

Does Fertility Respond to Financial Incentives?

Guy Laroque¹ Bernard Salanié^{2,3}

July 21, 2010

¹CREST-INSEE and University College London.

²Columbia University and CEPR, corresponding author; send correspondence to bsalanie@columbia.edu.

³This research benefited from the financial support of the Alliance program during a visit of Guy Laroque at Columbia University. We are grateful to Gary Becker, Didier Blanchet, Pierre-André Chiappori, 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 for their comments.

Abstract

In recent years more and more governments have resorted to financial incentives to increase fertility; yet the evidence on the impact of such measures is not very conclusive. We put forward an identification strategy that relies on the fact that variation of wages induces variation in benefits and tax credits among “comparable” households. We implement this approach by estimating a discrete choice model of female participation and fertility, using individual data from the French Labor Force Survey and a fairly detailed representation of the French tax-benefit system. Our results suggest that financial incentives have had a significant effect on fertility decisions in France, both for the first and for the third child. As an example, we simulate the effects of an additional, unconditional child credit of 150 euros per month; according to our estimates, this might raise total fertility by about 0.3 point.

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) and Willis (1973) 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). Yet in many countries (see Gauthier (1996)), family benefits are mostly designed as a way to ensure a minimum standard of living to families and children. There are notable exceptions. For instance, the policies towards families implemented in France at the end of the 1930s were in part based on the belief that family benefits increase fertility, and a 2004 reform explicitly mentions fertility as a concern. Sweden and Québec have also implemented pro-natalist policies; more recently, other countries with much more dire fertility issues such as Germany and Russia have enacted policy measures designed to bring them closer to the replacement rate.

Modeling fertility behavior is notoriously hard; commonly used models have low predictive power. Several approaches have been used so far (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) pair¹. This line of work relies on the reasonable assumption that there is variation in family policies over time and space that is not entirely explained by the other explanatory variables. On the other hand, it suffers from the usual problems of panel country regressions: it relies on hard-to-test functional form assumptions; the data omits many variables which may be correlated with the variation in policy parameters; and the studies sum up very complex tax-benefit systems in a couple of variables only.

To alleviate this difficulty, more recent work has used individual data, which alone can provide the analyst with a large number of observations and with enough information to properly measure financial incentives to fertility. There is a small body of econometric literature that aims at estimating the ef-

¹Thus 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.

fect of female and male wages on fertility. 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; but as these papers make no attempt to use the form of the tax-benefit system, the estimated coefficients of wages are hard to relate to effects of financial incentives.

A few 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; by comparing fertility of similar women residing in Québec and in other Canadian states over this period, he finds that this benefit strongly stimulated fertility. However, his approach cannot yield an estimate of a "price elasticity", which we would need in order to simulate the effect of alternative policies. Similar remarks apply to Kearney (2004), who estimates the fertility effect of the introduction of family caps in many US states following the 1996 reform of welfare². Finally, Cohen, Dehejia, and Romanov (2007) exploit a drastic reduction of child allowances in Israel that affected mostly large families in 2003; their results suggest a large negative impact on the fertility of poorer households and new immigrants. But the situation in Israel has quite specific features, most prominently the very high fertility of religious families whose labor supply is very low.

In contrast to these papers, we seek to identify behavioral elasticities. This raises two difficulties: the treatment of endogeneity, and accounting for the dynamic aspect of fertility decisions. Endogeneity results from what we call "fecundity" in this paper: the cumulative effect of the unobserved heterogeneity in preferences for children and in biological fertility. Women that are more fecund are likely to invest less in human capital, so that the market wages they may claim and the financial incentives they face are different from

²A number of papers that studied the effects of the American welfare system also report its estimated impact on fertility, although this is usually not their main focus (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.

that of less fecund women. This may make financial incentives endogenous in the fertility equation.

Dynamic considerations are also important. The birth of a child has consequences that extend beyond the duration of most benefits. Partly for this reason, the demographic literature stresses the distinction between the timing of births and completed fertility; it is commonly held that family benefits influence the former but have much weaker effects on the latter (in the jargon of demographers, “tempo” effects dominate “quantum” effects.)

Since Wolpin (1984) estimated the first dynamic model of fertility behavior, many papers have sought to solve these difficulties by setting up a full dynamic programming model. The role of financial incentives for fertility is usually one among other decision margins in this literature. To cite one of the more recent contributions, Keane and Wolpin (2007) estimate a dynamic programming model of the decisions of young women. Their simulations suggest that the EITC, whose amount is very low for single-child families, has increased the fertility of high school dropouts by about 0.2 child per woman. Todd and Wolpin (2005) study the PROGRESA subsidy program in Mexico; they report that “fertility outcomes are essentially invariant to the subsidies.”

The dynamic programming models deal with endogeneity concerns by allowing for permanent unobserved heterogeneity components (“types”) whose several dimensions include wage-earning ability, fecundity, and labor disutility among others. By construction, they allow for a clear distinction between timing of births and completed fertility effects. On the other hand, these models carry a heavy computational burden, with some negative consequences. First, it limits the analyst’s ability to experiment with alternative specifications. While this may not be a major issue for well-studied decision margins such as labor supply, we would argue that it is particularly detrimental when modelling ill-understood phenomena such as fertility. Second, the number of types is usually very small; e.g. Todd and Wolpin (2005) allow for three types only as the computational burden otherwise becomes unmanageable (see their footnote 33.) Finally and most importantly for our purposes, the representation of the tax-benefit system is heavily simplified in these papers. Keane and Wolpin (2007) marks a significant step forward by

using a 5-parameter benefit rule that they estimate on state-level panel data; but even so, this can only account for a small part of the variation across households.

Since the effects of the usual kind of financial incentives on fertility are likely to be small, they can be hard to isolate from changes in omitted variables; and so alternative approaches seem worth exploring. Useful variation in fertility incentives may come from policy changes; but drastic policy reforms are few and far between. Most of the useful variation in fact is to be found on cross-section data. It consists of the differences in fertility- and employment-contingent household incomes generated by different labor market opportunities and by the specificities of the tax-benefit system. Nonlinearities of tax-benefit schemes have been much used to study labor supply, see e.g. Blundell, Duncan, and Meghir (1998) and Laroque and Salanié (2002). A crucial identifying tool is a fairly detailed representation of the relevant taxes and benefits.

Section 2 presents and justifies our identification strategy, which is based on comparing the fertility outcomes of women who have similar characteristics except for their wage and the wage of their partner. As family benefits depend on the two parents' primary incomes, this provides variation in incentives that we use to identify the relevant elasticity. We also account for women's participation decisions, given that they are so closely linked to their fertility choices³. Our approach is based on the assumption that unexplained variation in wages is independent of unexplained fecundity, which is debatable; we discuss it later in the introduction.

Section 2 shows that such variation is sufficient to nonparametrically identify labor supply behavior. On the other hand, we need to add more structure in order to identify the joint labor supply and fertility decisions. We therefore introduce an additional assumption, which states that labor supply shocks,

³Lefebvre, Brouillette, and Felteau (1994) 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. However, we devote much more attention to identification issues. Similarly, our earlier work on this topic (Laroque and Salanié (2004a, 2004b)) did not give center stage to issues of identification; moreover, we included older women, which tended to contaminate our estimates.

conditional on realized fertility, are independent of the fertility shock. This orthogonality assumption is natural here, since the fertility decision usually takes place about one year before the conditional choice of labor supply. It yields a nonparametric generalization of the nested logit specification; this may be useful in other applications where the conditional independence assumption is economically reasonable.

Our paper omits a number of important factors. The most glaring one is dynamics: we reduce the complex dynamic decision process of couples to a static reduced form⁴. 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 neglect child-care subsidies, which are important both in theory⁵ and in practice⁶, as taking them into account would require modeling the choice of child care mode.

We apply the approach described above to the case of France. We focus on French data not only because we know it well, but also because France has a rather generous and diverse family benefit system. Its cost (including tax credits) is evaluated at about 0.8% of GDP, and it comprises some unconditional, some means-tested, and some employment-tested benefits. We set up a discrete choice model of fertility and participation decisions on the French Labor Force Surveys of 1997, 1998 and 1999. 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.

We present our estimation results in section 3. The overall procedure appears to work well and to yield robust estimates. They suggest that financial incentives play a notable role in determining fertility decisions in France, both

⁴We experimented with variants in which couples anticipate a flow of family benefits, based on the average fertility and labor market trajectories. The results did not change by much.

⁵See e.g. Apps and Rees (2004).

