Life cycle employment and fertility across institutional environments

(Article begins on next page)

17 April 2024

Contents lists available at ScienceDirect

European Economic Review

Life cycle employment and fertility across institutional environments ☆

Daniela Del Boca [a], Robert M. Sauer [b,*]

[a] *Department of Economics and Collegio Carlo Alberto, University of Turin, Italy*
[b] *Division of Economics, School of Social Sciences, University of Southampton, 58 Murray Building, Southampton SO17 1BJ, UK*

A R T I C L E   I N F O

A B S T R A C T

In this paper, we formulate a dynamic utility maximization model of female labor force participation and fertility choices and estimate approximate decision rules using data on married women in Italy, Spain and France. The estimated decision rules indicate that first-order state dependence is the most important factor determining female labor supply behavior in all three countries. We also find that cross-country differences in state dependence effects are consistent with the order of country-level measures of labor market flexibility and child care availability. Counterfactual simulations of the model indicate that female employment rates in Italy and Spain could reach EU target levels were French social policies to be adopted in those countries.

1. Introduction

The growth in women's participation in the labor market, especially among women with children, has been one of the most important economic and social phenomena of the last half century. The large scale movement of women into the labor market since the end of World War II has occurred in many different countries. However, the level of female employment rates across countries is still far from having converged, and the influence of social policies on female employment rates is not clearly understood. This has raised serious policy concerns, particularly in Europe, where the European Union (EU) has set quantitative targets for higher female employment rates for all member states.[1]

In order to try and better understand what underlies cross-country differences in female labor force participation rates, we formulate a general dynamic utility maximization model of female labor supply behavior and fertility choices, and estimate the approximate decision rules of the model separately for married women in Italy, Spain and France. The main focus is on measuring the differential relative importance of state dependence and unobserved heterogeneity in country-specific decision rules, and establishing a connection between the differential relative importance and variation across

---

* Corresponding author. Tel.: +44 2380592528.
E-mail address: r.m.sauer@soton.ac.uk (R.M. Sauer).

[1] At the Lisbon summit in March 2000, the European Council stated that all member states should set quantitative targets for higher employment rates in line with EU targets. These were set at 70% for total employment and 60% for women's employment, to be reached by the year 2010. In 2001, intermediate targets of 67% (total) and 57% (for women) were set to be reached by 2005.

countries in social policies. With this purpose in mind, we limit the set of countries in the analysis to only those with "similar" cultural characteristics—i.e., Italy, Spain and France. This helps distinguish social policy variation across countries from confounding factors related to culture, such as attitudes towards gender roles.

The reason for focusing on the differential relative importance of state dependence and unobserved heterogeneity in female work and fertility choices is that past research on female labor force participation has repeatedly shown that persistence is an important aspect of the labor supply decisions of married women (see, e.g., Heckman and Willis, 1977; Heckman, 1981; Nakamura and Nakamura, 1985; Eckstein and Wolpin, 1989). Persistence in participation status may be due to state dependence which arises from human capital accumulation or the costs of searching for a new job. The costs of searching for a new job are, in turn, affected by social policies such as the extent of employment regulation and the availability of child care. However, persistence can also be accounted for by permanent unobserved heterogeneity that reflects differences in mostly immutable preferences for work and/or productivity in the labor market. If unobserved heterogeneity is not properly accounted for in estimation, one may obtain spurious state dependence effects and make faulty inferences about the importance of adjustment costs, social policies, and the institutional environment.

Although several recent studies have also concentrated on disentangling state dependence from permanent unobserved heterogeneity in female labor supply (see e.g., Hyslop, 1999; Carrasco, 2001), to the best of our knowledge, there is no previous work that analyzes the differential relative importance of these factors across countries. Thus, no previous studies have examined the hypothesis that institutions governing social policies are important underlying sources of cross-country differences in state dependence. Institutions which make it more costly to adjust employment levels from one period to the next should generate more persistence and state dependence in female labor supply.

The approximate decision rules that we estimate indicate that state dependence, as opposed to unobserved heterogeneity, is clearly the most important factor determining persistence in labor market participation in all three countries. We also find that the order of state dependence effects across countries is correlated with the order in aggregate measures of labor market flexibility and child care availability. This is consistent with the existence of important differences in institutional environments. It also suggests that employment and child care policies, which affect participation adjustment costs, are additional causes of state dependence and hence cross-country variation in the level of female employment rates.

The estimated decision rules are also used to perform counterfactual simulations. The simulations show that female employment rates in Italy and Spain could reach EU target levels, at least 60% by 2010, were French-like social policies to be adopted in those countries. Under French parameters, Italian and Spanish female participation rates substantially converge towards the higher French female participation rate of 68%. We find that Italian participation rates increase from 53% to 63%, and Spanish female participation rates rise dramatically from 35% to 62%.

One caveat for our results is that they are based on approximate decision rules rather than exact ones. Adopting an exact solution approach would have been much more computationally intensive, but would also have better incorporated cross-equation and forward-looking restrictions implied by the dynamic decision problem. Thus, exact decision rules may look very different from approximate ones.

The rest of this paper is organized as follows. In the next section, we provide a brief background on the relationship between female labor market participation and fertility choices that motivates our model of joint decision-making. In Section 3, we describe the data. Section 4 outlines the life cycle model of labor market participation and fertility decisions. Section 5 discusses estimation of approximate decision rules. Section 6 presents the estimation results and assesses model fit. Section 7 correlates the estimated state dependence effects with aggregate measures of social policies, and reports the results of counterfactual simulations. The last section of the paper summarizes and concludes.

