of these (some of which may be hard to recover at the (country, year) level), these authors have used very restricted specifications. This is only valid, of course, if the time and space variation in the omitted variables is orthogonal to the variation in policy parameters. Moreover, these policy parameters are summed up by a couple of measures only; as shown in section 2.1, this can only be a very coarse summary of a complex system, which raises a thorny measurement problem². Thus the results obtained in these papers seem fragile.

Most recent work in fact has used individual data, which alone can provide the analyst with a large number of observations and with enough information to measure precisely financial incentives to fertility. Yet the first criticism given above extends to papers that use individual panel data but control for very few covariates, as in Whittington (1992). Other papers by Rosenzweig and Schultz (1985), Hotz and Miller (1988), Heckman and Walker (1989, 1990a, 1990b) confirm that as predicted by theory, fertility decreases with the woman's potential wage and increases with the other income of the household. The estimated effects are small, however³; as these papers make no attempt to use the form of the tax-benefit function R , it is not quite clear what the estimated coefficients of wages represent in any case.

²Ideally, one should also control for the endogeneity of policy variables, as argued by Heckman (1976) and more recently by Besley and Case (2000).

³Many papers have studied the effects of the American welfare system on fertility (see Moffitt (1998)); the results of Rosenzweig (1999) thus suggest that in the 70s, social transfers like the AFDC had a large effect on the probability that a young lower-class woman would become a single mother. Blau and Robins (1989) estimate the effect of child-care costs on fertility, using geographical variation in child-care availability.

Several recent studies have used the natural experiments approach to evaluate the fertility impact of a given family benefit. Thus Milligan (2004) studies the effect of a cash benefit given on the birth of a child in Québec in the 1990s; it finds that this benefit strongly stimulated fertility. His approach can be described by starting from (assuming linearity again, and letting R be one-dimensional for simplicity)

$$F_{it} = X_{it}\beta + \alpha R(X_{it}^E, w(X_{it}^W, \varepsilon_{it}); \pi_{it}) + u_t + u_{it}$$

where i indexes a woman and t is 0 (before the reform) or 1 (after the reform)⁴. If the child birth cash benefit $B(X_{i1})$ is not taxable and is the only policy change over this time period, then we can write the change in economic incentives over time as

$$\delta e_i = \delta_i^0 + B(X_{i1})\mathbf{1}(i \in Q),$$

where Q denotes Québec (as opposed to other Canadian provinces) and

$$\delta_i^0 = R(X_{i1}^E, w(X_{i1}^W, \varepsilon_{i1}); \pi_{i0}) - R(X_{i0}^E, w(X_{i0}^W, \varepsilon_{i0}); \pi_{i0})$$

is the change over time in financial incentives due to changes in wages and other relevant characteristics. Thus we could estimate in principle

$$F_{it} = (X_{it}\beta + \alpha R(X_{i0}^E, w(X_{i0}^W, \varepsilon_{i0}); \pi_{i0})) + \alpha \delta_i^0 \mathbf{1}(t = 1) + \alpha B(X_{i1})\mathbf{1}(t = 1)\mathbf{1}(i \in Q) + u_t + u_{it}, \quad (5)$$

where as usual the important coefficient α has for covariate an indicator that the observation is in Québec after the reform. Here the problem is that Milligan (2004) instead estimates

$$F_{it} = X_{it}\beta + \alpha B(X_{i1})\mathbf{1}(t = 1)\mathbf{1}(i \in Q) + \gamma \mathbf{1}(i \in Q) + u_t + u_{it}.$$

As emphasized in a similar context by Rosenzweig and Wolpin (2000), it appears from (5) that this strategy can only be valid if the variables omitted from Milligan's specification are orthogonal to the double indicator variable, that is if their mean is the same in Québec in $t = 1$ than in the pooled data. Since these variables include all variables that are in X^F but not in X (for instance the ages of children⁵) along with all variables that influence wages and the tax-benefit system, this is at least a debatable assumption. Similar

⁴We do not introduce an individual fixed effect as Milligan uses two cross-sections of data.

⁵Milligan only differentiates children according to whether they are older than 6.

criticisms apply to Kearney (2004), who estimates the fertility effect of the introduction of family caps in many US states following the 1996 reform of welfare. Of course, a rich enough dataset would allow the authors to control for the missing variables.

An alternative approach, which we pursue in this paper, is to rely on a structural model. This leads to a system of equations similar to (1), (2), (3) to which we add a labor supply equation

$$L = g(X^L, E, v), \quad (6)$$

along with functional forms for f , g and w , and distribution choices for the error terms. This method has several advantages. By specifying each equation of the model, we are less likely to omit relevant variables. Moreover, by construction it makes use of the full structure of the tax-benefit system through the known function R . Of course, it also has its drawbacks. One of them is that the resulting coefficient estimates are only as credible as the functional and distributional specifications they rely on. This comment applies very generally; but this particular model raises two more specific issues.

The first one is that the omitted variables problem is particularly galling for fertility, since we know little about its determinants. Even if we use a long list for the variables X , and even in a linear model, the estimate of f'_E will be biased unless all omitted variables are uncorrelated with the financial incentives E , conditionally on the included variables X . In practice, E depends on X^E , and therefore on many variables linked to past fertility; so if there is any fixed fertility effect (call it “fecundity”) that is not properly accounted by X , the estimator of f'_E may be biased. We have no way to test the uncorrelatedness assumption⁶; we can only mitigate the effect of omitted variable bias by using many variables in X . We will come back to this point when describing our specification.

Second, several recent papers, following Wolpin (1984), have used dynamic programming techniques to model fertility choices. Thus, Francesconi (2002) estimates a dynamic model of fertility and participation choices; Keane and Wolpin (2002a) simulate a dynamic model of fertility and welfare take-up. However, Francesconi (2002) sums up family benefits by transfers that are linear in the number of children. Thus his paper relies on a very coarse measure of the E variable, which is taken to be proportional to the number

⁶Existing evidence is mixed. Heckman and Walker (1989) find evidence for a strong unobserved heterogeneity component in fertility behavior; on the other hand, Heckman and Walker (1990a) describe unobserved fecundity as “empirically unimportant”. The latter paper uses more explanatory variables, but much fewer than we do in this paper.

of children. But the real tax-benefit system sets up complex interactions between household composition and income: in France for instance, many social and family transfers depend non-linearly on these two classes of variables. Moreover, the ages of children play a large role in determining benefits: ideally, a good measure of E should take into account the consequences of a birth on income flows during all years before the child comes of age. Such a dynamic simulation is made even more complicated by the fact that future income depends on women's participation decisions today, which in turn depend on household composition and thus on past births. Keane and Wolpin (2002a) do describe taxes and transfers much more realistically, but they only estimate a reduced form of their model in Keane and Wolpin (2002b).

We chose in this paper to emphasize a precise description of the tax-benefit system, along the lines of our previous work (Laroque and Salanié (2002)). We rely on a microsimulation model of the main taxes and benefits that approximates the complex R function reasonably well. The drawback of this approach is that this function R depends on a large array of X^E variables: the number of children, their ages, the wage of the partner (if any), the type of dwelling. . . If we were to plug this model into a dynamic model of fertility and participation, the vector of state variables would become impossibly large: we would need, for instance, the number of children for each age. None of the simulation and estimation techniques developed so far seems to be able to cope with such a state space.

We have tried to think of realistic ways of simplifying the function R so as to reduce the dimension of the state space, but complexity seems to be of the essence in this problem; our earlier work on female participation suggests that the detail of taxes and benefits matters a lot. We choose here a different tack by restricting the choice space of women. Thus we assume that when a woman at date t plans her fertility and labor supply behavior at date $(t+1)$, she assumes that (i) she will have no more children from $(t+1)$ on; and (ii) her employment status will stay unchanged after $(t+1)$. Even so, computing the discounted sum of incomes is still very costly, as the function R is complex; thus we simplify further by only discounting the family benefits (and not, e.g. the housing subsidies or the income tax). While these may seem to be very gross assumptions, we note that Francesconi (2002) finds that a fully dynamic model (with a very simplified R function) does not fit the data much better than a myopic model. It is indeed possible that the decision rules induced by these two models look similar—although their parameters clearly do not have the same interpretation.

We focus on French data not only because we know it well, but also because France has a rather generous and diversified family benefit system. Its cost is evaluated at about 0.8% of GDP, and it comprises some uncondi-

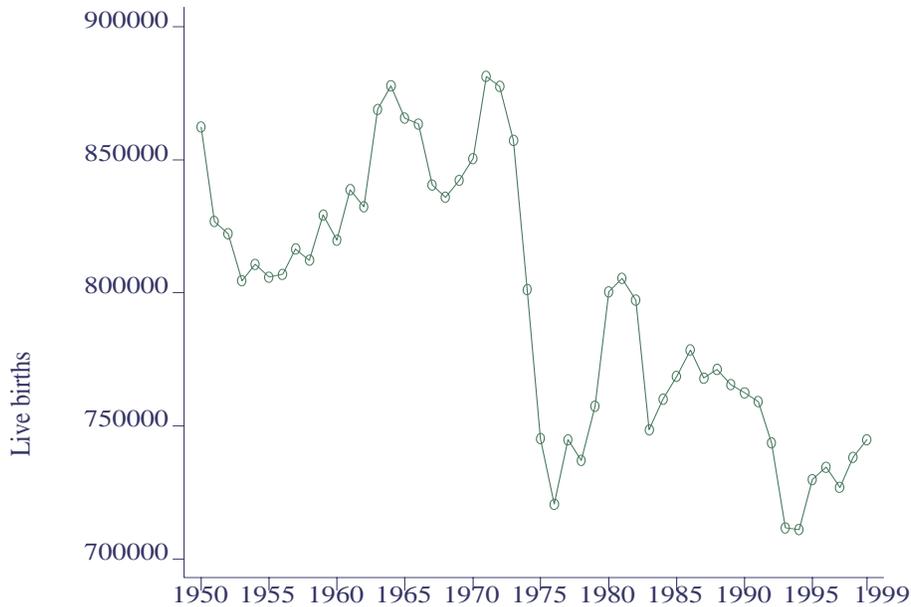


Figure 1: The Annual Number of Births in France from 1950 to 1999

tional, some means-tested, and some employment-tested benefits. In particular, the recent rebound in the number of births (see figure 1) has often been linked in the media with various family policy measures which date from the same period, particularly with the extension of the “Allocation Parentale d’Éducation” (APE) to the second-born in July 1994 (see the Appendix for more details).