⁶Blau and Robins (1989) estimate the effect of child-care costs on fertility in the US, using geographical variation in child-care availability. See also del Boca (2002) for Italy and Choné, Leblanc, and Robert-Bobée (2004) for France.

for the first and for the third child. Our key assumption is that unobserved heterogeneity (fecundity) does not induce a correlation between wages and fertility, conditional on the observables. We are confident that it holds here, for two reasons. First, our specification uses a much larger set of covariates than typical dynamic programming models; in particular, we condition on the existing number of children and their ages. Since more fecund women are likely to be more fecund at any age, this probably helps reduce the unobserved heterogeneity. Second, in order to test that the unexplained part of wages is exogenous to fertility, we add the generalized residuals of wages in the fertility equation; and conversely, we include the generalized residuals of fertility in the wage equation. In both cases, the estimated coefficients of these additional variables are very small and non-significant. All this concurs to validate the assumption that wages can in fact be taken as exogenous. We also test the implied separability of the functional form along lines similar to Ichimura and Taber (2002); this is only rejected for the second child.

We simulate the increase in births that results from adding to the existing tax-benefit system a child subsidy of 150 euros per month. Such a policy reform would have a direct cost of about 0.3% of GDP; we take this as a reasonable upper bound on what is politically feasible. According to our estimates, it would raise total fertility by about 0.3 point, while reducing female labor supply by about 0.5 point. Larger changes would presumably require acting on other dimensions such as child care and parental leave⁷.

Our paper completely skirts normative questions. Governments may subsidize fertility because they feel that current policies deter childbearing, or to “save Social Security”, and of course for a host of non-economic reasons. These are beyond our horizon here. On the other hand, it would be possible to use our estimated model to analyze the redistributive effects of any reform of family benefits; we leave this for further research.

⁷For instance, Lalive and Zweimueller (2009) shows that successive (large) changes in the duration of parental leave in Austria have had clear effects on both timing and number of births.

2 Identifying Incentive Effects

Those readers who are only interested in the specification and estimation of the model may want to skip this section and directly move to Section 3.

2.1 The Statistical Problem

Our aim here is to use data on fertility outcomes F_i , employment decisions L_i and explanatory variables Z_i for various observational units to infer the effect of financial incentives on fertility. Financial incentives to fertility depend on the tax-benefit-system, which we represent as a function that maps the labor supply and fertility decisions of the observation unit L_i and F_i into a net household income $R_{L_i F_i}(Z_i)$ that also depends on its characteristics Z_i . Thus we want to evaluate (both ex post and ex ante) the impact of changes in the function R on fertility outcomes. In France, the function R is very complicated, highly nonlinear and not even differentiable; but at least it is public information and it can be evaluated for all values of its arguments.

It is hard to discuss identification at this level of generality; let us therefore focus on the simple setup we actually use in our application, in which an observational unit is a household at a given date, the decision variable $L_i = 0, 1$ is the woman's extensive margin decision of working or not, and the fertility outcome $F_i = 0, 1$ is the occurrence of a birth.

The vector of explanatory variables Z_i may comprise several classes of variables. Since a woman at a given period is our observational unit, Z_i may contain:

- socio-demographic variables x_i such as education, region, past fertility, marriage status, and so on, which may impact current fertility directly or because they contribute to determining taxes and benefits;
- the market wage of this woman w_i , which we take to be her earning potential if she does not actually work;
- the actual wage income of her partner (if any), which we denote w_i^m .

Several of these variables are clearly not strongly exogenous: the unobserved “taste for children”, for instance, probably reduces human capital investment and therefore education; and past fertility can only be weakly exogenous. We abstract from endogeneity concerns in this section since we want to focus on identification: we assume that $Z = (x, w, w_m)$ is *strongly exogenous*. We also suppose for now that w is observed for all women, whereas of course it has to be imputed for non-working women.

If a woman’s labor supply decision is L , then her wage income wL enters the R function; but so does L itself since some benefits are employment-tested. Thus the tax-benefit system can be described in our application by four functions which correspond to the four alternatives $L = 0, 1; F = 0, 1$:

- net household income when the woman does not work and no birth occurred $R_{00}(Z)$;
- net household income when the woman does not work and a birth occurred $R_{01}(Z)$;
- net household income when the woman works and no birth occurred $R_{10}(Z)$;
- net household income when the woman works and a birth occurred $R_{11}(Z)$.

Note that when the woman does not work, her market wage w does not enter the net household income; thus $R_{00}(Z)$ and $R_{01}(Z)$ do not depend on w , but only on $X = (x, w^m)$.

Given this notation, let Q be the distribution of the endogenous variables L and F ; we can write it in all generality as

$$Q(L, F | R_{00}(X), R_{01}(X), R_{10}(Z), R_{11}(Z), Z)$$

and our aim is to evaluate changes in the distribution of F when any of the four functions R_{ij} shifts.

In this setup, we can only identify nonparametrically from the data the distribution of L and F conditional on Z , which we denote $P(L, F | Z)$. Under

our exogeneity assumptions, this allows us to answer questions such as:

- what is the effect of more education on fertility?
- what is the effect of higher male or female wages on fertility?

On the other hand, it does not allow us to evaluate policy-induced changes in R without further non-testable restrictions.

2.2 Identifying Assumptions

The basic problem we face is that since there are four alternatives, we can identify nonparametrically only three functions of Z from the data. But knowing Q requires identifying three functions of both Z and four other arguments (the R_{ij} 's.) Before we discuss possible sources of exclusion restrictions, we introduce our approach to identification on the simple labor supply model.

2.2.1 Identifying Labor Supply Behavior

As in our application, we focus here on the extensive margin⁸. In this case we observe a probability of working $P(Z)$; and our aim is “only” to recover a probability $Q(R_0(X), R_1(Z), Z)$, where R_1 (resp. R_0) is household net disposable income when the woman works (resp. does not work.) Here one simple way to achieve identification is to assume that

Assumption 1 *The woman’s market wage w only enters Q through its effect on R_1 (net household income when the woman works).*

Assumption 2 *The wage income of her partner (if any) w^m only enters Q through its effect on R_1 and on R_0 (net household income when the woman does not work).*

Recall our notation: $X = (x, w^m)$ and $Z = (X, w)$. Given assumptions 1 and 2, we can write $P(Z)$ as $Q(R_0(X), R_1(Z), x)$. Then changes in w identify

⁸The identification problem is similar on the intensive margin, unless of course (often implicit) assumptions are made to constrain the possible causes of variation in hours.

the effect of R_1 on labor supply Q ; and changes in w^m , compensated by changes in w so as to leave R_1 constant, identify the effect of R_0 on labor supply.

More formally, fix x and let $\mathcal{W}_2(x)$ be the support of the observed distribution of (w, w^m) given x . Then the effect of a change of the tax schedule (R_0, R_1) into (R'_0, R'_1) is identified as follows.

Take any (w, w^m) such that:

- $(w, w^m) \in \mathcal{W}_2(x)$;
- there exists a solution $(\tilde{w}, \tilde{w}^m) \in \mathcal{W}_2(x)$ to

$$\begin{cases} R'_0(w^m, x) & = R_0(\tilde{w}^m, x) \\ R'_1(w, w^m, x) & = R_1(\tilde{w}, \tilde{w}^m, x). \end{cases}$$

Then the effect of such a policy change on the labour supply of an individual with characteristics (w, w^m) is just

$$P(\tilde{w}, \tilde{w}^m, x) - P(w, w^m, x).$$

(and under assumptions 1 and 2, it clearly does not depend on the choice of the solution (\tilde{w}, \tilde{w}^m)).

When the range of $\mathcal{W}_2(x)$ is too small, identification may be sought through further restrictions, which often involve x . E.g. let there be a variable s that only enters Q through its effect on (R_0, R_1) , and denote $x = (s, t)$; then for given $Z = (w, w^m, s, t)$, we only need to find a solution $(\tilde{w}, \tilde{w}^m, \tilde{s})$ to

$$\begin{cases} R'_0(w^m, s, t) & = R_0(\tilde{w}^m, \tilde{s}, t) \\ R'_1(w, w^m, s, t) & = R_1(\tilde{w}, \tilde{w}^m, \tilde{s}, t) \end{cases}$$

in the support of the distribution of (w, w^m, s) given t , and the effect of the change then is

$$P(\tilde{w}, \tilde{w}^m, \tilde{s}, t) - P(w, w^m, s, t).$$

While this is enough to identify the effects of changes in taxes and benefits, we need more if we want to identify the underlying utility functions. From now on, the whole analysis is done conditional on x , which we drop from the notation. Let us set up a standard binomial model. Labor supply is the outcome of a utility maximization where

Assumption 3

- *the utility of the woman is $U_1(R_1) - d$ when she works and her household has disposable income R_1 ;*
- *it is $U_0(R_0)$ when she does not work and her household has disposable income R_0 ,*

where d measures the random disutility of labor, unobserved by the econometrician, which has a distribution G that does not depend on w or w^m . Similarly, U_0 and U_1 only depend on (w, w^m) through R_0 and R_1 .