2. Background

There is a vast literature that examines the relationship between fertility and female labor force participation. A negative effect of the number of children on female labor supply is often found. But the effect may not be causal because women with stronger preferences for motherhood may also be those with lower unobservable skills and motivation in the labor market. The endogeneity of fertility has been addressed in the past by looking at sources of unplanned births, e.g., the presence of twins (Rosenzweig and Wolpin, 1980), and variation in the availability and cost of contraceptives (Rosenzweig and Schultz, 1985). Angrist and Evans (1998) suggest using the sibling-sex composition as an instrument for fertility outcomes. However, this latter approach is not practical with European data since the number of women with at least two children is typically very small.

Instead of postulating and exploiting sources of exogenous variation in birth outcomes, there are a number of studies that attempt to directly test for the exogeneity of fertility within simple labor supply models. For example, Mroz (1987) tests the sensitivity of the parameters of the labor supply equation of married women with respect to a number of assumptions, including the exogeneity of fertility. Conditional on participation, he finds that fertility is exogenous to women's labor supply. However, using panel data and controlling for individual effects, Jakubson (1988) arrives at opposite conclusions. Hyslop (1999), following Browning (1992) and Chamberlain (1984), tests for the exogeneity of fertility via a discrete choice correlated random effects model. His results indicate that when dynamic factors such as state dependence or serial correlation are excluded, fertility is endogenous. However, in dynamic specifications with either first-order state

dependence or AR(1) serial correlation, he finds no evidence against the exogeneity of fertility hypothesis. Directly testing for the exogeneity of fertility has yielded very mixed results.

An additional strand in the literature looks at the effect of fertility within more explicit behavioral models of life cycle labor supply. Within this strand, there are both static lifetime models that assume fertility profiles are exogenous (e.g., Heckman and MaCurdy, 1980), and sequential (dynamic) models that treat fertility choices as predetermined (e.g., Eckstein and Wolpin, 1989). There are also dynamic models that explicitly take into account the contemporaneous jointness of fertility and labor supply decisions (e.g., Moffitt, 1984; Hotz and Miller, 1988; Del Boca, 2002; Francesconi, 2002; Keane and Wolpin, 2006). In this paper, we follow this latter approach by formulating a dynamic programming model of joint labor market participation and fertility decisions. However, following Keane and Wolpin (2002), we do not structurally estimate exact decision rules but rather estimate approximate decision rules. We also build on the work of Hyslop (1999) and Keane and Wolpin (2002) by developing a tight connection between the estimation of approximate decision rules and a dynamic discrete choice model with a complex error structure. In addition, we employ a relatively new simulated maximum likelihood (SML) algorithm that corrects for possible biases due to classification errors in reported participation and birth outcomes (see Keane and Wolpin, 2001; Keane and Sauer, 2005).

3. Data

The data used in this study are drawn from the European Community Household Panel (ECHP). The ECHP is a standardized multi-purpose longitudinal survey designed and coordinated by the Statistical Office of the European Communities (Eurostat). The survey is conducted annually on a representative panel of households in each member state of the EU. The survey covers a wide range of topics on living conditions such as income, employment, poverty and social exclusion, housing, health and migration. The unit of analysis in the ECHP is the family, and information is gathered on all individuals within the household that are 16 of age or older. However, it is also possible to recover information on family members that are younger than 16.

The ECHP began in 1994 (wave 1), following a two-wave pilot survey. Wave 1 covered about 60,000 households and 130,000 individuals in all 12 EU member states (Belgium, Denmark, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Portugal, Spain and the UK). Austria joined the survey in 1995 (wave 2), Finland joined in 1996 (wave 3) and Sweden joined in 1997 (wave 4). The last year the ECHP was administered was 2002 (wave 9). Eurostat terminated the project in 2003 and replaced it with a new instrument, the EU-SILC (Statistics on Income and Living Conditions), in order to focus more attention on the determinants of poverty and social exclusion.

We analyze ECHP data from Italy, Spain and France, between the years 1994 and 2000 (waves 1–7). Birth outcomes in 2001 (wave 8) of the survey are not observed due to a censoring problem. Hence, 2001 is excluded from the estimation sample. We limit the set of countries to Italy, Spain and France because they differ substantially in terms of social policies but are generally thought to have "similar" cultural environments (e.g., attitudes towards gender roles). Differences in social policies across these three countries will be empirically demonstrated in Section 7 below.[2]

The sample from each country that we analyze contains women who are between the ages of 21 and 45, who are continuously married or cohabitant with partners, that are continuously employed throughout the sample period, and who have complete employment and fertility histories. These restrictions are common in the female labor supply literature and we adhere to them. They exclude women who might still be enrolled in school or retired and who have a low probability of being fecund. The restriction that all women have complete employment and fertility histories excludes women in the ECHP who could not be contacted or refused to cooperate subsequent to being interviewed in wave 1, as well as women who entered the survey after wave 1 (see Nicoletti and Peracchi, 2004). The final estimation sample contains 830 women from Italy, 713 women from Spain and 993 women from France observed over seven years (1994–2000). The extent of sample selection generated by the sample exclusion restrictions is roughly similar in each country.[3]

Table 1 presents descriptive statistics by country. The means in the table are calculated by first computing average values over the seven-year panel for each woman, and then calculating averages over all women in the country-sample. The statistics show large differences in female education levels between countries. For example, in Italy only 8% of the women have tertiary education levels, while in Spain and France the proportions are much higher, 20% and 28%, respectively. The proportion of women whose youngest child in the household is three years old or younger is similar in Italy and Spain but relatively higher in France. France also has the highest mean annual partner earnings (in thousands of 2001 Euros), female labor market participation rate, and annual birth rate. The raw data display a positive correlation across countries in work and fertility outcomes.