We use the French Labor Force Surveys of 1997, 1998 and 1999 and our microsimulation model of taxes and benefits in order to estimate α . Our estimation also accounts for women’s participation decisions, given that they are so closely linked to their fertility choices⁷. In our model, every woman is characterized by her productivity, her disutility for work and her net utility for a new child. If her productivity is smaller than the minimum wage, the woman cannot take a job. Otherwise, she can take a job paid at her productivity if she wishes to do so. She then jointly decides on participation and fertility, depending on her individual characteristics.

Even within a myopic model, we have had to make simplifying assumptions. Thus we ignore part-time work, even though we know that the birth of a child often leads a woman to go from full-time to part-time. We also

⁷A paper by Lefebvre, Brouillette, and Felteau (1994) already estimated a nested logit of fertility and participation on individual Canadian data; their results suggest that family benefits have a non-negligible impact on fertility.

neglect child-care subsidies, which are important both in theory (see Apps and Rees (2001)) and in practice (del Boca (2002) and Choné, Leblanc, and Robert-Bobée (2002)), as taking them into account would require modeling the choice of child care mode.

The second section of this paper briefly describes the workings of the family-linked taxes and benefits in France. We illustrate them by microsimulations on some benchmark cases and on a sample of women drawn from the 1999 Labor Force Survey in section 2. Section 3 presents some reduced form evidence on the fertility effects of family benefits. Then section 4 presents our econometric model, and the main estimation results are shown in the following section. Section 6 concludes the paper by simulating the effects on fertility of various actual or possible policy reforms.

2 Fertility Incentives in France

Our objective here is to illustrate how the family composition affects disposable household income in France (the main taxes and benefits which we simulate are described briefly in the Appendix). To do this, we first look at some benchmark cases; we then turn to the distribution of financial fertility incentives in the population under study.

2.1 Benchmark Cases

We focus here on a married couple that lives in public housing and has 0, 1, 2 or 3 children. Each child is assumed to be between 3 and 6 years old. The woman earns by assumption an hourly wage equal to the minimum wage; we assume that she fulfills the eligibility conditions required to perceive the APE if she has a parity 2 or higher birth and she stops working. In case 1, the husband is unemployed; in case 2, he has a full time job that pays the median wage, or about 1,400 euros per month.

Let us start with case 1. Figure 2 compares the budget constraints of four households with respectively 0, 1, 2 and 3 children⁸. This figure shows the well-known inactivity trap: a half-time minimum wage job gives no net income gain to a couple that is eligible to the minimum income guarantee (the RMI, see Appendix), and the net gain from a full time minimum wage job is very small. When the woman does not work, the household's income consists in two-thirds RMI and one-third housing subsidy; both increase in value

⁸The graph plots the net disposable income of the household as a function of the weekly number of hours worked by the woman, after all taxes and benefits. The two vertical lines show full time and half time.

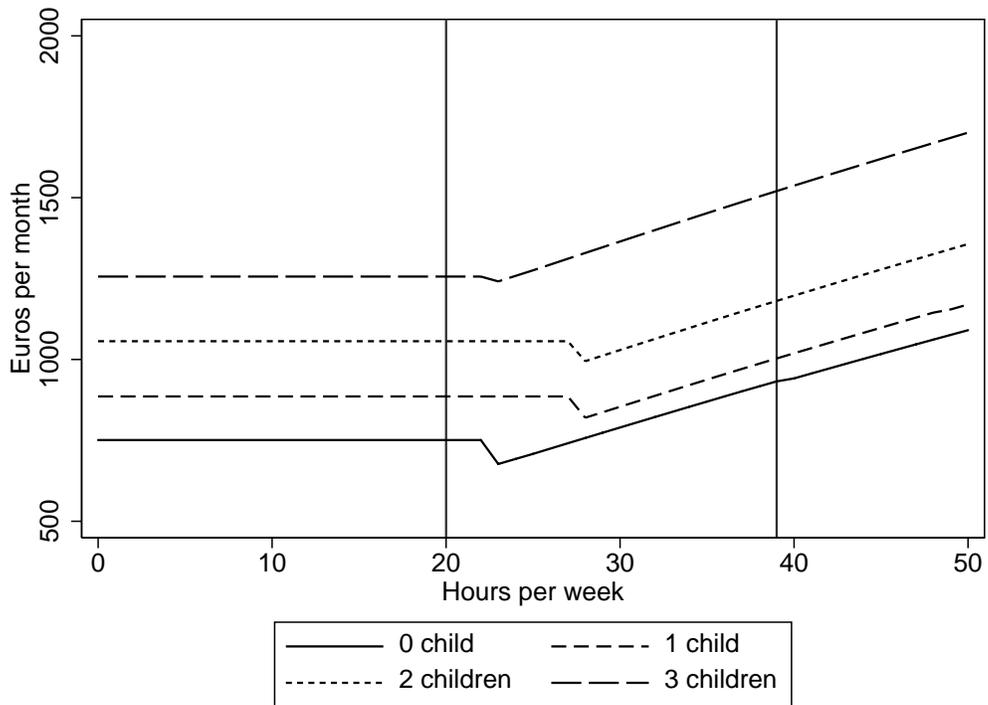


Figure 2: Case 1: net disposable income

Table 1: Income Gain with a Birth (case 1)

Birth	No job	Full time	Stop Working
Parity 1	135	220	-50
Parity 2	170	330	50
Parity 3	200	350	75

with the number of children. Moreover, this couple perceives the *allocations familiales* when it has at least two children, and the *complément familial* when it has three. When the woman works full time, the household does not get any RMI money any more; but its income is below the income tax exemption level, and it still perceives housing subsidies. The combination of these transfers has some probably unintended consequences; thus, the difference between the net disposable incomes of a 2-children and a 3-children household is much higher when the woman works than when she does not (340 vs 200 euros per month).

Still in case 1, let us now focus on the short-term financial effect of a birth. Table 1 shows the increase in net disposable income. In the first two columns, the woman does not change her hours⁹. If this woman does not work, a first birth gives her the APJE (150 euros per month) and a higher parity birth gives her the APE (465 euros per month)¹⁰; but these transfers are partly compensated by a decrease in perceived RMI, so that the net income gain is fairly small. If she works full time, the couple does not get any RMI and thus the income gain is much higher for births of parity 2 or more. The third column assumes that as is often the case, the woman stops working when giving birth. Then her household loses very little money for a first birth, and actually gains a little for a higher parity birth: the loss in wages (about 820 euros) is more than compensated by the increase in transfers, the APE in particular.

Let us now turn to case 2, in which the husband has a median wage full time job. Then there are three major changes: even when the woman does not work, the couple is not eligible to the RMI; it only gets housing subsidies if it has at least three children; and it pays the income tax in some configurations. Thus the family composition now affects net disposable income through taxes as well as transfers.

As shown in Table 2, a birth of parity 2 or higher now has very different

⁹Thus, the household of a woman who works full time and has a second child gains 330 euros per month.

¹⁰APE and APJE cannot be cumulated.

Table 2: Income Gain with a Birth (case 2)

Birth	No job	Full time	Stop working
Parity 1	220	195	-500
Parity 2	650	270	-40
Parity 3	690	370	55

effects than in case 1 when the woman does not work: her gain from the APE is not reduced by a decrease in the RMI, so that the income gain is much larger. On the other hand, it is of the same magnitude as in case 1 when the woman works. When she stops working for a first birth, the loss in wages dwarfs the small decrease in the income tax and the perception of the APJE. But for a higher parity birth, the APE fills the gap and the net gain in fact is positive for a parity 3 birth¹¹.

2.2 Simulations on a Representative Sample

The benchmark cases presented above illustrate that the short run financial gain from a birth varies a lot with income and employment status, even conditioning on the family composition. But they are not completely convincing, as the situations they depict may be rare in the population. Therefore they should be complemented by simulating the operation of the tax-benefit system on a representative sample of the population.

For this purpose, as well as for estimation purposes later on, we rely on data from the French Labor Force Survey. We need precise information on household income in order to simulate benefits; and we also want to model women's participation decisions as simply as possible. This leads us to notably reduce the number of observations: we leave aside households which contain an adult retiree, or a self-employed person, or a civil servant¹²; we also focus on women who have left school for more than two years, so as to avoid dealing with the schooling vs work decision. Finally and in this section only, we focus on fertile women, aged less than 40 years (the econometric study concerns women aged 20 to 50). The simulations retraced here on the

¹¹A well-known feature of the French income tax is the *quotient familial*, a form of income splitting which has highly regressive effects. At the levels of income in cases 1 and 2, the income tax bill is small and the *quotient familial* in fact has smaller effects than, say, the dependence of housing subsidies on the number of children. The benefits of the *quotient familial* reach a maximum at much higher income levels, of about 15,000 euros per month.

¹²Civil servants have their own benefit system, which we did not try to simulate.