Then we can reparameterize the probability of working as $Q(R_0, R_1)$. Identification of the binomial choice model has been studied by Matzkin (1992); but for our purposes, the following result in the line of Manski (1985,1988) is more convenient:

Theorem 1 *Assume the true model satisfies Assumptions 1, 2 and 3. Moreover assume that*

1. *G , U_0 and U_1 are differentiable, with positive derivatives;*
2. *denote \mathcal{R}_2 the image of \mathcal{W}_2 by the pair of functions (R_0, R_1) ; then \mathcal{R}_2 has a non-empty interior in \mathbb{R}^2 ;*
3. *there exists a point $(\tilde{r}_0, \tilde{r}_1)$ in \mathcal{R}_2 such that $Q(\tilde{r}_0, \tilde{r}_1) = 0.5$;*
4. *$G(0) = 0.5$.*

Take any triple of functions (u_0, u_1, g) that is consistent with the data generated by (U_0, U_1, G) . Then under these assumptions, there exist numbers a (positive) and b such that for all (r_0, r_1) in the interior of \mathcal{R}_2 ,

$$u_0(r_0) = aU_0(r_0) + b \quad \text{and} \quad u_1(r_1) = aU_1(r_1) + b;$$

and defining the function $\delta(r_0, r_1) = u_1(r_1) - u_0(r_0)$ over \mathcal{R}_2 ,

$$g(y) = Q(r_0, r_1) \text{ if } y = \delta(r_0, r_1).$$

Proof: Take (r_0, r_1) inside \mathcal{R}_2 . Since $G' > 0$, simple differentiation in the interior of \mathcal{R}_2 yields

$$-\frac{\frac{\partial Q}{\partial R_0}}{\frac{\partial Q}{\partial R_1}}(r_0, r_1) = \frac{U'_0(r_0)}{U'_1(r_1)}. \quad (D)$$

It follows that any possible (u_0, u_1) must satisfy

$$\frac{U'_0(r_0)}{u'_0(r_0)} = \frac{U'_1(r_1)}{u'_1(r_1)} = a,$$

where a is constant, since it cannot depend on r_1 by the first equality, nor on r_0 by the second equality (and a is positive since utilities are increasing.) Integrating gives

$$\begin{cases} u_0(r_0) = aU_0(r_0) + b_0 \\ u_1(r_1) = aU_1(r_1) + b_1. \end{cases}$$

Moreover, $G(0) = 0.5$ by assumption, so that

$$U_1(\tilde{r}_1) \equiv U_0(\tilde{r}_0).$$

This equality must also hold for u_0 and u_1 ; but $G(U_1(\tilde{r}_1) - U_0(\tilde{r}_0)) = 0.5$ can also be rewritten as

$$G\left(\frac{u_1(\tilde{r}_1) - u_0(\tilde{r}_0)}{a} - \frac{b_1 - b_0}{a}\right) = 0.5,$$

which implies

$$G\left(-\frac{b_1 - b_0}{a}\right) = 0.5$$

and so $b_0 = b_1 = b$. QED.

Note that this model has testable “nonparametric” implications:

1. Q must increase in R_0 and decrease in R_1 ;
2. $\log(-Q'_0/Q'_1)$ must have a zero cross derivative.

While the first one is fairly weak, the second one may provide an informative check on the model. Recall Assumption 3: conditional on x , the distribution of the disutility of labor d does not depend on (w, w_m) , that is, on (R_0, R_1) . If for instance the conditional variance of d depends on R_0 and R_1 in any way, then the second restriction above will not hold. If the conditional mean of d depends on (R_0, R_1) , tests based on the second restriction will also detect it unless this conditional mean is separably additive in R_0 and R_1 .

2.2.2 Missing Wages

To be complete, we should recall that wages w are only observed for working women. Selection may come from preference shocks or be linked to the minimum wage. The issue is akin to the treatment effect problem: to achieve nonparametric identification, we need one variable in the selection equation (into employment) that only enters the outcome equation (the wage equation) through its effect on the propensity score. The partner’s wage w^m (if present) can play such a role; in our specification we also use family composition variables to that effect: they figure in the disutility of work ($U_1 - U_0$), but not in the wage equation.

2.2.3 Identifying Fertility Behavior

For the more complicated case of joint fertility-participation behavior, such restrictions will not suffice. Pursuing the strategy in section 2.2.1, we would

start by assuming that w only enters Q through R_{10} and R_{11} and that w^m only enters Q through the R_{ij} 's. Then we would want to identify the change from $(R_{00}, R_{01}, R_{10}, R_{11})$ to $(R'_{00}, R'_{01}, R'_{10}, R'_{11})$ on a woman (w, w^m, x) by solving a system of equations of the form

$$\begin{cases} R'_{00}(w^m, x) &= R_{00}(\tilde{w}^m, \tilde{x}) \\ R'_{01}(w^m, x) &= R_{01}(\tilde{w}^m, \tilde{x}) \\ R'_{10}(w, w^m, x) &= R_{10}(\tilde{w}, \tilde{w}^m, \tilde{x}) \\ R'_{11}(w, w^m, x) &= R_{11}(\tilde{w}, \tilde{w}^m, \tilde{x}). \end{cases}$$

Since we have four equations, this can only work if we can find a subset s of the variables in x that only enters Q through the R_{ij} 's; and the dimension of s must be at least two. This is not a very appealing strategy: we can make a reasonable case for excluding some variables from the disutility of labor or the utility of income, but it is much harder for the costs and benefits of children.

A possible source for such variation in s is policy reforms that change the net cost of children, preferably in different ways for different types of households. As an example, the recent rebound in fertility in France (see figure 1) has sometimes 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” to the second-born in 1994. Piketty (2005) attempted to evaluate the consequences of this reform on fertility; he concludes that they are very hard to isolate, but that the 1994 reform cannot account for more than 20–30% of the observed increase in births.

We adopt a quite different approach: we resort to structural restrictions, which hopefully only act as identifying constraints. Again, we condition on x , and we drop it from the notation. We study the decision to work or not ($L = 1, 0$), with a newborn child or not ($F = 1, 0$). The choice among the four (L, F) pairs is derived from utility maximizing, where the utilities are




Figure 1: Children per woman from 1901 to 2007

given by

$$\begin{aligned}
 \mathcal{U}_{00} &= U_{00}(R_{00}) \\
 \mathcal{U}_{01} &= U_{01}(R_{01}) - c_1 \\
 \mathcal{U}_{10} &= U_{10}(R_{10}) - \zeta_{10} \\
 \mathcal{U}_{11} &= U_{11}(R_{11}) - c_1 - \zeta_{11}.
 \end{aligned} \tag{1}$$

Given biological constraints, decisions are taken sequentially: a woman decides on her fertility, then, given realized family composition, on her participation. The solution, as in dynamic programming, is computed backwards: first one compares \mathcal{U}_{0F} and \mathcal{U}_{1F} for a given value of $F = 0, 1$, then the expected values of $\max(\mathcal{U}_{0F}, \mathcal{U}_{1F})$ to choose the optimal F .

The utility indices U_{LF} are deterministic functions of household incomes, which are assumed to be perfectly anticipated. Without loss of generality, we take 00 to be the reference situation for which the disturbance is normalized to zero. The other random terms are written in such a way that ζ_{10} and ζ_{11} represent the labor supply disturbances in the second stage, given realized fertility, while c_1 represents a shock on the cost of a newborn child.

It is convenient to assume that the three terms $(\zeta_{10}, \zeta_{11}, c_1)$ are mutually independent.⁹

Assumption 4

- *the women anticipate perfectly their own wage as well as the wage of their partner;*
- *the shocks on the disutilities of labor given the presence of a newborn or not, ζ_{10} and ζ_{11} , and the shock on the cost of a child c_1 are independently distributed.*

We denote G_0 (resp. G_1) the unknown distribution of ζ_{10} (resp. ζ_{11}), and K the unknown distribution of c_1 . It turns out that these assumptions are (almost) enough to identify the true utility functions U_{00} , U_{01} , U_{10} and U_{11} and the distributions G_0 , G_1 and K up to the usual scale and location normalization.