In order to see how work and fertility choices change over time, Figs. 1 and 2 display the annual labor market participation and birth rates over the sample period in each country. Fig. 1 illustrates that participation rates over the sample period are always highest in France, second highest in Italy, and lowest in Spain. In the latter part of the 1990s,

---

[2] Limiting the number of countries included in the analysis also substantially reduces computational burden.

[3] Employment and birth outcomes are assigned according to whether the woman is working at the time of the interview and whether she has had a child in the year of the interview. There is no variable indicating whether a woman is on maternity/parental leave. We expect that women who are on leave state that they work since the relevant employment question is about "normal" activity status.

Table 1
Descriptive statistics by country

	France	Italy	Spain
Age	36.64	37.81	37.27
	(5.79)	(5.19)	(5.51)
Secondary education	.73	.52	.49
	–	–	–
Tertiary education	.28	.08	.20
	–	–	–
Youngest child 0–3	.15	.11	.12
	–	–	–
Youngest child >3	.73	.81	.81
	–	–	–
Husband's earnings	20.10	16.61	17.34
	(20.09)	(7.35)	(9.73)
Birth rate	.06	.04	.05
	–	–	–
Employment rate	.66	.48	.34
	–	–	–
N	993	830	713

Note: Individual means for each women over seven years are calculated, and then the means are averaged over all women in the country sample. Husband's earnings are in thousands of 2001 Euros. Standard deviations of continuous variables are in parentheses.
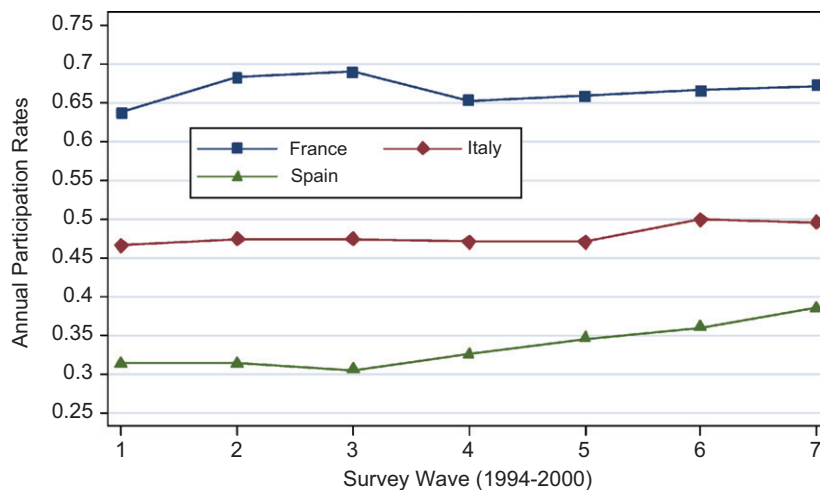

Fig. 1. Annual participation rates by country (1994–2000). Note: Survey wave 1 corresponds to the year 1994 and survey wave 7 corresponds to the year 2000.

participation rates in Spain begin to converge to those in Italy, while Italian participation rates remain mostly constant. French participation rates fluctuate a bit more than those in Italy and Spain.

Fig. 2 graphically illustrates that birth rates are consistently highest in France over the sample period. Spanish birth rates start out quite high, exceeding those in France as well as Italy in 1995, but fall relatively rapidly over time (as participation rates increase). Towards the end of the sample period Spanish birth rates roughly equalize with those in Italy, and both are nearly half the birth rates in France. The birth rates in each country fall over time as the women in the sample age.

Persistence in female labor supply is illustrated in Table 2, which displays the distribution of years worked over the sample period, separately by country. In Italy, the proportion of women who always work and who never work are quite
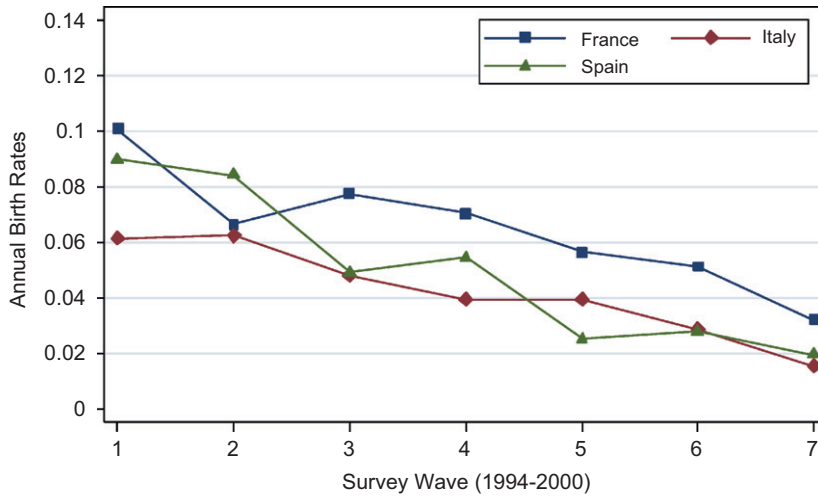
Fig. 2. Annual birth rates by country (1994–2000). *Note*: Survey wave 1 corresponds to the year 1994 and survey wave 7 corresponds to the year 2000.