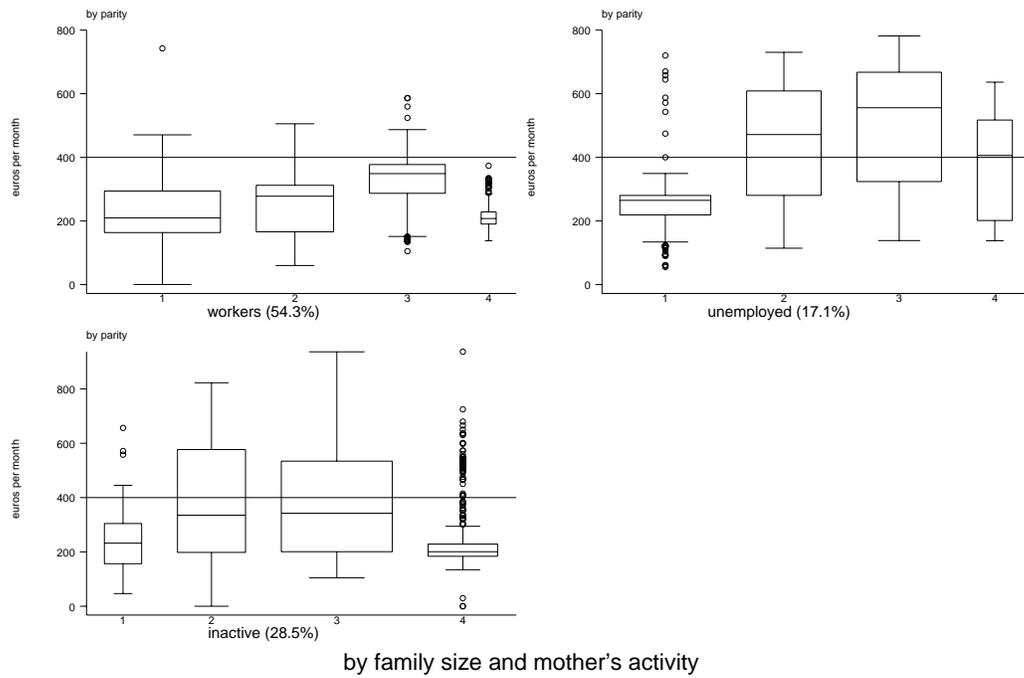


Figure 3: Increase of income after a birth, unchanged hours

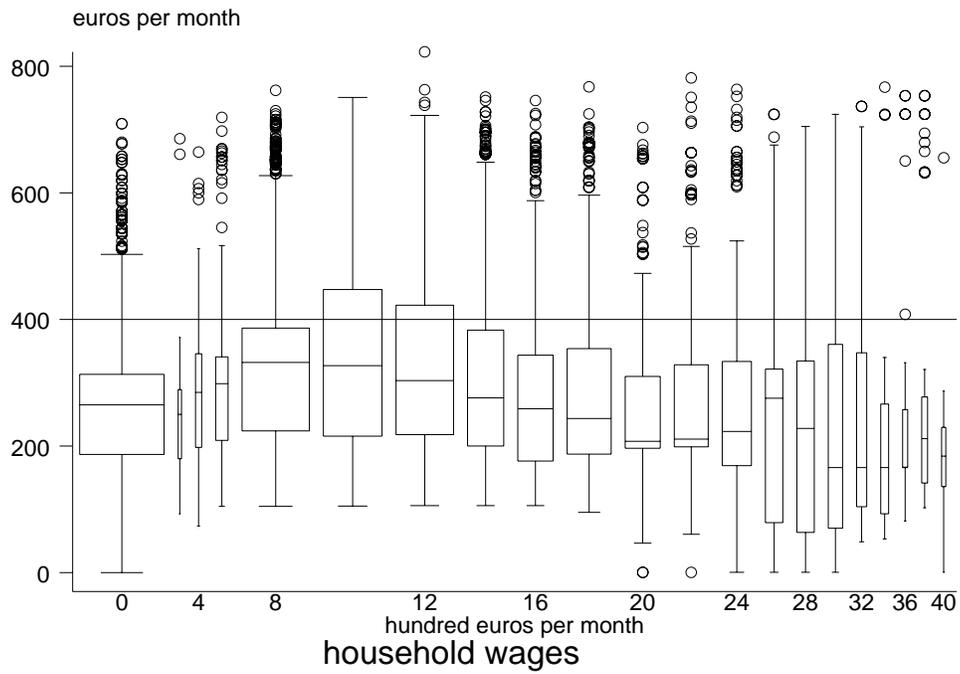


Figure 4: Increase of income after a birth, unchanged hours

1999 survey results bear on a sample of 10,333 observations that represent 3,400,000 women, whose ages are roughly uniformly distributed between 20 and 40. 80% of these women have a partner. 57% are married, 37% are single and 6% are divorced. 28% are childless, 26% have one child, 31% have two, 12% have three, and 3% have at least four. These women are less skilled than the average French woman: we left aside the civil servants, who are often teachers. 23% have no diploma, 36% have a diploma below the high school diploma, 17% finished high school, 14% completed two years of study after high school, and 10% went beyond that.

We observe for each woman in this sample the composition of her family, whether she is eligible to unemployment benefits, her type of housing (public or not); we also observe her wage when she works, and her partner's wage (if any). We use these characteristics to reconstitute the net disposable household income in 1999 by simulating the various taxes and benefits (see the Appendix). The result is the **reference income** RI .

To examine the *mechanical* effect of a birth on income, we first assume that each adult member of the household keeps the same working hours and wages as those observed in the 1999 survey and we compute the impact on income of a birth that occurs during 1998. Thus we simulate a birth in all households that did not witness one and denote BI_+ the income after that birth; and we subtract a newborn in households that did witness a birth in 1998: the resulting income we denote BI_- . Thus the variation in income due to a birth ΔI is either $BI_+ - RI$ or $RI - BI_-$.

Figures 3 and 4 use Tukey boxes to plot the distribution of ΔI in the sample for various values of the employment status of the mother, the parity of the birth and the total wages in the household. In Figure 3, the top left box presents the distribution of ΔI for a birth of parity 1 to 4 for wage-earning women. The other boxes concern unemployed and inactive women. Figure 3 suggests several observations. First, the birth of a first child gives an income gain that varies little in the population: the median ΔI is slightly above 200 euros/month, and the interquartile interval is only about 100 euros/month, irrespective of the woman's employment status. Next, the second child has both larger and more variable effects, especially for women who are unemployed or inactive, as many of them benefit from the APE. When an active woman has a third child, her household receives an even larger ΔI , due to an increase in family benefits. For many unemployed women, the second child is three or older, so that the third birth brings the APE; while for inactive women, the second child is often younger and thus the household already received the APE. The fourth child brings a smaller ΔI , independently of the employment status of the mother. Finally, fertility seems to be less encouraged by the tax-benefit system for wage-earning mothers than for

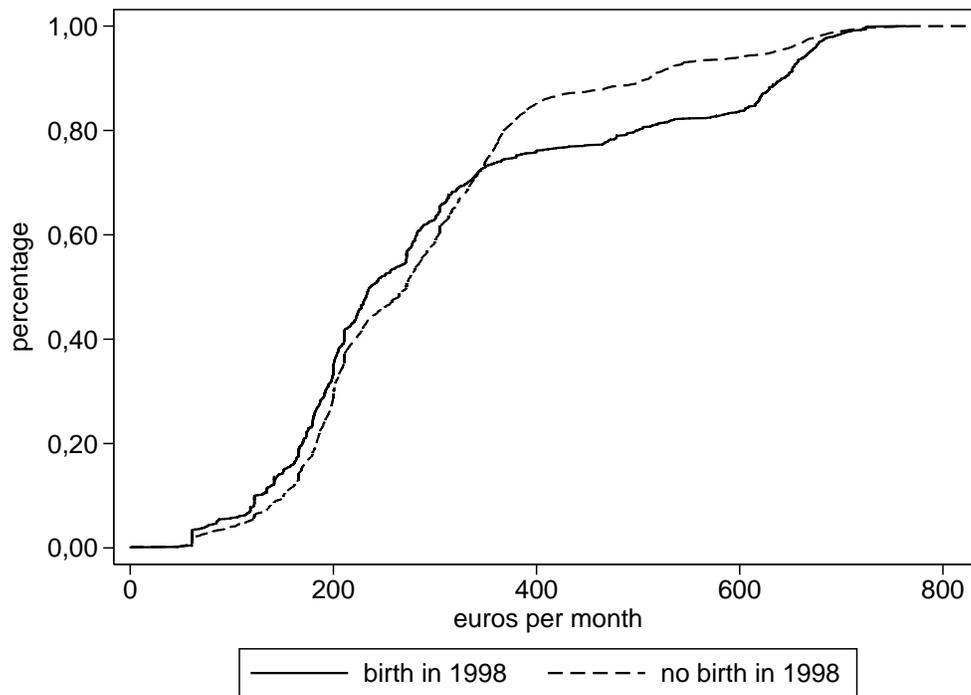


Figure 5: Cdf of ΔI for unchanged hours, depending on whether a birth took place in 1998

unemployed or inactive women; this is largely due to the APE.

Figure 4 shows how the distribution of ΔI varies with the total wages in the household¹³. The n th wage group goes from $100n$ to $100(n + 2)$ euros. The figure only reaches up to 4,200 euros, since there are very few observations above. Birth-induced income gains seem to vary much less with total household wages than with birth parity and/or employment status (as given by the previous figure). There is a local maximum around 1,000 to 1,500 euros/month, and then the median decreases slowly. As explained in section 2.1, the effects of the *quotient familial* are more spectacular but only appear for much higher incomes.