First consider labour supply decisions given the realization of the “new-born” variable F for a woman of characteristics (w, w^m) . Given the realization of F , the woman decides to work iff $\mathcal{U}_{1F} > \mathcal{U}_{0F}$. So she works with probability

$$P_F(w, w^m) = G_F(U_{1F}(R_{1F}(w, w^m)) - U_{0F}(R_{0F}(w^m))).$$

This is of course the same problem as in section 2.2.1: a nice simplification from our choice of stochastic specification is that c_1 cancels out in the labour supply decision problem. Reparameterize $P_F(w, w^m)$ as $Q_F(R_{0F}, R_{1F})$ as before. Then we can apply Theorem 1, adapted to the new notation, for given $F = 0, 1$.

Theorem 2 *Assume that the true model satisfies assumptions 1 to 4, and in addition:*

1. *G_F , U_{0F} and U_{1F} are strictly increasing and differentiable;*

⁹In fact we only need to assume that the couple (ζ_{11}, c_1) is independent of ζ_{10} . But it would be messier since then we have to use the distribution of ζ_{11} conditional on c_1 when describing the labor supply decision after a birth.

2. $G_F(0) = 0.5$;

3. the support \mathcal{R}_{2F} of the distribution of $(R_{0F}(w^m), R_{1F}(w, w^m))$ has a non-empty interior;

4. it contains a $(\tilde{r}_{0F}, \tilde{r}_{1F})$ such that $U_{1F}(\tilde{r}_{1F}) = U_{0F}(\tilde{r}_{0F})$;

then the utility functions are identified up to scale and location. More precisely, for any representation (g_F, u_{LF}) consistent with the data, there exist numbers a_F (positive) and b_F such that for all $(r_0, r_1) \in \mathcal{R}_{2F}$,

$$u_{0F}(r_0) = a_F U_{0F}(r_0) + b_F \quad \text{and} \quad u_{1F}(r_1) = a_F U_{1F}(r_1) + b_F;$$

and defining the function $\delta_F(r_{0F}, r_{1F}) = u_1(r_{1F}) - u_0(r_{0F})$ over \mathcal{R}_{2F} , we have

$$g_F(y) = Q_F(r_{0F}, r_{1F}) \text{ if } y = \delta_F(r_{0F}, r_{1F}).$$

Note that this implies that on the range of δ_F , $g_F(y) = G_F(y/a_F)$.

Thus labor supply decisions in themselves give us some information about utility functions; but this is not enough to allow us to identify incentive effects of the R functions on fertility: changes in the ratio a_1/a_0 and in the cdf of the cost of children c_1 , on which we know nothing as yet, can give totally different estimates of the effect of fertility incentives.

To identify these remaining parameters, we of course need to study fertility decisions. Nature comes to our help, in that fertility decisions must be made in advance. The *true* indirect utility of a fertility decision $F = 0, 1$ is

$$V_F - c_F \equiv E \max(U_{0F}(R_{0F}(w^m)), U_{1F}(R_{1F}(w, w^m))) - \zeta_{1F} - c_F$$

where the expectation is only taken over the true cdf G_F of the shock ζ_{1F} , whose value is known to the woman at the time she takes her fertility decision (and we normalized $c_0 = 0$.)

Now for any variable ξ with cdf S , denote

$$H(S, \alpha, \beta) = E_\xi \max(\alpha - \xi, \beta).$$

Then we can write

$$V_F = H(G_F, U_{1F}(R_{1F}), U_{0F}(R_{0F})).$$

Therefore the probability of a newborn $F = 1$ is

$$P = \Pr(V_1 - c_1 > V_0) = K(V_1 - V_0) \equiv K(\Omega), \quad (K)$$

with

$$\begin{aligned} \Omega(w, w^m) = & H(G_1, U_{11}(R_{11}(w, w^m)), U_{01}(R_{01}(w))) \\ & - H(G_0, U_{10}(R_{10}(w, w^m)), U_{00}(R_{00}(w))). \quad (2) \end{aligned}$$

Labour supply data yield a representation (g_F, u_{LF}) equal to the true one (G_F, U_{LF}) up to an intercept b_F and a scale factor a_F :

$$\begin{cases} u_{LF}(y) & \equiv a_F U_{LF}(y) + b_F \\ g_F(a_F y) & \equiv G_F(y). \end{cases}$$

Any such choice of (a_0, a_1, b_0, b_1) generates its own representation of Ω , call it ω . The data gives us the probability $P(w, w^m)$ that a woman has a baby, conditional on her wage and her partner's. We must still solve the functional equation

$$P(w, w^m) \equiv k(\omega(w, w^m))$$

with unknowns k , (a_0, a_1) , and (b_0, b_1) .

To go further, we need some regularity assumptions. We first need to define an “expected marginal rate of substitution” for $F = 0, 1$:

$$M_F(w, w^m) = \frac{\frac{\partial V_F}{\partial w}}{\frac{\partial V_F}{\partial w^m}}(w, w^m);$$

We now state

Assumption 5 *There exists (\tilde{w}, \tilde{w}^m) in the interior of \mathcal{W}_2 such that $M_0(\tilde{w}, \tilde{w}^m) \neq M_1(\tilde{w}, \tilde{w}^m)$.*

Assumption 5 is very weak. It only requires that in expected terms (over the shocks on the disutility of labor ζ_{10} and ζ_{11}), the birth of a child changes the utility value of an increase in the woman's potential wage, relative to an increase in her partner's wage. There are several reasons why we expect this to hold, with $M_0 < M_1$ in fact for most values of (w, w^m) . Let us just mention two:

- since women spend more time with young children than men, an increase in their wage-earning potential is (relatively) less valuable if $F = 1$ since a birth reduces participation;
- some very significant family benefits, such as the APE in France, are much more attractive for low earners.

Note that this assumption is testable: in fact, given utility functions u_{LF} obtained after estimating labor supply, and estimates $P_F(Z)$ of the probability of working after fertility outcome $F = 0, 1$, the marginal utility of a wage increase is just the average of the marginal utilities when working and when not working, weighted by the corresponding probabilities¹⁰ which can be directly estimated from the data.

Define

$$v_F = H(g_F, u_{1F}(R_{1F}), u_{0F}(R_{0F})).$$

Then for $t = w$ or $t = w^m$, we can write:

$$\frac{\partial v_F}{\partial t} = P_F \frac{\partial u_{1F}}{\partial t} + (1 - P_F) \frac{\partial u_{0F}}{\partial t}; \quad (V)$$

and this is enough to estimate M_F since the derivatives of U_{LF} are equal to those of u_{LF} , known up to a scale factor a_F which cancels out in the ratio. We will present estimates of M_0 and M_1 later on.

We can now state

Theorem 3 *Let all assumptions of Theorem 2 hold, along with Assumption 5. Assume in addition that :*

¹⁰Given our assumptions, no woman is ever at the margin of participation.

1. K is differentiable and strictly increasing at $A(\tilde{w}, \tilde{w}^m)$;
2. the R_{LF} functions are differentiable at (\tilde{w}, \tilde{w}^m) .

Then

1. the marginal utilities are identified locally up to a common scale parameter. More precisely, any representation u_{LF} consistent with the data must satisfy

$$u'_{LF}(R_{LF}(w, w^m)) = \lambda U'_{LF}(R_{LF}(w, w^m))$$

over \mathcal{W}_2 for some $\lambda > 0$.

2. the cdfs of ζ_{10} , ζ_{11} and c_1 are also identified up to the choice of λ : any representation (G_0, G_1, K) consistent with the data must satisfy locally

$$\begin{cases} g_F(y) = G_F(y/\lambda) & \text{for } F = 0, 1; \\ k(y) = K(y/\lambda). \end{cases}$$

3. as a consequence, fertility effects of the tax-benefit system are exactly identified at every point of \mathcal{W}_2 .