Table 2
Distribution of panel years worked by country (column percentages)

Years worked	France	Italy	Spain
0	.180	.395	.492
1	.045	.053	.086
2	.039	.042	.060
3	.050	.037	.048
4	.058	.032	.036
5	.068	.034	.027
6	.099	.037	.032
7	.459	.369	.219
	1.000	1.000	1.000

similar, 37% and 39%, respectively. These two modal points account for more than three-quarters of the distribution. In Spain, relatively less women always work than in Italy, 22%, but many more women never work, 49%. The percentages in France are quite different: A larger proportion of women always work, 46%, and a smaller proportion never work, 18%. In all three countries, the two modal points in the distribution are at the corners.[4]

Strong persistence in female labor supply can also be discerned from Table 3, which presents average rates of transition between employment state in year $t-1$ and employment state in year t. The diagonals of the matrices in each country show that both persistence in participation, and persistence in nonparticipation, is relatively highest in Italy. In France, it is much more common for women to move from nonparticipation to participation, while in Spain there is more movement from participation to nonparticipation. The patterns in the transition matrices are broadly consistent with a negative association between participation rates and persistence.[5]

4. Model

In this section we specify a general dynamic utility maximization model of female labor supply and birth decisions, and derive the model's approximate decision rules. In the next section, we discuss the details of the approximation, the estimation technique employed, and identification.

4.1. Basic structure

Consider a married women i who maximizes remaining discounted lifetime utility by choosing, in each year t, whether or not to participate in the labor market, h_{it}, and whether or not to give birth, b_{it}. We abstract from the part-time, full-time

---

[4] The descriptive statistics displayed in this section for France, Italy and Spain can be easily compared to similar statistics for the US reported in Hyslop (1999), and for Germany reported in Croda and Kyriazidou (2004).

[5] Azmat et al. (2004) have shown that in Italy and Spain, where more women are unemployed relative to men, females are more likely to move from employment to unemployment and less likely to enter from unemployment to employment, compared to males.

Table 3
Employment transitions by country (row percentages)

France	Work in t	
Work in $t-1$	0	1
0	.839	.161
1	.073	.927

Italy	Work in t	
Work in $t-1$	0	1
0	.932	.068
1	.064	.936

Spain	Work in t	
Work in $t-1$	0	1
0	.928	.072
1	.112	.888

(hours) margin and possible fertility complications by assuming that planned live births can occur with certainty within the same year t.

Remaining lifetime utility at time t for woman i is given by

$$V_{it}(S_{it}) = \max_{\{h_{it}, b_{it}\}} \mathrm{E}\left[\sum_{t=\tau}^{T} \delta^{\tau-t} U_{it}(\cdot)\Big| S_{it}\right] \tag{1}$$

where τ is the theoretical start of the decision process, T is the end of the decision horizon, δ is the subjective discount factor, and S_{it} is the state space at time t. S_{it} contains all the factors, exogenous and endogenous, known to the individual at time t affecting current utility or the probability distribution of future utility. $U_{it}(\cdot)$ is the utility flow at time t and $V_{it}(S_{it})$ is the value function.

The maximization problem in (1) can be cast in terms of alternative specific value functions, $V_{it}^{bh}(S_{it})$, each of which follow Bellman's equation, i.e.,

$$V_{it}(S_{it}) = \max[V_{it}^{00}(S_{it}), V_{it}^{10}(S_{it}), V_{it}^{01}(S_{it}), V_{it}^{11}(S_{it})]$$
$$V_{it}^{hb}(S_{it}) = U_{it}(\cdot) + \delta\mathrm{E}(V_{i,t+1}(S_{i,t+1})|h_{it}, b_{it}, S_{it}), \quad t < T$$
$$= U_{iT}(\cdot), \quad t = T \tag{2}$$

In $t < T$, the value of each choice combination is the current period utility flow plus discounted expected lifetime utility in time $t+1$. The state space in time $t+1$, $S_{i,t+1}$, is updated according to the laws of motion in the endogenous state variables to be described below. In $t = T$, there is no future component to the alternative specific value functions.

The utility flow $U_{it}(\cdot)$ in (1) and (2) is assumed to be

$$U_{it}(\cdot) = U_{it}(h_{it}, b_{it}, C_{it}, h_{i,t-1}, N_{it}, L, K, \varepsilon_{it}^{h}, \varepsilon_{it}^{b}) \tag{3}$$

where C_{it} is current period consumption, $h_{i,t-1}$ is last period's employment status, N_{it} is the current stock of children, L is an aggregate measure of labor market flexibility, K is an aggregate measure of child care services, and ε_{it}^{h} and ε_{it}^{b} are time t preference shocks to working and giving birth, respectively. Note that the presence of $h_{i,t-1}$, L and K in (3) explicitly recognizes that utility is not intertemporally separable in leisure (see Hotz et al., 1988), and that there are economy-wide factors (social policies) that affect work and birth choices.

To further justify the set of arguments in (3), consider the following linear utility function as an example:

$$U_{it}(\cdot) = C_{it} + (\gamma_{0h} + \gamma_{1h}C_{it} + \gamma_{2h}h_{i,t-1} + \gamma_{3h}N_{it} + \varepsilon_{it}^{h})h_{it} + (\gamma_{0b} + \varepsilon_{it}^{b})b_{it} \tag{4}$$

where γ_{2h} is a linear function of the aggregate measures L and K,

$$\gamma_{2h} = \gamma_{2h0} + \gamma_{2h1}L + \gamma_{2h2}K \tag{5}$$

The parameters in (4) and (5) have straightforward interpretations. γ_{0h} is the marginal utility of working in year t, reflecting the utility cost of work effort. γ_{1h} measures the extent to which the marginal utility of consumption varies with participation status. γ_{2h} captures the utility cost saved from not having to adjust last period's participation status. γ_{3h} is the

marginal utility of an additional child when participating in the labor market relative to not participating. γ_{0b} is the marginal utility (or disutility) of giving birth in year t.[6]

Note that γ_{2h} is a key state dependence parameter and is further modeled in (5) to be a (deterministic) function of the institutional environment. This specification explicitly recognizes that the cost of adjusting last period's participation status is influenced by job search costs, which are in turn influenced by the extent of labor market flexibility L (e.g., the supply of part-time jobs and temporary employment contracts), and child care costs, which are in turn influenced by the supply of child care services K.