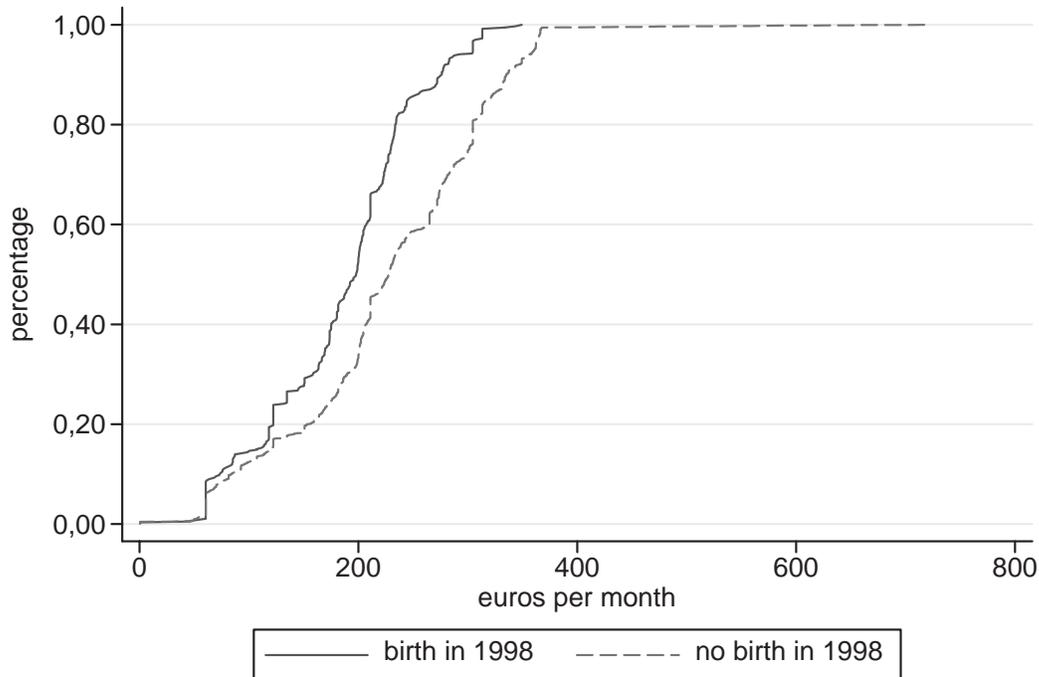


Figure 6: Cdf of ΔI per birth status, first child

3 Reduced-form Evidence

Figures 3 and 4 mix households that had a child or not in 1998. But it is useful to separate them to see whether financial incentives have an impact on fertility. Fertility decisions often go with a change in employment status; but for those women who keep the same working hours, we expect the distribution of $RI - BI_-$ (for households who had a child in 1998) to first-order stochastically dominate that of $BI_+ - RI$ (for those who did not have a child in 1998). Figure 5 plots the cdfs of these two distributions. The two curves are very close on the first three quartiles of the distribution; then the distribution “with birth” indeed clearly dominates the other, by about 200 euros/month. The corresponding women are often non-working women who receive the APE. This is confirmed by Figures 6 to 8, which show a stronger evidence for births of parities 2 and 3, but no such thing for the first birth.

While this evidence is suggestive, it may be due to subpopulation differences. One could imagine, for instance, that less skilled women have more children. Since they also receive more family benefits, this could explain the

¹³Of course, wages are correlated with age, family composition and so on.

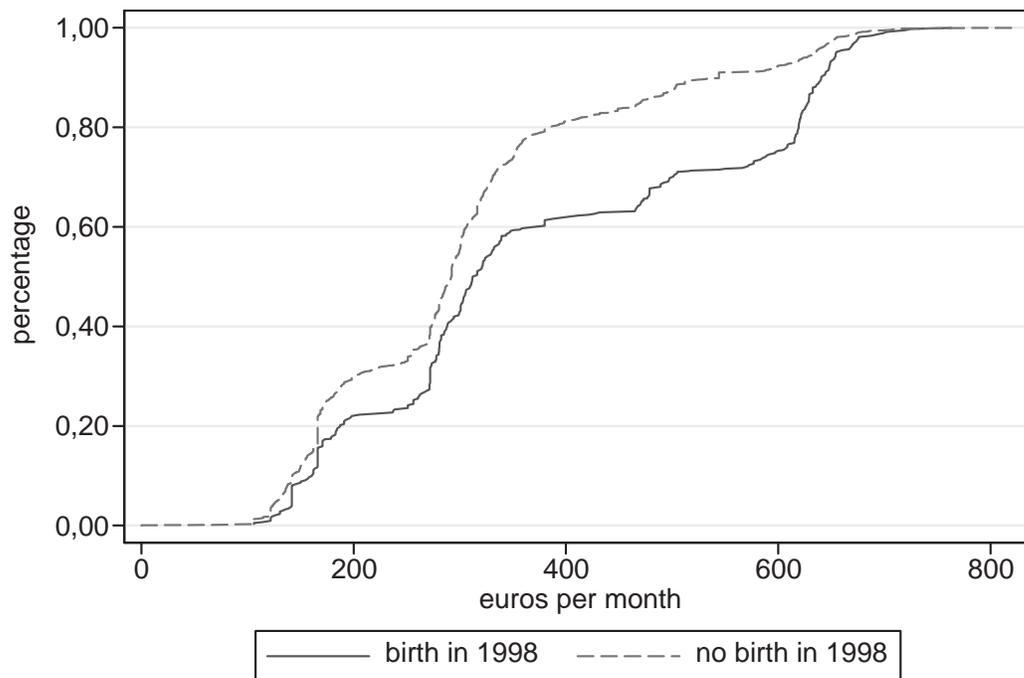


Figure 7: Cdf of ΔI per birth status, second child

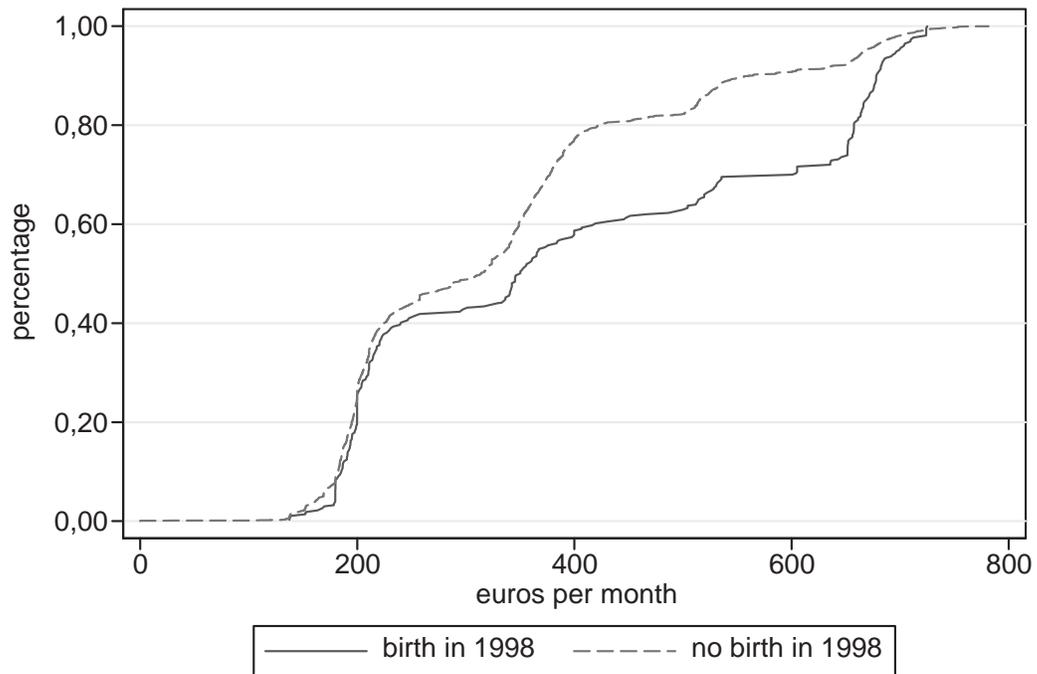


Figure 8: Cdf of ΔI per birth status, third child

stochastic dominance in these figures. We need to condition on observable characteristics. To do this, we estimate a probit model for the occurrence of a birth:

$$\Pr(\text{birth}|\mathcal{X}) = \Phi(\mathcal{X}\beta + E\alpha)$$

where \mathcal{X} contains all explanatory variables in X , X^E , X^W and X^L (about a hundred variables¹⁴), Φ denotes the cumulative distribution function of a standardized centered normal, and E measures financial incentives for fertility.

Let R_{ij} be the net disposable incomes of the household in the four possible contingencies:

- R_{00} : non employment in year $(t + 1)$ and no birth in t
- R_{01} : non employment in $(t + 1)$ and birth in t
- R_{10} : employment in $(t + 1)$ and no birth in t
- R_{11} : employment in $(t + 1)$ and birth in t .

These incomes are the disposable incomes of the household after taxes and transfers, which depend on the wages of the members of the household (when they work) and on family composition. The labor supply decision changes wage income and the fertility decision modifies family composition, and consequently disposable income.

These four incomes give rise to three measures of short-term¹⁵ fertility incentives:

- $E_1 = R_{01} - R_{00}$, the income variation when a jobless woman has a child.
- $E_2 = R_{01} - R_{10}$, the income variation when a woman has a child and quits her job.
- $E_3 = R_{11} - R_{10}$, the income variation when a working woman has a child.

We take E to be the matrix (E_1, E_2, E_3) , interacted with indicators of parity 1, 2 and 3. Note that R_{10} and R_{11} depend on the woman's wage W , which we only observe when the woman is actually working in our sample; to evaluate the missing values, we estimated a participation model and used the simulated values for W .

¹⁴These are the same variables that we use in the structural model in section 4.

¹⁵The fit does not improve when we discount family benefits as explained in the introduction.

Table 3: Probit-based Simulated Effects

Incentive measure	Parity 1	Parity 2	Parity 3
E_1	6.0 (7.7)	-5.7 (3.0)	11.6 (3.6)
E_2	2.9 (3.6)	0.3 (2.5)	6.1 (3.1)
E_3	7.8 (14.1)	-9.1 (7.9)	-4.8 (12.5)

Table 3 summarizes the results we obtained when estimating this probit model on the 9,297 women of the 1997, 1998 and 1999 Labor Force Surveys who are in the peak fertility period (ages 25 to 35). The figures in the table represent the (sample average) increase in the probability of a birth when one of these three incentive measures increases by 100 euros per month:

$$100 \frac{\partial \ln \Pr(\text{birth}|\mathcal{X})}{\partial E}.$$

The standard errors of these figures are between parentheses. Only two cells in the table are statistically significant. They both relate to women who already have two children. They suggest that providing 100 euros/month more generous incentives to non-working women (resp. women who stop working) who already have two children would increase their birth rate by 11.6% (resp. 6.1%). However, there are also three cells with negative effects, one of which (for non-working women with one child) is largeish and close to being statistically significant. Another disappointing feature is that when we do not interact the incentive measures with parity, the results are quite different: E_2 still has a sizable positive (and statistically significant) effect on fertility, but E_1 has a negligible effect and E_3 has a fairly large (but insignificant) negative effect. Dropping the interaction with parity is, admittedly, not a very credible restriction; and it is very strongly rejected by the data. But all in all, this probit-based approach does not give very solid evidence of fertility effects.