Proof: under these assumptions, let (u_{LF}) , (a_F) , (b_F) , (g_F) be a representation compatible with the labour supply data. Now use (V) to define for $Z = (w, w^m) \in \mathcal{W}_2$, and denoting again $t = w$ or $t = w^m$:

$$\frac{\partial v_F}{\partial t}(Z) = P_F(Z) \frac{\partial u_{1F}}{\partial t} + (1 - P_F(Z)) \frac{\partial u_{0F}}{\partial t};$$

Clearly,

$$\frac{\partial v_F}{\partial t} = a_F \frac{\partial V_F}{\partial t}.$$

Now we can write

$$\frac{\partial P}{\partial t} = K'(K^{-1}(P)) \left(\frac{\partial V_1}{\partial t} - \frac{\partial V_0}{\partial t} \right) \quad (Q).$$

It follows that

$$\frac{\frac{\partial P}{\partial w}}{\frac{\partial P}{\partial w^m}} = \frac{\frac{\partial V_1}{\partial w} - \frac{\partial V_0}{\partial w}}{\frac{\partial V_1}{\partial w^m} - \frac{\partial V_0}{\partial w^m}},$$

or, with the alternative representation of the indirect utilities,

$$\frac{\frac{\partial P}{\partial w}}{\frac{\partial P}{\partial w^m}} = \frac{\frac{\partial v_1}{\partial w} - \frac{a_1}{a_0} \frac{\partial v_0}{\partial w}}{\frac{\partial v_1}{\partial w^m} - \frac{a_1}{a_0} \frac{\partial v_0}{\partial w^m}}.$$

This can be rearranged into

$$\frac{a_1}{a_0} \left(\frac{\partial v_0}{\partial w^m} \frac{\partial P}{\partial w} - \frac{\partial v_0}{\partial w} \frac{\partial P}{\partial w^m} \right) = \frac{\partial v_1}{\partial w} \frac{\partial P}{\partial w^m} - \frac{\partial v_1}{\partial w^m} \frac{\partial P}{\partial w}.$$

Everything in this equation is known at this stage, except for the ratio a_1/a_0 . So it determines this ratio uniquely, provided that the term that multiplies it is not zero in some t . But if this term were zero, then so would the term on the right-hand side; and simple algebra shows that then both M_0 and M_1 would be equal to the marginal rate of substitution of P . Taking $Z = (\tilde{w}, \tilde{w}^m)$ shows that $\frac{a_1}{a_0}$ is identified since by assumption 5, $M_0 \neq M_1$ in this point at least.

Now fix $a_0 = \lambda > 0$, which determines a_1 and fixes v_0 and v_1 ; and take any $Z \in \mathcal{W}_2$. Equation (Q) directly gives the density of k at $k^{-1}(P(t)) = v_1(t) - v_0(t)$. Now suppose we change the tax-benefit system so that R_{LF} becomes $R_{LF} + dR_{LF}$; then the probability of a birth changes by

$$dP = k'(k^{-1}(P)) \left(\frac{\partial v_1}{\partial R_{01}} dR_{01} + \frac{\partial v_1}{\partial R_{11}} dR_{11} - \frac{\partial v_0}{\partial R_{00}} dR_{00} - \frac{\partial v_0}{\partial R_{10}} dR_{10} \right); \quad (3)$$

but everything in this formula is identified since, e.g.,

$$\frac{\partial v_1}{\partial R_{01}} = (1 - P_F) \frac{\partial u_{01}}{\partial R_{01}}.$$

It is easy to see that a different choice of λ would just rescale k' and the u_{LF} 's in opposite directions, QED.

3 Specification and estimation results

3.1 Data

The model is estimated on data from the French labor force surveys of 1997, 1998 and 1999. We discard women who are civil servants¹¹, who are over forty years old, or who finished their studies less than two years before they were surveyed¹².

We decided to only keep women who were born French. Immigrant women (defined by being born non-French and outside of France) represent about 9% of fertile-age women in France, but 13% of births each year are to an immigrant mother. This higher fertility is partly due to a different age profile of immigrant women, and also to births being delayed until the wife of an immigrant worker comes to live in France. Women who immigrated to France before puberty in fact have very similar fertility to French-born women. Thus the fertility of immigrant women has quite different determinants than that of French-born women.

Furthermore, in the results that are shown below, we only consider births of parity 1, 2 or 3, and women who live with a partner. Just as for immigrant women, the behavior of both large families and single parent families is significantly different, and elucidating it would require another paper. This leaves us with a sample of 16,891 observations, whose distribution into employment and fertility status is as in Table 1.

Our method relies heavily on the variation in financial incentives induced by variation in wages w and w^m . Is there enough such variation in our dataset? A simple way to measure it is to define three measures of fertility incentives as

$$\begin{cases} E_1(w_m) & = R_{01}(w_m) - R_{00}(w_m) \\ E_2(w, w_m) & = R_{01}(w_m) - R_{10}(w, w_m) \\ E_3(w, w_m) & = R_{11}(w, w_m) - R_{10}(w, w_m) \end{cases}$$

¹¹Civil servants face different benefits schemes and labor market incentives.

¹²The latter is to avoid modeling the decision of pursuing studies or going into the labor force.

Fertility Employment	0	1	Total
0	6,844	1,751	8,595
1	7,430	866	8,296
Total	14,274	2,617	16,891

Table 1: The sample structure

Table 2: Breaking down the variance in incentives

Incentive measure	Standard error (euros/month)	Exogenous	w	w_m
E_1	270	88%	0%	12%
E_2	1,000	47%	51%	2%
E_3	90	45%	24%	31%

Thus E_1 refers to a woman who does not work before or after a birth, E_2 to a woman who stops working, and E_3 to a woman who works before and after. Then we regress employment, w and w_m on all explanatory variables; we simulate the estimated equations to generate the three incentive measures in the absence of idiosyncratic shocks. The results in Table 2 show that shocks to wages account for a large proportion of the variance in the E_i 's across households, which buttresses our approach to identification.

The number of observations with a newborn child (2,617) is fairly small. Therefore we shall have to be parsimonious when specifying the labor supply of women who had a baby in the past year.

3.2 Specification

Our model of fertility is based on the identification analysis of section 2; but we give a self contained presentation here. First consider a woman's decision to work, in the two events of a new birth $F = 1$ or not $F = 0$. The value of working is denoted \mathcal{U}_{LF} , with $L = 1$ in case of work and $L = 0$ otherwise. Given fertility F , the woman works if and only if $\mathcal{U}_{1F} - \mathcal{U}_{0F}$ is positive, and

she achieves a utility level

$$\bar{V}_F = \max(\mathcal{U}_{0F}, \mathcal{U}_{1F}).$$

Now go back in time for the fertility decision. The household decides to have a child when the expected value of $\bar{V}_1 - \bar{V}_0$, given the available information, is positive.

3.2.1 Labor supply

In order to specify the decision to work for a given fertility F , we let μ_F denote its sensitivity to financial incentives ($R_{1F} - R_{0F}$), and d_F the disutility of labor. We use a Probit specification:

$$\Pr(L = 1|F) = \Phi(\mu_F(R_{1F}(w, w^m) - R_{0F}(w^m)) - d_F). \quad (4)$$

In the absence of a new born ($F = 0$), μ_0 is taken to be a linear function of yearly age dummies, whether the woman has a college degree, and family size. μ_1 is equal to μ_0 plus a constant, a linear function of age, and a dummy for families with three kids.

We include income effects in the disutility of labor:

$$d_F = \nu_F R_{0F} + \pi_F.$$

ν_0 has dummies for a child aged 1 or 2, and for families with 2 or 3 kids; ν_1 is equal to ν_0 plus a constant and a dummy for families with three kids. Finally π_F stands for the nonfinancial determinants of labor supply (π_0 has yearly age dummies, diploma broken down in three categories, family size and age of the younger children; π_1 is equal to π_0 plus a constant and a linear combination of five variables: age and its square, a dummy for a college degree, a dummy for an unmarried couple, and a dummy for families with three kids.)

3.2.2 Fertility

Our specification for labor supply is compatible with any specification of the underlying utilities \mathcal{U}_{LF} such that

$$\mathcal{U}_{1F} - \mathcal{U}_{0F} = a_F \left(R_{1F} - R_{0F} - \frac{d_F + \varepsilon_L}{\mu_F} \right),$$

where a_F is positive and the shock on the disutility of labor ε_L is a centered standard normal.

In order to specify the fertility decision, we still need to model the sensitivity parameters a_F and to set the level of the utilities \mathcal{U}_{LF} . The latter implies specifying the cost of children and the random shocks to the fertility decision.

We specify a_F flexibly. We include a cubic spline of the age of the mother, dummies for a college-educated mother, and dummies for the ages of older children; and we allow all these coefficients to vary freely with the parity of the birth. The cost of children depends on the age of the mother, her education, and family composition. Finally, we use a simple logit specification for the random shocks to fertility.

Note that in terms of the system of equations (1) of the identification section, our specification amounts to

$$U_{LF}(R_{LF}) = a_F \left(R_{LF} - \frac{d_F}{\mu_F} \mathbf{1}(L = 1) \right) - C_1 \mathbf{1}(F = 1)$$

along with

$$\zeta_{1F} = a_F \frac{\varepsilon_L}{\mu_F}$$

and with c_1 a standard logistic.

3.3 Estimation results

We built a microsimulation program that represents the French tax benefit system and allows us to compute household disposable income for any values of wages and household characteristics. As usual, we do not observe the wages

of non-working women. For simplicity of implementation, we use a two-step strategy. We first estimate a model of labor supply coupled with a Mincer wage equation, accounting for the selection due to the minimum wage (as in Laroque and Salanié (2002)). Then we take the simulated wages as given to estimate the joint model of labor supply and fertility, using maximum likelihood¹³.