It is maintained throughout the analysis that L and K are taken as given by the individual decision-maker. This can be justified by considering that the supply of part-time jobs, and other flexible employment contracts, as well as the supply of child care services in the countries that we examine, are mostly determined within the political process. Lags and inefficiencies in public sector responses to changes in demand, as well as slowly changing institutions in general, reduce the extent of the endogeneity problem. Although we will not use proxies for L and K directly in estimation, due to a lack of sufficient individual and time-series variation, Eq. (5) is important for clarifying the logic of our indirect approach to inferring the effect of social policies on female labor supply behavior. In short, social policies influence the cost of adjusting participation status, and hence the degree of state dependence in female participation.

The per-period budget constraint in the lifetime utility maximization model is simply specified as

$$y_{it}^f h_{it} + y_{it}^m = C_{it} + C_n N_{it} \tag{6}$$

where y_{it}^f is the woman's labor market earnings in year t, y_{it}^m is the partner's labor market earnings in year t (a proxy for "transitory" nonlabor income), and C_n represents the goods-cost per child. y_{it}^m will not be further modeled but will rather appear directly in estimation (in logs). The average of y_{it}^m over the sample period (a proxy for "permanent" nonlabor income) will affect the distribution of individual preference effects (see below). This helps partially account for the endogeneity of nonlabor income.[7]

Let the wife's labor market earnings y_{it}^f be further determined by

$$y_{it}^f = g(x_i', H_{it}, r_{it}, t, \varepsilon_{it}^f) \tag{7}$$

where $g(\cdot)$ is a general function of the covariate vector x_i', accumulated actual work experience during the sample period H_{it}, region of residence in the current period r_{it}, a time effect t, and a productivity shock ε_{it}^f.

The vector of covariates x_i' in (7) contains proxies for accumulated human capital prior to the start of the sample period. The start of the sample period for each individual i is denoted as $t = \tau_i$. More specifically, the vector x_i' is $x_i' = \{E_{i\tau_i}, a_{i\tau_i}, a_{i\tau_i}^2\}$ where $E_{i\tau_i}$ is the education level by the start of the sample period, and a_{τ_i} is the woman's age upon entry to the sample (a proxy for potential experience). As with nonlabor income, the endogeneity of education level is partially accounted for in estimation by allowing education to affect the distribution of individual preference effects. Note that the first period of observed data for each woman $t = \tau_i$, which is the age of entry into the sample, will generally not be the start of the theoretical decision process $t = \tau$. We deal with this initial conditions problem by using Heckman's approximate solution, which we will describe in more detail in the next section.

The laws of motion for the endogenous state variables in the dynamic program are

$$\begin{aligned} N_{i,t+1} &= N_{it} + b_{it} \\ H_{i,t+1} &= H_{it} + h_{it} \end{aligned} \tag{8}$$

where the initial conditions are $N_{it} = N_{i\tau_i}$ and $H_{i\tau_i} = 0$. The initial (reported) stock of children at the start of the sample period is augmented by one for each birth during the sample period. Similarly, accumulated actual work experience during the sample period is augmented by one for each period a woman chooses to work. Note that both N_{it} and H_{it} are potential sources of state dependence in addition to adjustment costs via $h_{i,t-1}$.

4.2. Additional error structure

The work and birth preference shocks ε_{it}^h and ε_{it}^b in the general utility function in (3) are given additional structure as follows:

$$\begin{aligned} \varepsilon_{it}^h &= \alpha_i^h + \xi_{it}^h \\ \varepsilon_{it}^b &= \alpha_i^b + \xi_{it}^b \end{aligned} \tag{9}$$

---

[6] Linear utility functions similar to the one specified in (4) have been empirically implemented in the related literature (see, e.g., Eckstein and Wolpin, 1989; Francesconi, 2002; Keane and Wolpin, 2002). Utility functions of this form are generally identifiable in structural estimation.

[7] The budget constraint could also include L and K. For example, one could specify two female wage functions (full-time and part-time) where the level of those wages, and probability of obtaining offers, are a function of L. Similarly, the goods-cost per child C_n could be a function of K. Because we are only estimating approximate decision rules, whether L and K enter the utility function or the budget constraint, or both, will not impact on the estimation results.

where α_i^j, $j = h, b$ are time-invariant individual preference effects, and ξ_{it}^j, $j = h, b$ are transitory preference shocks. The time-invariant effects capture permanent unobserved heterogeneity related to immutable (or slowly changing) preferences for a career and for a family. The α_i^j's induce serial correlation in the error terms.

Rather than imposing a parametric distribution on the α_i^j's, we assume that they are random with a discrete distribution that has three mass points. That is,

$$\alpha_i^h = \theta_1^h A_{1i} + \theta_2^h A_{2i}$$
$$\alpha_i^b = \theta_1^b A_{1i} + \theta_2^b A_{2i} \tag{10}$$

where A_{1i} is a dummy variable for unobserved "type" 1, A_{2i} is a dummy variable for unobserved "type" 2, and A_{0i} is a dummy for unobserved "type" 0, which is the base type. The structure in (10) allows for distinct career and family preference effects through the coefficients θ_1^j and θ_2^j, $j = h, b$. The exact form of the discrete mixing distribution will be given below.[8]

In order to capture possible persistence in unobserved wage offers, the productivity shock ε_{it}^f is also assumed to be serially correlated. ε_{it}^f is given an AR(1) structure,

$$\varepsilon_{it}^f = \rho \varepsilon_{i,t-1}^f + \nu_{it} \tag{11}$$

where ν_{it} is a serially uncorrelated transitory error component.