A pessimistic view of these results would conclude that the effects we are looking for just cannot be found in the data. This would indeed be true if we had a huge number of observations and we used a nonparametric regression. Neither is true (or possible) here; moreover, common sense suggests several reasonable variations of this reduced-form model. For instance, fertility incentives for working women should matter more for women who are more likely to be employed; this suggests interacting E_2 and E_3 with a proxy for

the probability of participation. But there is little distance from there to a structural discrete choice model of participation of fertility, to which we now turn.

4 A Model of Labor Supply and Fertility

We work on three successive cross sections. We aim to get an estimate of the effect of financial incentives on fertility through the large variability of situations in the data. Indeed, if financial incentives have a strong effect on fertility, we expect more numerous births in households which, all other things equal, receive more money after a birth.

We jointly model fertility and labor supply. These two decisions are indeed linked by the design of the benefits: the APE is conditional on reducing labor supply, and several other family benefits, the RMI (the minimum income guarantee) and housing benefits, jointly depend on the family composition and on income (therefore indirectly on employment status). We rely on a detailed representation of the transfers (the function R in the introduction); we slightly deviate from the presentation in the introduction in that employment not only depends on labor supply but also on the level of the minimum wage as in Laroque and Salanié (2002)¹⁶. The employer cost of the minimum wage is rather large in France and our previous modeling experience shows that it would be a mistake to neglect this fact.

4.1 The Discrete Choice Model

As mentioned in the Introduction, we limit the financial expectations to one year in this paper, except for family benefits that are smoothed over an 18-year period. We also leave aside part-time work for simplicity. We assume that the women make their fertility decisions in year t in a rather myopic way: they forecast their status in year $(t + 1)$, in four possible contingencies, being employed or not, with a newborn baby or not. They decide on having a baby in year t when they find it worthwhile in this four-way comparison for year $(t + 1)$. We also make the (somewhat unrealistic) assumption that the participation decision in year t does not influence the circumstances in date $(t + 1)$, so that, as far as date t is concerned, the participation decision and the fertility decision can be considered separately.

Wage income can only be observed for working women; thus we add a wage equation, which describes the labor cost for a full time (39 hour) work

¹⁶Clearly, this implies that the wage schedule cannot be identified nonparametrically below the minimum wage.

week:

$$\log C = X\beta + \sigma\varepsilon$$

where ε has a standard centered normal distribution.

Let d be the labor disutility and v the net (of associated costs) benefit of a birth for a woman of the sample, forecasted for year $(t + 1)$. We denote $\alpha_d > 0$ the sensitivity of labor supply to financial incentives and $\alpha_v > 0$ the sensitivity of fertility to financial incentives. We normalize the utilities associated with the four states by setting the coefficient of the income variable always equal to $\alpha_d\alpha_v$. Then the utilities to be compared are

$$\begin{cases} U_{00} &= \alpha_d\alpha_v R_{00} \\ U_{01} &= \alpha_d\alpha_v R_{01} + \alpha_d v \\ U_{10} &= \alpha_d\alpha_v R_{10} - \alpha_v d_0 \\ U_{11} &= \alpha_d\alpha_v R_{11} - \alpha_v d_1 + \alpha_d v \end{cases}$$

We assume that the labor disutility d_i , for $i = 0, 1$, is given by

$$d_i = Z_i\gamma + \rho\varepsilon + \eta_d$$

where Z_i gathers variables describing non wage income, family composition (age and number of children), the matrimonial status, and a variable that describes whether the woman has a newborn or not, so that labor disutility differs when the woman just had a baby (d_1) or not (d_0). The variable ε stands for an unobserved heterogeneity on wages, which may be correlated with labor disutility. The net benefit v of a birth is specified as

$$v = V\delta + \eta_v$$

where V includes variables describing the family composition, matrimonial status and the age of the woman. The random pair (η_d, η_v) is assumed to have a joint centered standard normal distribution. It is normalized by setting both standard errors to be one, with a free correlation coefficient c .

A woman in the population knows her characteristics, both observable (V, Z) and unobservable $(\varepsilon, \eta_d$ and $\eta_v)$ to the econometrician. Barring any obstacle to employment, she chooses to have a baby when either U_{01} or U_{11} is the largest of the four numbers U_{00} , U_{01} , U_{10} and U_{11} , and not to have a baby otherwise.

Participation in date t is consistent with the forecast made earlier in date $t - 1$. Given the family composition, in particular the possible presence of a new born in the household, the woman participates when the benefits from work exceed the cost, evaluated with the same parameters d , v , α_d and α_v as above.

Remark 1: From the point of view of the econometrician, who does not know the value of the idiosyncratic shocks but assumes a normal distribution, note that if Φ is the c.d.f. of the standard centered normal distribution, then conditionally on the decision of being unemployed, the probability of a birth is

$$\Phi(V\delta + \alpha_v(R_{01} - R_{00}))$$

which justifies our heuristic interpretation of α_v as a measure of the sensitivity of fertility to financial incentives. Similarly, the probability of working in the next year given a birth is

$$\Phi(\alpha_d(R_{11} - R_{01}) - Z_1\gamma - \rho\varepsilon)$$

and α_d does measure the sensitivity of labor supply to financial incentives.

Remark 2: In the notation of model (1) in the Introduction, E consists here of the 4-tuple $(R_{00}, R_{01}, R_{10}, R_{11})$ and X contains all variables in Z_0, Z_1 and V . R_{10} and R_{11} depend on the woman's wage, which we only observe if she works; otherwise we use her simulated wage, as is standard in participation models.

4.2 Empirical Implementation

We work on three cross-sections of women, interviewed at the Enquêtes Emploi (Labor Force Surveys) in March 1997, March 1998 and January 1999. The exogenous variables of the model are observed in the survey. The endogenous variables are the employment status at the time of the survey, the wage in case of employment¹⁷, and the occurrence of a birth during the year, as observed in the following survey. We select observations as described in section 2.2; this gives us 22,996 observations to work with.

We programmed the tax and transfer schedule in 1997, 1998 and 1999, through a function R which maps the detailed characteristics of the household, labor costs of man and woman, number and age of the children, etc., into its disposable income. To keep notation simple, we only make explicit two arguments so that $R(C, F)$ denotes the disposable income of the household under the then current legislation, when the labor cost of the woman is C (equal to 0 when she does not work) and the family composition of her household is F . As explained in the introduction, we adopt a semi-myopic formulation: a woman who plans her fertility and participation behavior in year t only takes into account the consequences on her household's disposable income in $(t + 1)$, except for family benefits b_f . If $B_t(C, F)$ measures

¹⁷Remember that we only consider two possible situations on the labor market: non employment or full-time employment.

the benefits accruing to a family of characteristics (C, F) according to the legislation at date t , family benefits are smoothed over eighteen years:

$$b_{ft} = \frac{1}{1 - \beta^{19}} \sum_{\tau=t}^{t+18} \beta^{\tau-t} B_t(C_t, F_{t\tau}).$$

This requires assumptions about what happens after $t + 1$): we assume that participation then is as in $(t + 1)$, and no more children are born. These are clearly very strong assumptions: but something like this is required to keep computations tractable, and at least this specification allows us to differentiate between benefits that are only given to very young children (like the APE) and benefits that last until children come of age (as the *allocations familiales*).

We assume a myopic expectation of the legislation, so that the same function B_t is used at $\tau = t + i$, $i = 1, \dots, 18$. Also all household characteristics are supposed to stay constant¹⁸, except for the natural aging (without deaths nor new births) of the family: $F_{t\tau}$ denotes the naturally forecasted age structure of the family from the observed composition of the household F_{tt} at date t . The smoothing factor β is a parameter of the model; it is estimated at 0.58, with a standard error of 0.06.

According to the previous description, the model then operates as follows. In year t ($t = 1997, 1998, 1999$), the woman decides on her labor supply by comparing $\alpha_d R(C, F_{tt})$ and $(\alpha_d R(0, F_{tt}) + d)$. A woman who has decided to participate finds a job if her labor cost C exceeds the cost of the minimum wage \underline{C} . At the same time, the woman plans on possibly giving birth during t , by anticipating her situation in $(t + 1)$ and comparing the utilities of the four possible states discussed above. A couple more remarks are in order:

- the less skilled women, for whom $C < \underline{C}$, in fact only compare the two situations of non employment, since the minimum wage excludes them from the labor force.
- the decision to have a child is not necessary followed by a birth, and a birth is not always voluntary. However, the data do not allow us to separate these circumstances, except by functional form assumptions. We therefore identify willingness to have a child with birth.

¹⁸These include the woman's market wage: thus we neglect the negative effect of fertility on the wages-experience nexus.

5 Estimation Results

The model is estimated by full information maximum likelihood¹⁹. The wage and disutility of labor equations have standard specifications. The explanatory variables X of the wage equation include the age at the end of studies and its square, experience²⁰ and its square, and indicator variables for the six levels of diplomas. The variables Z which explain the disutility of labor are the numbers of children less than one year old, less than three years old, between 3 and 6, and between 6 and 18, the age of the woman, an indicator of marital status, income when not working and its interactions with family composition and marital status. As explained in section 4, the disutility of labor should also be allowed to vary when the woman has just given birth; thus we also include interactions of the dummy for a newborn child with the number of older children and $R(0, F_{tt})$. We specify the sensitivity of labor disutility with respect to financial incentives α_d as a quadratic function of age.

Specifying the fertility equation is not as straightforward. The willingness to give birth certainly depends on several groups of variables, such as diplomas, matrimonial status, family composition and age. There is no obvious way to parsimoniously combine these variables, and experimenting is costly. Moreover it is likely that the sensitivity of fertility to financial incentives, if it exists, varies both with age and with the rank of the child.

For all these reasons, the specification we consider here is not fully satisfactory. We chose to err on the side of a fairly general specification and included a total of 75 explanatory variables in V . We specify one equation for each birth parity: this amounts to conditioning on past fertility, which we hope alleviates the endogeneity problem mentioned in the Introduction²¹. Each such equation contains marital status, the age breakdown of children, diploma and unearned income $R(0, F_{tt})$, all interacted with a quadratic function of the woman's age (truncated at 40)²².

¹⁹Individuals stay in the Labor Force Survey panel for three years, so that some of the individuals in one of our three cross-sections can also be found in another cross-section. Ideally, we would model this correlation by a factor model and integrate over the common element. As this would have greatly increased the computation time, we chose to treat all observations as independent.