All the tables are gathered at the end of the paper. Incomes are measured in hundreds of euros. To fix ideas, the median monthly household disposable income in our sample (as measured by the observed value of R_{LF}) is 1,950 euros.

3.3.1 Labor supply

We specify and estimate in a flexible way the utilities of working from (4), which are identified through the choice of a standard normal disturbance.

The coefficients of μ_F , which measures the sensitivity of the labor supply of women to financial incentives to work ($R_{1F} - R_{0F}$), are shown in Table 6. For women without babies, who are in relatively larger number, μ_0 is represented as the sum of indicator functions, of annual age, of college education and of the size of the family (parity 2 or 3). To interpret the magnitude of the coefficients, one must bear in mind the fact that on average half of the women of the sample work: the shape of the normal distribution implies that at the average point a change of 0.1 in the deterministic utility $u_{1F} - u_{0F}$ corresponds to a change of 4 points in the probability of working.

The sensitivity to financial incentives seems to be slightly decreasing with age, from around 0.27 to 0.20 for a low skilled woman without children (Table 6). An increase of 100 euros in the monthly income of such a household would increase the woman's probability of working by approximately 10 points. The sensitivity to incentives is slightly lower for educated women and/or for women with two children, as a_0 is reduced by .04 in both cases.

Given that only 15% of women in our sample have a baby in a given year, we specify μ_1 as the sum of μ_0 , a linear function of age, and an indicator

¹³We use the same labor supply specification in both steps.

of a family with two children. The age term is insignificant, but the mean estimate of these additional terms is negative; this suggests that the labor supply of women with a newborn is less sensitive to financial incentives. On the other hand, having two children increases $\mu_1 - \mu_0$ by 0.08: so after the birth of a third child women are more sensitive to financial incentives.

The coefficient ν_F measures the income effects on labor supply and is presented in Table 7. Leisure is a normal good for women without children: an increase of 100 euros per month of other income (i.e. of her husband income) reduces their probability of working by 4 points on average. The effect is reversed for women with a baby, as already seen in Laroque and Salanié (2002): one interpretation is that having a husband with a larger income allows the woman to hire child care and continue with her professional career. Indeed for a second born, an increase of other income of 100 euros per month increases the probability of getting a job by 7 points on average.

Finally the coefficient π_F describes the constant term. For flexibility, π_0 is represented as the sum of indicator functions, of age dummies, of diplomas and of household structure (see Table 8). The constant is modified for women with a young baby, by adding a quadratic polynomial in age, with three indicators of education and family structure. The results are by and large unsurprising:¹⁴ high diplomas make women more likely to work, having two children and/or a baby deters from work.

Given the estimated utility functions, we can compute the “expected marginal rates of substitution” M_F , equal to $(\partial V_F / \partial w) / (\partial V_F / \partial w^m)$, for each observation of our sample. Figure 2 plots M_1 against M_0 for parities 0, 1 and 2. Our identifying assumption 5 only requires that the points do not cluster too close to the diagonal, conditional on the exogenous variables that affect fertility. The graphs above prove that M_0 and M_1 are rather different for a large subset of values of x . Moreover, the great majority of observations has $M_0 > M_1$ as expected. Of course, the only way to check that we have enough variation for identification is to estimate the fertility model.

¹⁴The only unusual sign is the negative one that appears in front of the ‘unmarried’ indicators. Women living with a partner, without being legally married, appear to have a lower labor supply, all other things equal. We do not have a ready explanation for this result.

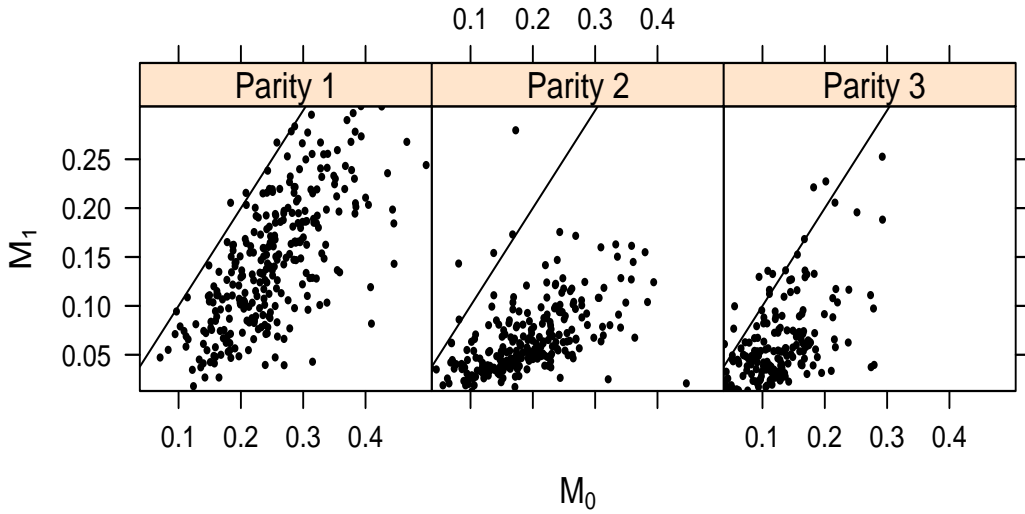


Figure 2: Scatter plots of M_1 against M_0

3.3.2 Fertility decisions

As already mentioned, fertility is specified as a logit, and the specification and coefficients differ across parities. The probability of fertility can be rewritten as

$$\Pr(F = 1) = \frac{1}{1 + \exp[-(a_1(V_1 - V_0) - (a_0 - a_1)V_0 - C_1)]}, \quad (5)$$

with C_1 the cost of children.

The results are presented in Tables 9, 10 and 11 respectively for parities 1, 2 and 3.

The main variable of interest is the difference of the expected utilities when having a baby or not, $(V_1 - V_0) = \Delta V$. This variable is interacted with age education and family composition, but since 15% of the women in the sample have a baby, we do not have enough degrees of freedom to use a full set of age dummies. Instead we posit a spline function of age, with three knots located at the 15%, 50% and 85% quantiles of the age distribution of the sample for each parity.




Figure 3: Semi elasticity of the probability of fertility with respect to $V_1 - V_0$

We renormalized the estimated utility indices V_F so that they can be interpreted in hundreds of euros per month. Thus given the functional form of the logit, a change of 100 euros per month in $(V_1 - V_0)$ changes the probability Q of fertility by $Q(1 - Q)\theta$, where θ is the coefficient of ΔV that appears in the tables. Take for example the change in Q for an average 24 year old woman without children associated with an increase of financial incentives of 100 euros. Since Q is equal to 0.17 on average for these women, the increase in the probability of a birth is $0.17 \times 0.83 \times 0.21 = 0.030$; these 3.0 points represent a 17% increase in the fertility rate. Another way to phrase is that this 0.17 figure measures the average semi-elasticity of fertility with respect to financial incentives for this group of women:

$$\frac{\partial \ln Q}{\partial (V_1 - V_0)} = 0.17.$$

Given the very diverse household compositions and incomes, the semi-

elasticities vary a lot in the population. This appears clearly in figure 3, which plots the distribution of the estimated semi-elasticity with respect to $(V_1 - V_0)$ in our sample¹⁵. The first striking fact is perhaps that for second births (the middle panel), the estimated semi-elasticities cluster around zero: only 30% are larger than 0.1, and 20% are in fact negative (albeit small). Thus financial incentives seem to play little role in the decision to have a second child.

On the other hand, almost all estimated semi-elasticities are positive for the third birth, and half of them are larger than 0.3. These are large effects: according to our estimates, an increase of 100 euros per month would bring the probability of a third birth from 0.08 to 0.11 for a woman with two children at age 32 (which is the median age for a third birth.)

The picture is more mixed for first births, as shown in the left panel of figure 3. At the median age of 28 for a first birth, almost all semi-elasticities are positive, and half of them are above 0.16. The effect is twice larger on average for younger women; on the other hand, it becomes negative for the majority of women of age 30 and older. On the whole population of childless women, 80% of estimated semi-elasticities are positive; and the median is 0.11, which is much smaller than for third births.

As the discussion above makes clear, the estimated effects are quite sensitive to age. The profile of the estimated coefficients of financial incentives $(V_1 - V_0)$ for each age is represented on Figure 4, along with the 95% confidence band. Other factors also play an important role. Note for instance from Table 9 that childless college graduates seem to be more sensitive to financial incentives than women without a college education. This is partly compensated by the fact that college graduates of course have children later, when the baseline effect turns negative; still, they tend to have a slightly larger semi-elasticity than other women. This is somewhat surprising; and Table 11 confirms it for the third birth.