As will become more apparent below, because we estimate approximate decision rules, and do not use data on female wages in estimation, the empirical estimates will be robust to several variations on the error structure described above. The important point is that the error structure above provides a more explicit theoretical grounding for introducing random effects and AR(1) serially correlated errors into the estimation of approximate decision rules.

4.3. Approximate decision rules

Instead of adopting specific functional forms for $U_{it}(\cdot)$ and $g(\cdot)$, and numerically solving the dynamic programming model for exact decision rules, we characterize the approximate decision rules of the general model. An important reason for adopting the approximate solution technique in this context is that the computational burden of calculating an exact solution for each country would be severe.[9]

The first step in the characterization of approximate decision rules involves substituting the female wage into the budget constraint and then substituting the budget constraint into the choice specific $U_{it}(\cdot)$'s.[10] This yields a distinct choice-specific state space S_{it}^{hb} for each combined work and birth alternative.

For example, consider the basic structure of the model, the linear utility function in (4), and the state dependence equation in (5). Substitution yields the following S_{it}^{hb}'s:

$$S_{it}^{00} = \{N_{it}, y_{it}^m\}$$
$$S_{it}^{10} = \{N_{it}, h_{i,t-1}, h_{i,t-1}L, h_{i,t-1}K, x_{it}', H_{it}, r_{it}, t, y_{it}^m, \varepsilon_{it}^h, \varepsilon_{it}^b, \varepsilon_{it}^f\}$$
$$S_{it}^{01} = \{N_{it}, y_{it}^m, \varepsilon_{it}^h, \varepsilon_{it}^b\}$$
$$S_{it}^{11} = \{N_{it}, h_{i,t-1}, h_{i,t-1}L, h_{i,t-1}K, x_{it}', H_{it}, r_{it}, t, y_{it}^m, \varepsilon_{it}^h, \varepsilon_{it}^b, \varepsilon_{it}^f\} \tag{12}$$

Note that the distinct S_{it}^{hb}'s in (12) does not imply that the decision rules for each choice combination will be a function of a unique subset of covariates. Rather, each decision rule will depend on the entire state space S_{it} because the value of any particular choice combination is computed by comparing it to the values of all other choice combinations. S_{it} is the union of the S_{it}^{hb}'s, where

$$S_{it} = \{N_{it}, h_{i,t-1}, h_{i,t-1}L, h_{i,t-1}K, x_{it}', H_{it}, r_{it}, t, y_{it}^m, \varepsilon_{it}^h, \varepsilon_{it}^b, \varepsilon_{it}^f\} \tag{13}$$

To see this more clearly, and without loss of generality, consider the myopic version of the model. In the myopic version there is no future component to the alternative specific value functions for $t \leqslant T$, i.e., $\delta = 0$.[11] Denote $d_{it}^{hb} = 1$ if alternative (h_{it}, b_{it}) is chosen and $d_{it}^{hb} = 0$, otherwise. Utility maximization implies the following comparison of utility flows in

---

[8] In preliminary estimations, three mass points were found to fit the data better than two. Four points did not produce a significant increase in the value of the log-likelihood function.

[9] See the discussion in Keane and Wolpin (1997) on different approaches to structural estimation and especially Keane and Wolpin (2002). See Buchinsky and Gotlibovski (2006) for a different type of approximation.

[10] The overwhelming majority of related studies on labor market participation substitute out for the wage and do not incorporate observed wage data in estimation (see, e.g., Magnac, 2000). Eckstein and Wolpin (1989) is an exception.

[11] In the dynamic version of the model, each alternative specific value function at time t has S_{it} as an argument simply because the expected maximum future returns component compares the values of all choice combinations in the future.

each period t:

$$d_{it}^{00} = 1 \quad \text{when } U_{it}(0,0) - U_{it}(j,k) > 0, \quad (j,k) \neq (0,0) \text{ or } F_{it}^{00}(S_{it}) > 0$$
$$d_{it}^{10} = 1 \quad \text{when } U_{it}(1,0) - U_{it}(j,k) > 0, \quad (j,k) \neq (1,0) \text{ or } F_{it}^{10}(S_{it}) > 0$$
$$d_{it}^{01} = 1 \quad \text{when } U_{it}(0,1) - U_{it}(j,k) > 0, \quad (j,k) \neq (0,1) \text{ or } F_{it}^{01}(S_{it}) > 0$$
$$d_{it}^{11} = 1 \quad \text{when } U_{it}(1,1) - U_{it}(j,k) > 0, \quad (j,k) \neq (1,1) \text{ or } F_{it}^{11}(S_{it}) > 0 \tag{14}$$

Estimating approximate decision rules is accomplished by choosing a functional form for $\Pr(d_{it}^{hb} = 1) = \Pr(F_{it}^{hb}(S_{it}) > 0)$. The contribution of the optimization model is to provide a firm theoretical grounding for the common set of covariates, S_{it}, that affect each choice probability. Note that the model we develop implies that both $h_{i,t-1}$ and H_{it} should be included in estimation. This is in contrast to previous research that imposes more arbitrary restrictions on the distributed lag process (e.g., a first-order Markov process as in Hyslop, 1999).

As will be seen more clearly in the next section, we estimate the approximate decision rules via a bivariate probit model. A bivariate probit model is an appropriate statistical representation of the choice problem only under certain conditions. In particular, the utility specification cannot include a birth and work interaction term. In addition, the approximation approach will be robust to other alternative structures of the model (e.g., the incorporation of job offer probabilities and/or permanent layoff probabilities) only if these alternative structures do not imply additional observable covariates should be included in the state space, or there should be a different error structure than the one specified above.