²⁰Measured, lacking better information, as the time since the end of studies.

²¹Any estimation bias on the crucial α_v coefficients due to unobserved “fecundity” is proportional to its covariance with financial incentives, conditional on our 75 variables and parity. We think this is likely to be small.

²²For births of parity 3, we also include an indicator that the first two children are girls—this seems more relevant in France than the more symmetrical “same sex” variable used by Angrist and Evans (1998).

If it is hard to choose the variables to be included in V , it is even harder to decide what α_v should be allowed to depend on. It seems natural to allow α_v to vary with birth parity. It is also likely that the sensitivity of fertility to financial incentives varies with the level of fertility itself. To take this into account, we specify for each birth parity:

$$\alpha_v = \alpha_v^0 + \frac{\alpha_v^1}{\phi(V\delta)} + \alpha_v^2 \frac{\Phi(V\delta)}{\phi(V\delta)}$$

To understand this formula, remember that, conditionally on being unemployed, the probability of a birth is (approximately) $P_N = \Phi(V\delta + \alpha_v \Delta R)$, where ΔR represents the income gain from a birth. If only α_v^1 is nonzero, then a first-order expansion in $\Delta R = 0$ gives

$$\frac{\partial P_N}{\partial(\Delta R)} = \alpha_v^1$$

so that α_v^1 measures the sensitivity of fertility in levels. Similarly, if only α_v^2 is nonzero, we find

$$\frac{\partial \log P_N}{\partial(\Delta R)} = \alpha_v^2$$

and α_v^2 measures a semi-elasticity. Finally, if only α_v^0 is nonzero, then

$$\frac{\partial P_N}{\partial(\Delta R)} = \alpha_v^0 \phi(V\delta)$$

and since P_N is much smaller than one half, the sensitivity of fertility to incentives increases with its level.

We first comment briefly on the labour market estimates. Concerning the wage equation, the inclusion of the fertility equation in the model only marginally modifies the results of our previous studies. Most coefficients of the disutility of labor d_i are significantly different from zero and have the expected sign: the presence of children increases the disutility of labor, especially when the children are young. The disutility of labor increases with the wage of the partner, through $R(0, F_{tt})$, and with age. The matrimonial status has the expected sign: married women are less eager to get a job on the labor market. The estimated value of ρ is 0.25 (standard error 0.04), which implies that a high productivity is associated with a somewhat higher labor disutility, everything else equal. The correlation c between η_d and η_v turns out to be negligible at -0.01 (0.04). The presence of a newborn child has a very strong disincentive effect on participation (this effect is less pronounced for women whose partner earns high wages).

Table 4: $1000 \times \alpha_d$

Variable	Coefficient	Standard error
Age	196.6	27.4
Age squared	-2.9	0.3
Constant	-1.6	0.6

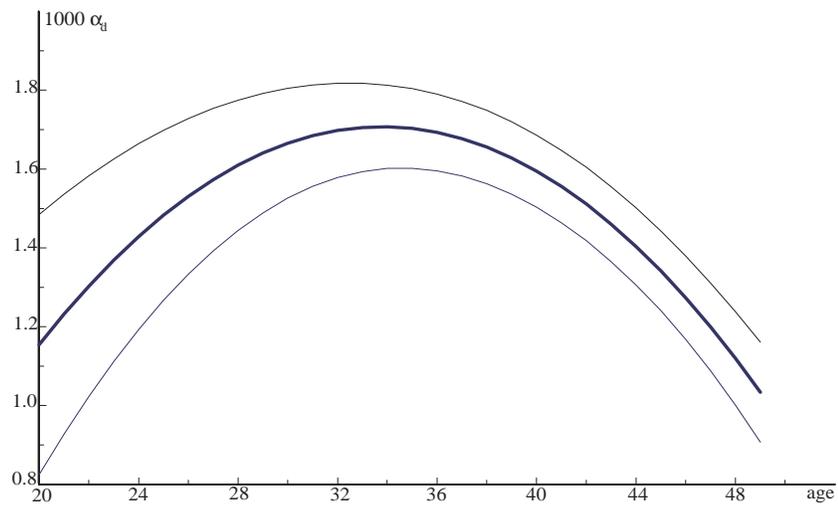


Figure 9: Sensitivity of Labor Supply to Financial Incentives

Table 5: Fertility equation: significance tests for groups of variables

Variables	Degrees of freedom	$\chi^2(d)$ (p-value)		
		Parity 1	Parity 2	Parity 3
Income out of work	6	5.6 (47%)	4.7 (59%)	0.9 (99%)
Diploma	3	22.0 (0.0%)	9.1 (2.8%)	2.3 (51%)
Financial incentives	2	12.4 (0.2%)	61.3 (0.0%)	42.2 (0.0%)
Two daughters	3	–	–	6.6 (8.5%)

Table 4 and Figure 9 present the estimates of the sensitivity of labor supply to financial incentives. The sensitivity of participation to financial incentives varies somewhat with age; it is highest for women aged 35. The average estimated value of α_d is in line with our previous results; as we will see, it suggests a notable effect of financial incentives on participation.

The variables appearing in the fertility equation are very numerous. The coefficients of a number of these variables and of their interactions are not significantly different from zero. Table 5 shows that income out of work, together with its square is non significant: any influence of this variable on fertility seems to come through the participation decision of the mother. Diploma significantly affects the timing of the first birth (studies delay the arrival of the first-born), but does not play any further role in the enlargement of the family. On the other hand, financial incentives appear to be present at all parities, and are strongly significant for births of parity 2 or higher. Finally, families with two daughters seem to be more likely to want a third child than others, at the 10% significance level.

It is not easy to read the effects on fertility from the estimated coefficients, given all the interactions. Thus we prefer to show plots of the predicted fertility rate $\Phi(V\hat{\delta})$ of a nonworking woman with zero income gain from a birth ($\Delta R = 0$) for some benchmark cases.

The top left graphs in Figure 10 plot fertility rates as a function of age for first and second births, depending on the age of the first born in the latter case. Our estimates suggest that the maximal fertility rate occurs at 26: then childless women have close to a 0.5 probability of having a child during the coming year. The presence of a newborn considerably reduces this figure, while when this child reaches 2, the fertility rate increases to about 0.4. This

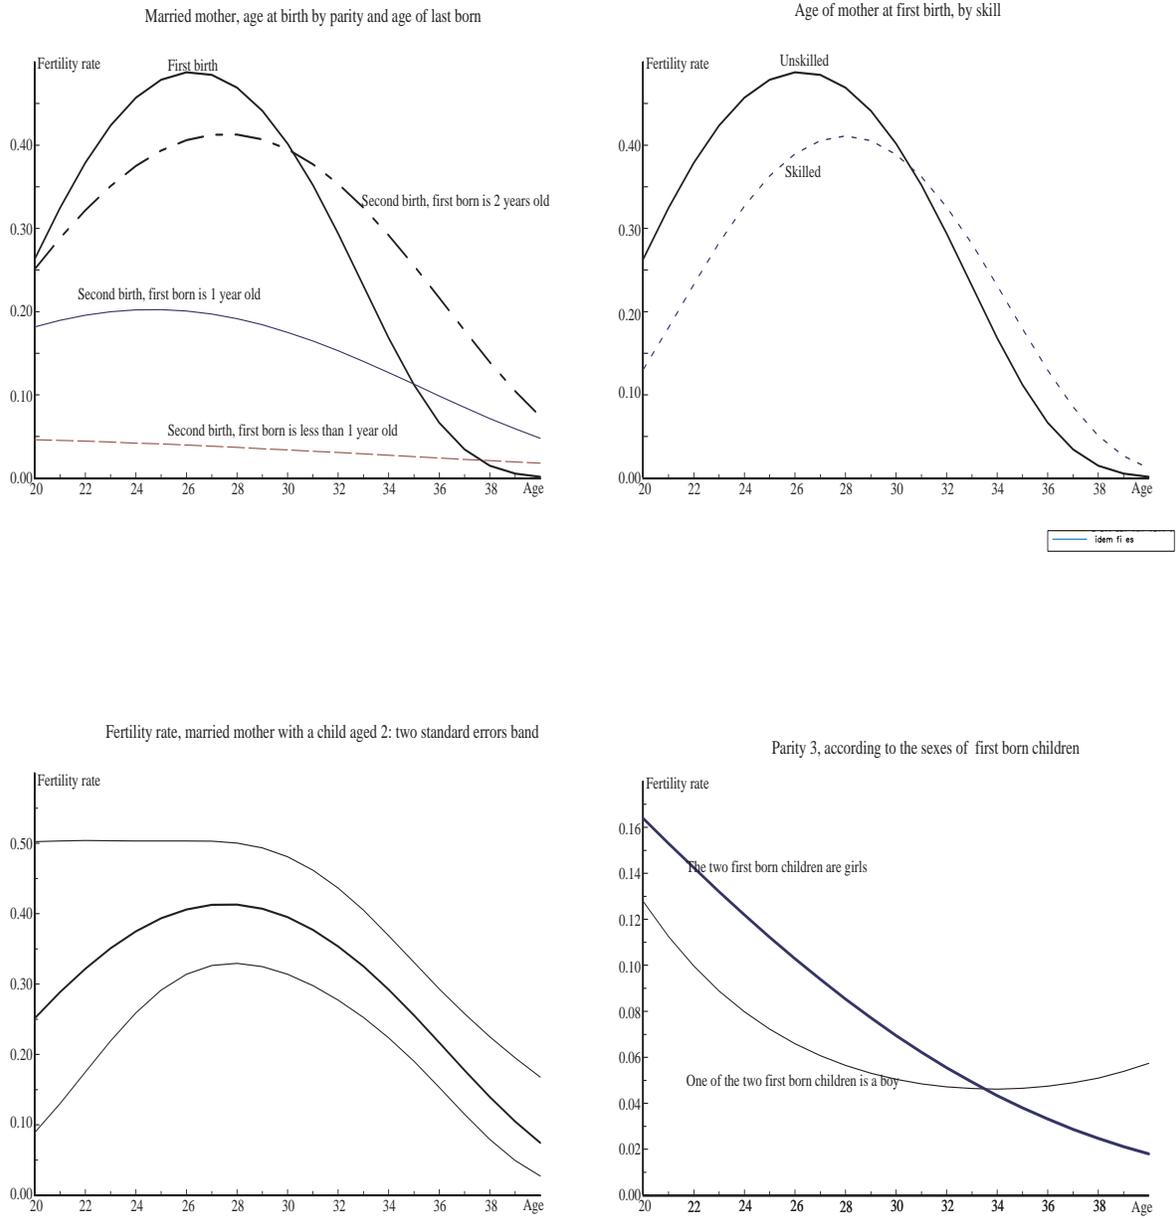


Figure 10: Estimated Fertility Rates

points towards an optimal spacing behavior. To illustrate the precision of our results, the bottom left panel of Figure 10 gives the 95% confidence interval of the predicted fertility rate for second births when the older child is 2. The top right of the figure plots the first birth rates according to skill, which are significantly different as already seen in Table 5: less skilled women have their first child at a younger age, on average more than two years before more skilled women. Finally, the bottom right graph shows the (not very large, but significant at the 10% level) difference in fertility between two-children families, depending whether the two children are girls, or not. The difference vanishes, and even changes sign, between ages 33 and 34.