¹⁵Figures 3, 6, 4 and 5 focus on women of ages 24 to 36. Only 15% of first births occur before age 24, and 15% of third births occur after age 36.




Figure 4: Estimated effects of $(V_1 - V_0)$ on fertility

Our specification also allows for income effects, through the interactions of V_0 with age, education and family composition. The age profile, keeping fixed education and family composition, is represented on Figure 5. The overall effect on fertility is negative for women with lower education, and close to zero for those with a college education. Figure 6 shows a striking symmetry to Figure 3: for each parity, the estimated densities are close mirror images with respect to the zero vertical line.

A quadratic variant. We checked the robustness of the results by estimating a slightly different model, involving a quadratic income term in the utility function:

$$\tilde{U}_{LF}(R) = U_{LF}(R) + d_F R^2.$$




Figure 5: Estimated effects of V_0 on fertility




Figure 6: Semi elasticity of the probability of fertility with respect to V_0

We allow the coefficient d to differ according to F , as well as to parity. Compared with the reference specification, there are six new coefficients to be estimated. We also modify the normalization to keep close with the initial V_F into

$$\tilde{V}_F = \frac{1}{\mu_F + d_F R_{0F}} E \max[\tilde{U}_{0F}, \tilde{U}_{1F}].$$

The point estimates of all the coefficients are positive, pointing towards a convex utility function. The coefficients are largest for parity 2, smaller for parity 1, and close to zero for parity 3. While the sum $\mu_F + d_F R_{0F}$ stays positive, some of the μ_F coefficients become negative. A likelihood test indicates that the quadratic terms are highly significant: the difference between the two values of the log-likelihood is 19.6, for six degrees of freedom. However we have not pursued this specification as the improvement of the fit concerns labor supply, while the adjustment of fertility deteriorates.

The importance of financial incentives. We estimated the reference model, constraining all the thirty-six coefficients of the value functions V_F in the fertility equation to be zero. The log-likelihood is reduced by 86.9 points, which is significant at any level. Thus in statistical terms, financial incentives are highly significant in the fertility decision.

3.4 Testing our specification

We found that we needed a large number of covariates x to control for fertility decisions, with parity-specific fertility equations, age of the woman, ages of existing children. . . This makes it infeasible to go beyond parametric methods given the available data, even though the structure of our model and the continuous nature of the key variables R_{LF}, w, w^m naturally suggests using an average derivative estimation method.

On the other hand, we can test the assumptions that underlie our identification of fertility responses. We shall use the U_{LF} obtained at the first stage to estimate V_F up to the scale factors a_F . This allows us first to compute the expected marginal rates of substitution M_0 and M_1 and to check that they

differ from each other. Moreover, formula (K) implies that

$$\ln \left(-\frac{\frac{\partial Q}{\partial V_1}}{\frac{\partial Q}{\partial V_0}} \right) \quad \text{has a zero cross-derivative in } (V_0, V_1).$$

Our specification in equation (5) satisfies the zero cross-derivative property by construction. To test it, we extend the specification by including the term $(a_1V_1 - a_0V_0)^2$ and additional quadratic terms (V_0V_1) and V_1^2 ; and we test that the coefficients of the two last terms are zero. This is similar to the test of the index assumption in Ichimura and Taber (2002).

When we include six such coefficients (one of each for each of the three parities), the test that they are jointly zero is rejected, with a p -value of 1%. However, this is due to parity 2: the null hypothesis is not rejected for parities 1 and 3.

Our approach to identification relies heavily on the exogeneity of wages in the fertility equation. A simple way to test it is to compute generalized log-wage residuals from the estimated Mincer equation¹⁶; and to plug them as an additional covariate in the fertility equation. If wages are exogenous, the corresponding coefficient should be zero. We obtained the results in Table 3. None of the three coefficients is significant at the usual 5% level¹⁷. Moreover, the point estimates are rather small. Take for instance a woman with one child whose estimated probability of a second birth is 0.2. Then simple computations show an increase of one standard deviation of the log-wage residual reduces this probability by 0.004. The effects for parities 1 and 3 are even smaller.

We also ran the converse test, including generalized fertility residuals in the Mincer equation. The additional variable again has a very small and non-significant estimated coefficient (0.003, with standard error 0.004.)

¹⁶Recall that generalized residuals are equal to the usual residuals for working women, and to the conditional expectation of residuals for non-working women.

¹⁷Likelihood ratio tests of the joint significance of the three additional variables confirm it: the p -value is 0.60.

Parity	Estimate	Standard error
1	0.016	(0.165)
2	-0.075	(0.047)
3	-0.031	(0.061)

Table 3: Testing exogeneity of wages

4 Simulating a child credit

Parity	Unconditional			Employment tested		
	Fertility	Work	Work/newborn	Fertility	Work	Work/newborn
1	2.6	-0.6	-1.2	1.6	-1.7	-7.5
2	1.1	-0.5	-1.0	0.7	-1.3	-5.4
3	3.0	-0.5	-0.9	2.5	-0.8	-5.1
All	2.2	-0.5	-1.0	1.6	-1.2	-6.0

Table 4: Changes in fertility and participation associated with a child credit policy

Another way to describe the magnitude of the sensitivity of fertility to financial incentives is to mimic a reform of family benefits. We approximate such a reform by increasing the R 's in case of birth by 150 euros per month. This additional transfer is akin to a universal child subsidy, neither means tested nor employment tested. This would be very costly—about 0.3% of GDP—and therefore we take it as an upper bound of realistic policy changes.

We can evaluate the effects of such a reform by adding it to our tax-benefit computation module, simulating its effects on labor supply, recomputing V_0 and V_1 , and simulating fertility according to the above model. We contrast the situations where the transfer is conditional on stopping work (R_{01} is increased by 150 euros, not R_{11}) or unconditional (both R_{01} and R_{11} are changed). The aggregate results, shown in points of fertility and labor force participation, appear in Table 4, with box-plot graphs of the impacts on Figures 7 and 8.

The exercise confirms that the effects of financial incentives on fertility are sizeable. The unconditional child credit increases the fertility rate in




Figure 7: Child credit: impact on labor supply

the sample by 2.2 points, from 15.5% to 17.7%. The increase is pronounced for the first and third children, smaller for parity 2. It is associated with a modest decrease of .5 point in labor participation, from 49.1% to 48.6%.

When the child credit is conditioned on not working, the increase in fertility is only 1.6 points, which is still large. The brunt of the reduction bears on first births: indeed women without children have the largest participation rate, 63.1%, to be compared with 47.8% for women with one child and 37.3% for women with two children. Unsurprisingly, the overall reduction in participation is larger than when the child credit is unconditional, -1.2 point instead of -0.5. This effect is of course concentrated on women with a newborn, for which it is huge: their participation goes down by 6.0 points¹⁸.

¹⁸This is in line with the results found in France after the extension of the *Allocation parentale d'éducation* to the second born in 1994: its value amounted to 450 euros, and it may have reduced participation of this group by 15 points.




Figure 8: Child credit: impact on fertility

Conclusion

Our paper adopts an approach that differs from much of the literature on evaluation. This choice is driven by the specific features of the problem we study. Given that fertility behavior depends so crucially on ages (of parents and of existing children), and that drastic reforms are few and far between, most of the useful variation is to be found on cross section data. It consists of the differences in fertility- and employment-contingent household incomes generated by differences of (potential, conditional) wages and the specificities of the French tax-benefit system. Nonlinearities of tax-benefit schemes have been much used to study labor supply, see e.g. Blundell, Duncan, and Meghir (1998) and Laroque and Salanié (2002). Indeed, these variations are sufficient to nonparametrically identify labor supply behavior; but they are not enough to identify the joint labor supply and fertility decisions. We therefore introduce an additional assumption, which states that labor supply shocks, conditional on realized fertility, are independent of the fertility shock. The

implied separability of the functional form can be tested along lines similar to Ichimura and Taber (2002).

This orthogonality assumption is natural here, since the fertility decision usually takes place about one year before the conditional choice of labor supply. It yields a nonparametric generalization of the nested logit specification; this may be useful in other applications where the conditional independence assumption is economically reasonable. As an example, work in progress by Chiappori, Salanié, Tillman and Weiss uses a similar approach to examine endogamy by ethnic group and education level.

In order to reduce the incidence of possible unobserved “taste for children” factors, we used a very detailed specification by age and parity. The overall procedure appears to work well and to yield robust results. Fertility is quite sensitive to financial incentives for the first and the third birth, but it hardly responds to them for the second birth.