5. Estimation

In this section, we expand upon estimation of the approximate decision rules, describe the SML algorithm employed, and briefly discuss identification.

5.1. Estimating approximate decision rules

Estimation of the approximate decision rules of the optimization model proceeds by specifying $\Pr(F_{it}^{hb}(S_{it}) > 0)$ in the following way:

$$\Pr(d_{it}^{00} = 1) = \Pr(F_{it}^{00}(S_{it}) > 0) = \int_{-\infty}^{0} \int_{-\infty}^{0} f(H_{it}^*(S_{it}), B_{it}^*(S_{it})) \, dH_{it}^* \, dB_{it}^*$$
$$\Pr(d_{it}^{10} = 1) = \Pr(F_{it}^{10}(S_{it}) > 0) = \int_{0}^{\infty} \int_{-\infty}^{0} f(H_{it}^*(S_{it}), B_{it}^*(S_{it})) \, dH_{it}^* \, dB_{it}^*$$
$$\Pr(d_{t}^{01} = 1) = \Pr(F_{it}^{01}(S_{it}) > 0) = \int_{-\infty}^{0} \int_{0}^{\infty} f(H_{it}^*(S_{it}), B_{it}^*(S_{it})) \, dH_{it}^* \, dB_{it}^*$$
$$\Pr(d_{t}^{11} = 1) = \Pr(F_{it}^{11}(S_{it}) > 0) = 1 - \sum_{hb \in \{(00),(10),(01)\}} \Pr(F_{it}^{hb}(S_{it}) > 0) \tag{15}$$

where $f(\cdot)$ is the bivariate normal density. The choice probabilities in (15) are those of a bivariate probit model.[12]

To be consistent with bivariate probit choice probabilities, $H_{it}^*(S_{it})$ and $B_{it}^*(S_{it})$ in (15) must be distributed bivariate normal. This is accomplished by specifying $H_{it}^*(S_{it})$ and $B_{it}^*(S_{it})$, for $t > \tau_i$, as linear functions of S_{it},

$$H_{it}^*(S_{it}) = \alpha_{0h} + \alpha_{1h} y_{it}^m + \alpha_{2h} N_{it} + \alpha_{3h} h_{i,t-1} + \alpha_{4h} H_{it} + \alpha_{5h} x_i' + \alpha_{6h} r_{it} + \alpha_{7h} t$$
$$\qquad + \alpha_{8h} A_{1i} + \alpha_{9h} A_{2i} + \alpha_{10h} \varepsilon_{i,t-1}^f + \nu_{it} + \eta_{it}^h$$
$$B_{it}^*(S_{it}) = \alpha_{0b} + \alpha_{1b} y_{it}^m + \alpha_{2b} N_{it} + \alpha_{3b} h_{i,t-1} + \alpha_{4b} H_{it} + \alpha_{5b} x_i' + \alpha_{6b} r_{it} + \alpha_{7b} t$$
$$\qquad + \alpha_{8b} A_{1i} + \alpha_{9b} A_{2i} + \alpha_{10b} \varepsilon_{i,t-1}^f + \nu_{it} + \eta_{it}^b \tag{16}$$

where the η_{it}^j's, $j = h, b$, not introduced earlier, are assumed to be independent of ν_{it}, and distributed bivariate normal with zero means and unit variances. The η_{it}^j's are error components that capture deviations of the approximate decision rules from the exact ones, which could arise from linearization. They also contain the omitted variables L and K, which are common to $H_{it}^*(S_{it})$ and $B_{it}^*(S_{it})$. Note that L and K are the only state variables that appear in the decision model but not in the approximate decision rules. They are omitted from the approximate decision rules because, as mentioned earlier, their direct effects are not identifiable, due to a lack of sufficient individual and time-series variation in aggregate proxies for labor market flexibility and child care services. Hence, L and K can be thought of as components of the η_{it}^j's.

Because the choice probabilities in (15) are conditional on unobserved type, they must be weighted by type probabilities in order to obtain unconditional likelihood contributions for each individual. The three mass point probabilities, which

---

[12] Keane and Wolpin (2002) adopt a specification for $\Pr(F_{it}^{hb}(S_{it}) > 0)$ such that a multinomial logit is generated. Bivariate models are more parsimonious than multinomial models, but they are more restrictive as mentioned above (see Weeks and Orme, 2004). Thus, bivariate and multinomial models could yield different approximate decision rules.

constitute the mixing distribution of the individual effects, are specified as

$$
\Pr(A_{1i}) = L_1^A(y_{ip}^m, E_{i\tau_i})
$$
$$
\Pr(A_{2i}) = L_2^A(y_{ip}^m, E_{i\tau_i})
$$
$$
Pr(A_{0i}) = 1 - \Pr(A_{1i}) - \Pr(A_{2i}) \tag{17}
$$

where $L_k^A(\cdot)$, $k = 1, 2$, are logistic in form with different coefficients for each k. The logistic function ensures that each mass point probability remains between zero and one during iterations. As mentioned earlier, the type probabilities are functions of nonlabor income and education. More specifically, "permanent" nonlabor income y_{ip}^m enters the type probabilities, rather than transitory nonlabor income y_{it}^m, where y_{ip}^m is the average of y_{it}^m over the sample period. Entering the endogenous variables y_{ip}^m and $E_{i\tau_i}$ into the type probabilities is motivated by the correlated random effects model (Chamberlain, 1984).[13]

In order to explicitly address the initial conditions problem in the dynamic bivariate probit, we make use of the Heckman approximate solution (Heckman, 1981). The Heckman approximate solution to the initial conditions problem entails specifying $H_{it}^*(S_{it})$ and $B_{it}^*(S_{it})$ functions in the initial sample period $t = \tau_i$ without lagged endogenous state variables, and with errors that are correlated with the error terms in $t > \tau_i$.