There are a number of ways to describe the quality of the fit of the model. We focus on fertility. Given the estimated coefficients, it is easy to compute the probability of a birth in each household of the sample, conditional on the exogenous variables and on the (endogenous) employment status and wage. This probability indeed is a by-product of the computation of the likelihood function. The average of this probability for the households where a birth occurred is 0.22, while it is equal to 0.05 for the other households. The pseudo R^2 , ratio of the variance of the above probability to the variance of the birth indicator variable, is 0.25. Table 6 compares the observed fertility rates with those estimated for the whole sample and for various subgroups. The model correctly predicts that lower-skilled woman and (spectacularly) women without a stable partner have fewer children; it also gives a reasonably good fit for the age profile of women who give birth.

The main novelty in the paper is the estimation of the sensitivity of fertility to financial incentives α_v . Table 7 gives the estimation results. For several coefficients, the maximization hit on the zero boundary and stayed there²³. Note that while the two non-zero coefficients attached to the first birth appear separately non different from zero, the estimators are correlated and the χ^2 test of Table 5 shows that they are jointly significant. To see what their values imply, we may first go back to the illustrative calculations above, and remember that the approximate effect of financial incentives on the probability of giving birth is

$$\frac{\partial P_N}{\partial(\Delta R)} = \alpha_v^0 \phi(V\delta) + \alpha_v^1 + \alpha_v^2 \Phi(V\delta).$$

Take a non-working woman with the same characteristics as those used for Figure 10. Then an increase in financial incentives of 100 euros/month yields an increase from the reference fertility rate of

$$0.02 + 0.065 \times \text{reference fertility rate in percent,}$$

²³This is robust to changes in initial parameter values.

Table 6: The Fit of Fertility on the Sample

Criterion	Observations		Fit
	Number	Fertility rate	
All	722,625	0.066	0.065
Parity of the birth			
Parity one	245,262	0.056	0.056
Parity two	281,420	0.090	0.091
Parity three or more	195,944	0.055	0.053
Presence of a young child			
$C_{-1} = 0$	697,355	0.069	0.068
$C_{-1} > 0$	25,271	0.030	0.033
Age of the mother			
20-24	66,007	0.200	0.194
25-26	84,287	0.166	0.168
27-28	119,730	0.164	0.172
29-30	139,275	0.151	0.143
31-35	206,690	0.087	0.092
36-40	90,932	0.040	0.034
41-49	15,705	0.004	0.004
Mother's diploma			
Graduate	62,169	0.088	0.082
Undergraduate	103,970	0.087	0.086
High school	132,477	0.084	0.085
Basic technical training	186,615	0.061	0.065
Junior high school	49,791	0.050	0.054
No diploma	187,604	0.054	0.049
Partner status and employment status			
Alone	48,383	0.021	0.022
With partner	674,243	0.077	0.076
Salaried in t	299,486	0.057	0.067
Unemployed	423,140	0.073	0.064

Table 7: $1000 \times \alpha_v$

Parameter	Estimate	Standard error
Parity 1		
α_v^0	0.130	0.428
α_v^1	0.000	–
α_v^2	0.223	0.582
Parity 2		
α_v^0	0.000	–
α_v^1	0.002	0.001
α_v^2	0.655	0.222
Parity 3 and higher		
α_v^0	0.000	–
α_v^1	0.003	0.002
α_v^2	0.886	0.372

or

$$0.03 + 0.089 \times \text{reference fertility rate in percent},$$

measured in percentage points, respectively for births of parity 2 and of parity three or more. The impact is essentially nil, when the reference fertility rate is zero. For births of parity 2, the fertility rate belongs to the 20%–40% range for women in their early thirties: an extra 100 euros raises a 30% rate by $0.02 + 0.065 \times 30 = 2.0$ percentage point, or a 6.7 percent increase in fertility. For births of parity 3, the reference fertility rate is typically lower, around 20%. Increasing by 100 euros the monthly benefit associated with a birth then raises the fertility rate by $0.03 + 0.089 \times 20 = 1.8$ percentage point, a 9.0 percent increase in fertility. A similar computation for parity 1, at the peak of fertility when the probability of a birth is equal to 0.5, i.e. $V\delta = 0$, shows that an increase of 100 euros in financial incentives leads to an increase of $0.013/\sqrt{2\pi} + 0.022 \times 50 = 1.1$ percentage point, corresponding to a 2.2% increase of the fertility rate.

These numbers may seem high. Note, however, that a permanent income supplement of 100 euros per month and per child is a very costly measure: about 1% of GDP, or slightly more than the cost of the existing family benefit system. Our estimates are therefore in the same ballpark as those coming from the probits in section 3, as well as those from the country studies described in the introduction. But our model gives a much richer description of the interaction of family benefits with the labor market and the tax-benefit system; this makes it possible to simulate the effect on fertility of various

policy reforms more realistically (see section 6).

The identifying assumption in our model is that the women who get most financial benefits from a birth should have more children. We already estimated such models in Laroque and Salanié (2004a) and in Laroque and Salanié (2004b), with rather different results. There are two main differences between these two previous studies and the current paper. The first one is that we used then a fully myopic model, so that the duration of benefits was not well taken into account. The second one is that we included fewer variables in the fertility equation. In Laroque and Salanié (2004b), we then found a strong influence of financial incentives on fertility, but none for two-children or larger families, while in Laroque and Salanié (2004a), with a different set of variables, there is essentially no effect of financial incentives. The added variables we include here give a richer specification of the socio-demographic determinants of fertility; but they also influence taxes and benefits. We believe that our previous results were marred with omitted variables bias, and that the current paper, where we have purposely allowed for a large number of determinants in the fertility equation, has more reliable estimates. Indeed, while it is hard to separate the direct and the indirect (through financial incentives) effects of all of these variables, there seems to be enough information in the data to identify the two.

To put it more precisely, remember the data generating process in (1). Any hope to identify f'_E relies on the conditional variance $V(E|X)$ being large enough in the data. As E depends on X^E and w by (2) and most of X^E is included in X , we must therefore hope that $\text{cov}(E, w|X)$ is large enough. But this is not clear, as shown by Figure 4 in section 2. On the other hand, we have about 23,000 observations, so the practical strength of this argument is not obvious. To test it, we studied what happens to the estimates of α_v when the model is overparameterized (by including too many variables in V). To do this, we simulated a “reasonable” fertility model with only about 20 variables in V and economically significant values of α_v ; then we estimated it with the specification in this paper. To our surprise, the estimated values of α_v were fairly close to the values we used to simulate the pseudo-data. While this is hardly conclusive evidence, it suggests that there is enough variance of E given X to estimate α_v in our sample. A supporting argument comes from regressing the short-term values of the fertility incentives E_1 , E_2 and E_3 of section 3 on all explanatory variables of the structural model. The results are shown in Table 8. They show that all these variables still leave a lot of variability to be explained in fertility incentives.

Table 8: Explaining Fertility Incentives

Incentive measure (euros/month)	Standard error (euros/month)	R^2
E_1	222	0.54
E_2	343	0.71
E_3	79	0.48

6 Simulating the Model

We get another measure of the effect of family benefits on fertility by simulating policy reforms on our sample. Our first experiment is a rather artificial one: we increase the women’s market wages C or unearned income $R(0, \cdot)$ (which includes the wages of their partners, if any) as in comparative statics exercises. The resulting variations in fertility have the sign expected from theory: increasing market wages of women by 5% reduces fertility by 0.6%, and increasing unearned income by 5% increases fertility by 1.5%. The first effect is rather small. The second effect is in large part due to the fact that a higher unearned income reduces participation and thus makes child rearing more attractive.

We mentioned in the Introduction that the extension of the *Allocation Parentale d’Éducation* (APE) to births of parity 2 in 1994 is often cited as a cause of the recent rebound in fertility, an increase of around 7% of the number of births between 1995 and 2000. We can simulate its abolition in order to evaluate its contribution. It turns out that if we believe our estimates, the 1994 reform may have increased births of parity 2 by 10.9%; on the other hand, it made the birth of a third child relatively less attractive and may have reduced births of parity 3 by 2.4%. These contrasting effects result in an increase of 3.7% in the total number of births; this is about half of the observed rebound²⁴.

Our estimates also suggest that the fairly wide-ranging 2004 reform of family benefits called the *Prestation d’Accueil au Jeune Enfant* (PAJE), which had explicit pro-natalist objectives, may increase births by 4.7%. To go to the limit of what seems politically feasible, we also simulated the effect of creating a monthly childcare credit of 180 euros²⁵ per child younger

²⁴These figures must be interpreted cautiously. Indeed, what we estimate are probabilities conditional on family composition; and the family composition of our sample changes over time due to a policy reform. The simulations we discuss in this section do not take this into account; but these indirect effects are probably not very large.

²⁵This is the poverty line for a young child in France.

than 3. This would increase fertility by as much as 13.4%; but it would also be rather costly (about 0.3% of GDP, roughly half of the cost of the current family benefits).