References

- APPS, P., AND R. REES (2004): “Fertility, Female Labour Supply and Public Policy,” *Scandinavian Journal of Economics*, 106, 745—763.
- 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.
- 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.
- BLUNDELL, R., A. DUNCAN, AND C. MEGHIR (1998): “Estimating Labor Supply Responses Using Tax Reforms,” *Econometrica*, 66, 827–861.
- CHONÉ, P., D. LEBLANC, AND I. ROBERT-BOBÉE (2004): “Offre de travail féminine et garde des enfants,” *Economie et prévision*, 162.
- CIGNO, A. (1986): “Fertility and the Tax-Benefit System: A Reconsideration of the Theory of Family Taxation,” *Economic Journal*, 96, 1035–1051.
- COHEN, A., R. DEHEJIA, AND D. ROMANOV (2007): “Do Financial Incentives Affect Fertility?,” Discussion Paper 13700, NBER.
- 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.
- 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., 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.
- ICHIMURA, H., AND C. TABER (2002): “Semiparametric Reduced-Form Estimation of Tuition Subsidies,” *American Economic Review*, 92, 286–292.
- KEANE, M., AND K. WOLPIN (2007): “The Role of Labor and Marriage Markets, Preference Heterogeneity and the Welfare System on the Life Cycle Decisions of Black, Hispanic and White Women,” Discussion paper, University of Pennsylvania.
- 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.
- LALIVE, R., AND J. ZWEIMUELLER (2009): “How does Parental Leave Affect Fertility and Return to Work? Evidence from Two Natural Experiments,” *Quarterly Journal of Economics*, 124, 1363–1402.

- 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*, 50, 423–450.
- 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.
- PIKETTY, T. (2005): “L’impact de l’allocation parentale d’éducation sur l’activité féminine et la fécondité en France, 1982-2002,” in *Histoires de familles, histoires familiales*, ed. by C. Lefèvre. INED.
- 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.
- TODD, P., AND K. WOLPIN (2005): “Assessing the impact of a school subsidy program in Mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility,” *American Economic Review*, 96, 1384–1417.

Wage equation		
variable	estimate	std error
Age at end of study	0.118	0.011
id. squared	-2.072	0.272
Time since end of study	0.051	0.003
id. squared	-1.046	0.103
Diploma 1	0.664	0.016
Diploma 2	0.459	0.013
Diploma 3	0.283	0.012
Diploma 4	0.149	0.010
Diploma 5	0.142	0.013
1997 survey	7.237	0.118
1998 survey	7.217	0.118
1999 survey	7.221	0.118
σ	0.272	0.002

Table 5: The wage equation

WILLIS, R. (1973): “A New Approach to the Economic Theory of Fertility Behavior,” *Journal of Political Economy*, 81, S14–S64.

WOLPIN, K. (1984): “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality,” *Journal of Political Economy*, 92, 852–874.

Appendix

The estimation results for the wage equation appear in Table 5. They are very similar to those in our earlier papers and book.

μ_0		
Variable	estimate	std error
20-22	0.29	0.10
23	0.21	0.05
24	0.30	0.05
25	0.26	0.03
26	0.25	0.03
27	0.26	0.03
28	0.22	0.02
Age 29	0.23	0.02
30	0.21	0.02
31	0.18	0.02
32	0.22	0.02
33	0.22	0.02
34	0.19	0.02
35	0.20	0.02
36	0.18	0.02
37	0.20	0.02
38	0.19	0.02
39	0.23	0.02
College	-0.03	0.01
Parity 2	-0.01	0.01
Parity 3	-0.05	0.01

$$\mu_1 = \mu_0 + \underset{(.08)}{.16} - \underset{(.00)}{.01} \text{Age} + \underset{(.03)}{.09} \mathbf{1} \text{ Parity 3}$$

Table 6: The coefficients of financial incentives in labor supply

ν_0		
Variable	estimate	std error
Constant	.06	.01
Child aged 1	-.01	.01
Child aged 2	-.02	.01
Parity 2	-.00	0.01
Parity 3	-.03	0.01

$$\nu_1 = \nu_0 - \underset{(.01)}{.03} + \underset{(.02)}{.03} \mathbf{1} \text{ Parity 3}$$

Table 7: Income effects in labor supply

π_0		
Variable	estimate	std error
20-22	1.48	0.65
23	0.71	0.33
24	1.32	0.31
25	0.87	0.23
26	0.75	0.23
27	0.88	0.21
28	0.51	0.18
Age 29	0.56	0.17
30	0.37	0.16
31	0.17	0.16
32	0.45	0.16
33	0.47	0.17
34	0.25	0.16
35	0.27	0.16
36	0.20	0.14
37	0.35	0.15
38	0.44	0.15
39	0.69	0.17
1998 Survey	-0.13	.03
1999 Survey	-0.22	.03
Dip. 1	-0.25	.10
Dip. 2	-0.58	.09
Dip. 3	-0.52	.08
Child. ≤ 3	0.33	0.10
id. ≤ 6	0.05	0.03
Parity 2	0.08	0.12
Parity 3	0.39	0.14
Unmarried	-0.04	0.03

$$\pi_1 = \pi_0 + \underset{(1.90)}{3.64} + \underset{(0.07)}{0.02} \mathbf{1} \text{ College} + \underset{(0.32)}{0.49} \mathbf{1} \text{ Parity 3} + \underset{(0.06)}{0.22} \mathbf{1} \text{ Unmarried} - \underset{(0.12)}{0.15} \text{ Age} + \underset{(0.18)}{0.11} \frac{\text{Age squared}}{100}$$

Table 8: Constant terms in labor supply

Variable	Estimate	Standard Error
ΔV , age 24	0.21	0.07
ΔV , age 28	0.07	0.08
ΔV , age 32	-0.16	0.07
V_0 , age 24	-0.04	0.02
V_0 , age 28	-0.00	0.02
V_0 , age 32	0.01	0.02
ΔV , college graduate	0.17	0.08
V_0 , college graduate	-0.00	0.02
age 24	0.36	0.28
age 28	0.73	0.30
age 32	0.96	0.35
unmarried	0.47	0.65
unmarried \times age	0.01	0.02
college graduate	2.58	0.84
college graduate \times age	-0.09	0.03

Table 9: Fertility Effects for First Birth

Variable	Estimate	Standard Error
ΔV , age 26	-0.08	0.05
ΔV , age 30	0.06	0.05
ΔV , age 34	0.00	0.05
V_0 , age 26	-0.04	0.02
V_0 , age 30	-0.04	0.02
V_0 , age 34	-0.03	0.02
ΔV , college graduate	-0.00	0.06
V_0 , college graduate	0.05	0.02
ΔV first born is 1	0.19	0.09
ΔV first born is 2	0.11	0.06
ΔV first born is 3	0.13	0.06
V_0 first born is 1	0.00	0.02
V_0 first born is 2	0.01	0.02
V_0 first born is 3	0.01	0.02
age 26	-0.15	0.35
age 30	0.65	0.31
age 34	1.47	0.33
unmarried	2.74	0.62
unmarried \times age	-0.08	0.02
college graduate	1.65	0.90
college graduate \times age	-0.04	0.03
first born is 1	2.32	0.48
first born is 2	0.69	0.34
first born is 3	0.02	0.36
first born is 4	-0.90	0.13
first born is 5	-0.64	0.15
first born is 6	-0.21	0.17

Table 10: Fertility Effects for Second Birth

Variable	Estimate	Standard Error
ΔV , age 27	0.10	0.08
ΔV , age 32	0.14	0.06
ΔV , age 36	0.01	0.07
V_0 , age 27	-0.12	0.03
V_0 , age 32	-0.05	0.02
V_0 , age 36	-0.05	0.02
ΔV , college graduate	0.18	0.10
V_0 , college graduate	0.08	0.02
ΔV one child is 1	0.62	0.55
ΔV one child is 2	0.37	0.31
ΔV one child is 3	-0.03	0.07
V_0 one child is 1	-0.01	0.04
V_0 one child is 2	-0.03	0.02
V_0 one child is 3	-0.01	0.02
age 27	0.44	0.61
age 32	2.33	0.50
age 36	2.49	0.51
unmarried	4.27	1.00
unmarried \times age	-0.14	0.03
college graduate	0.20	1.72
college graduate \times age	-0.05	0.05
one child is 1	-2.07	1.30
one child is 2	-0.29	0.63
one child is 3	-0.82	0.50
one child is 4	-0.36	0.18
one child is 5	-0.41	0.18
one child is 6	-0.30	0.20
children between 7 and 10	-0.30	0.17

Table 11: Fertility Effects for Third Birth