All of these reforms also have side-effects on participation and employment. We find that the PAJE would have little effect on employment, which it would reduce by 0.3%. The 180 euro childcare credit also would be almost neutral on employment, since it is perfectly unconditional. On the other hand, the extension of the APE may have reduced participation by 1.7% according to our simulations. This is a very large figure if one remembers that the eligible women are only those with exactly one child and a job; yet it confirms estimates we obtained in our earlier work, as well as estimates based on a natural experiments method. This last experiment also serves to show two important phenomena. First, family benefits may have significant employment effects. Second, their effects on fertility are both direct (when $\alpha_v > 0$) and indirect; the latter are due to the interaction between employment and fertility. Thus the APE stimulates births in part because, by inciting some women to quit their jobs, it makes it easier for them to have children.

Are our figures small or large? Let us first take a policy viewpoint. According to our results, it is possible to design family benefits so as to increase fertility by a few percentage points; and the 1994 APE reform may have caused half of the current rebound in birth rates. On the other hand, it seems unlikely that such reforms can effectively reverse the effects of major changes in fertility such as the “baby bust” of 1965-1975. Returning to basic economics, our paper can be seen as an attempt to measure the price elasticity of the demand for children. Unfortunately, defining and evaluating the “cost of children” is not an easy task, both in conceptual and in practical terms. Browning (1992), for instance, distinguishes four quite different definitions in section 3.1 of his valuable survey; and none of them yields itself to direct measurement. Cutting to the quick, we resort to the OECD equivalence scale, which attributes 1 point to the first adult in a household, 0.5 for every other adult or child younger than 14, and 0.3 for every child younger than 14. According to this, a household with disposable income I and e points on the equivalence scale would need an income supplement of $0.3I/e$ to keep the same living standard after a new birth. We then take this income supplement to measure the cost of a child.

In our sample population, the average cost of a child is 310 euros per month, with a fairly large standard error of 140 euros/month. It is somewhat smaller for higher parities: 330 for parity 1, 300 for parity 2 and 280 for parity 3. We then simulate a uniform increase of 10% of the cost of a child. The resulting reduction in births is 2.2%, which points towards a price elasticity

of the demand for children of 0.2. This elasticity is small compared to that of, say, married women's participation; but it is comparable to that of men's participation decisions.

References

- ANGRIST, J., AND W. EVANS (1998): “Children and Their Parents’ Labor Supply: Evidence From Exogenous Variations in Size,” *American Economic Review*, 88, 450–477.
- APPS, P., AND R. REES (2001): “Fertility, Female Labour Supply and Public Policy,” Discussion Paper 409, IZA.
- BECKER, G. (1960): “An Economic Analysis of Fertility,” in *Demographic and Economic Change in Developed Countries*. NBER.
- (1991): *A Treatise on the Family*. Harvard University Press.
- BESLEY, T., AND A. CASE (2000): “Unnatural Experiments? Estimating the Incidence of Endogenous Policies,” *Economic Journal*, 110, F672–F694.
- BLANCHET, D., AND O. EKERT-JAFFÉ (1994): “The Demographic Impact of Family Benefits: Evidence from a Micromodel and from Macrodata,” in *The Family, the Market and the State in Ageing Societies*, ed. by J. Ermisch, and N. Ogawa. Clarendon Press.
- BLAU, D., AND P. ROBINS (1989): “Fertility, Employment, and Child-care Costs,” *Demography*, 26, 287–299.
- BROWNING, M. (1992): “Children and Household Economic Behavior,” *Journal of Economic Literature*, 30, 1434–1475.
- CHONÉ, P., D. LEBLANC, AND I. ROBERT-BOBÉE (2002): “Offre de travail féminine et garde des enfants,” forthcoming in *Economie et prévision*.
- CIGNO, A. (1986): “Fertility and the Tax-Benefit System: A Reconsideration of the Theory of Family Taxation,” *Economic Journal*, 96, 1035–1051.
- DEL BOCA, D. (2002): “The Effect of Child Care and Part-time Opportunities on Participation and Fertility Decisions in Italy,” *Journal of Population Economics*, 15, 549–573.
- EKERT-JAFFÉ, O. (1986): “Effets et limites des aides financières aux familles : une expérience, un modèle,” *Population*, 2, 327–348.
- FRANCESCONI, M. (2002): “A Joint Dynamic Model of Fertility and Work of Married Women,” *Journal of Labor Economics*, 20, 336–380.
- GAUTHIER, A. (1996): *The State and the Family*. Oxford University Press.

- GAUTHIER, A., AND J. HATZIUS (1997): “Family Benefits and Fertility: An Econometric Model,” *Population Studies*, 51, 295–346.
- HECKMAN, J. (1976): “Simultaneous Equation Models with Continuous and Discrete Endogenous Variables and Structural Shifts,” in *Studies in Non-linear Estimation*, ed. by S. Goldfeld, and R. Quandt. North Holland.
- HECKMAN, J., AND J. WALKER (1989): “Using Goodness of Fit and Other Criteria to Choose among Competing Duration Models: A Case Study of Hutterite Data,” in *Sociological Methodology*, ed. by C. Clogg. American Sociological Association.
- (1990a): “The Relationship between Wages and Income and the Timing and Spacing of Births: Evidence from Swedish Longitudinal Data,” *Econometrica*, 58, 1411–1441.
- (1990b): “The Third Birth in Sweden,” *Journal of Population Economics*, 3, 235–275.
- HOTZ, J., R. KLERMAN, AND R. WILLIS (1997): “The Economics of Fertility in Developed Countries,” in *Handbook of Population Economics*, ed. by M. Rosenzweig, and O. Stark. North Holland.
- HOTZ, J., AND R. MILLER (1988): “An Empirical Model of Lifecycle Fertility and Female Labour Supply,” *Econometrica*, 56, 91–118.
- KEANE, M., AND K. WOLPIN (2002a): “Estimating Welfare Effects Consistent with Forward-looking Behaviour, I: Lessons from a Simulation Exercise,” *Journal of Human Resources*, 37, 570–599.
- (2002b): “Estimating Welfare Effects Consistent with Forward-looking Behaviour, II: Empirical Results,” *Journal of Human Resources*, 37, 600–622.
- KEARNEY, M. S. (2004): “Is There an Effect of Incremental Welfare Benefits on Fertility Behavior? A Look at the Family Cap,” *Journal of Human Resources*, 39, 295–325.
- LAROQUE, G., AND B. SALANIÉ (2002): “Labor Market Institutions and Employment in France,” *Journal of Applied Econometrics*, 17, 25–48.
- (2004a): “Fécondité et offre de travail des femmes en France,” *Économie publique*, 13(2), 61–94.

- (2004b): “Fertility and Financial Incentives in France,” *CESIfo Economic Studies*, 3/04, forthcoming.
- LEFEBVRE, P., L. BROUILLETTE, AND C. FELTEAU (1994): “Les effets des impôts et des allocations familiales sur les comportements de fécondité et de travail des canadiennes : résultat d’un modèle de choix discrets,” *Population*, 2, 415–456.
- MILLIGAN, K. (2004): “Subsidizing the Stork: New Evidence on Tax Incentives and Fertility,” *Review of Economics and Statistics*, forthcoming.
- MOFFITT, R. (1998): “The Effect of Welfare on Marriage and Fertility,” in *Welfare, the Family, and Reproductive Behavior: Research Perspectives*, ed. by R. Moffitt. National Academy Press.
- ROSENZWEIG, M. (1999): “Welfare, Marital Prospects, and Nonmarital Childbearing,” *Journal of Political Economy*, 107, S3–S32.
- ROSENZWEIG, M., AND T. P. SCHULTZ (1985): “The Demand for and Supply of Births: Fertility and its Lifecycle Consequences,” *American Economic Review*, 75, 992–1015.
- ROSENZWEIG, M., AND K. WOLPIN (2000): “Natural ‘Natural Experiments’ in Economics,” *Journal of Economic Literature*, 38, 827–874.
- THÉLOT, C., AND M. VILLAC (1998): *Politique familiale : bilan et perspectives*. La Documentation Française.
- WHITTINGTON, L. (1992): “Taxes and the Family: The Impact of the Tax Exemption for Dependents on Marital Fertility,” *Demography*, 29, 215–226.
- WOLPIN, K. (1984): “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality,” *Journal of Political Economy*, 92, 852–874.

Appendix: The Tax-benefit Microsimulation Model and the Fertility Coefficients

Our simulations account for most of the French taxes and benefits that are likely to influence participation and fertility decisions. We give here a brief description of the taxes and benefits which depend on household composition.

- Family benefits are not taxable. They include
 - the *allocations familiales*, which are not means-tested and are given to all households with at least two children.
 - the *complément familial* and the *allocation pour jeune enfant* (APJE) which is also means-tested.
 - the *allocation parentale d'éducation* (APE). The APE is given to parents (mostly women) who reduce their labor supply after the birth of a child of parity at least two. It is slightly larger than the net minimum wage for a half-time job when the woman stops working, irrespective of her household income. Only women who worked at least two years in the five years before the birth (in the ten years for a birth of parity 3 or more) may perceive the APE. The APE, the APJE and the *complément familial* are mutually exclusive.
- The *revenu minimum d'insertion* (RMI) is a differential benefit; it guarantees a minimum income to households whose head is older than 25. It is about 60% of the net minimum wage for a childless couple; it increases with the number of children and is withdrawn at a 100% marginal rate when other income increases.
- Housing subsidies finance a part of the rent paid by low- and middle-income households. Their value depends on many criteria (family composition, income, rent, type of housing, town of residence) which our simulation models rather coarsely.
- Finally, the housing tax and especially the income tax also play an important role in our analysis, as they depend on family composition.

These benefits are granted for a variable duration. The APE and the APJE only extend until the youngest child is three, and even the *allocations familiales* expire when the child is 20. Ideally, we would want to compute the discounted variation in income due to a birth given all future benefits, which is very difficult since this birth also changes future labor supply and births.

We therefore adopt here a short-term approach: the variation in income is only evaluated for the year that follows the birth. We are aware that this is unsatisfactory, especially as we adopt a long-term approach for the RMI, neglecting transitory rules.