The $H_{i\tau_i}^*(S_{i\tau_i})$ and $B_{i\tau_i}^*(S_{i\tau_i})$ functions take the form

$$
H_{i\tau_i}^*(S_{i\tau_i}) = \lambda_{0h} + \lambda_{1h}y_{i\tau_i}^m + \lambda_{2h}x_i' + \lambda_{3h}r_{i\tau_i} + \lambda_{4h}A_{1i} + \lambda_{5h}A_{2i} + \nu_{i\tau_i} + \eta_{i\tau_i}^h
$$
$$
B_{i\tau_i}^*(S_{i\tau_i}) = \lambda_{0b} + \lambda_{1b}y_{i\tau_i}^m + \lambda_{2b}x_i' + \lambda_{3b}r_{i\tau_i} + \lambda_{4b}A_{1i} + \lambda_{5b}A_{2i} + \nu_{i\tau_i} + \eta_{i\tau_i}^b \tag{18}
$$

where the correlation between the initial period errors and the errors during the sample period arises through the time-invariant individual effects and the AR(1) serial correlation in $\varepsilon_{i,t-1}^f$ (see Eq. (11)). [14]

Eqs. (15)–(18) establish a close link between estimation of approximate decision rules and estimation of a dynamic bivariate probit model with nonparametric correlated random effects, AR(1) errors, and Heckman's approximate solution to the initial conditions problem. Estimation of the dynamic bivariate probit model above is difficult using classical maximum likelihood techniques. In particular, calculation of choice probabilities with AR(1) serially correlated errors requires multiple integration. For this reason we use SML. The relatively new SML algorithm that we employ is described in the next subsection.

5.2. The SML algorithm

The SML algorithm that we use to estimate the approximate decision rules was originally developed by Keane and Wolpin (2001) for estimating exact decision rules of dynamic programming problems and overcoming missing data problems. In Keane and Sauer (2005), the algorithm is shown to be also useful for estimating more general dynamic discrete choice models. The algorithm has a computational advantage in that it is easier to implement than other SML algorithms (e.g., GHK) since it relies on unconditional simulations of the model. The new algorithm also corrects for biases due to classification error in reported participation and birth choices.

The algorithm proceeds as follows. For each individual, and for each unobserved type, two independent standard normal deviates are drawn in the initial sample period. The standard normal deviates are transformed into bivariate normal deviates using the lower diagonal Choleski factor of the bivariate normal distribution. The standard normal deviates are also used for generating AR(1) serial correlation as in (11).

Given a vector of trial parameters, draws for the error terms, and the covariates reported in the data, $H_{i\tau_i}^*(S_{i\tau_i})$ and $B_{i\tau_i}^*(S_{i\tau_i})$ are calculated according to (18). The simulated participation and birth choices in the initial sample period are then determined by $h_{i\tau_i}^m = I(H_{i\tau_i}^*(S_{i\tau_i}) > 0)$ and $b_{i\tau_i}^m = I(B_{i\tau_i}^*(S_{i\tau_i}) > 0)$, where m represents the replication number. Given initial period simulated choices, the state space is updated according to (8).

Simulated choices for the subsequent six decision periods are then constructed by drawing bivariate normal deviates in each period, computing the values of $H_{it}^*(S_{it})$ and $B_{it}^*(S_{it})$ in (16), and then determining $h_{it}^m = I(H_{it}^*(S_{it}) > 0)$ and $b_{it}^m = I(B_{it}^*(S_{it}) > 0)$. The state space is updated accordingly. This procedure is repeated a total of M times ($M = 750$) for each unobserved type, and each individual, to form type-specific simulated choice sequences $\{\{\{h_{it}^m\}_{t=\tau_i}^T\}_{i=1}^N\}_{m=1}^M$ and $\{\{\{b_{it}^m\}_{t=\tau_i}^T\}_{i=1}^N\}_{m=1}^M$ for each woman in the sample.

Formation of the type-specific likelihood contribution for each woman i entails specifying the probability of a reported choice in period t (h_{it}^* and b_{it}^*) as a function of the simulated choice in period t for replication m (h_{it}^m and b_{it}^m). These

---

[13] Preliminary estimations indicated that separate permanent and transitory nonlabor income effects were hard to empirically identify when entered together into either (16) or (17). There are no other covariates that could be empirically identified in the type probabilities.

[14] The variance of $\nu_{i\tau_i}$ is adjusted so that the productivity shock is stationary. In addition, α_{10b} in (16) is set to zero for purposes of identification. Thus, AR(1) serial correlation affects only $H_{it}^*(S_{it})$.

probabilities are as follows:

$$\pi_{01tm}^h = \Pr(h_{it}^* = 1 | h_{it}^m = 0) = L_1^h(h_{it}^m)$$
$$\pi_{10tm}^h = \Pr(h_{it}^* = 0 | h_{it}^m = 1) = L_2^h(h_{it}^m)$$
$$\pi_{01tm}^b = \Pr(b_{it}^* = 1 | b_{it}^m = 0) = L_1^b(b_{it}^m)$$
$$\pi_{10tm}^b = \Pr(b_{it}^* = 0 | b_{it}^m = 1) = L_1^b(b_{it}^m) \tag{19}$$

where $\pi_{00tm}^j = 1 - \pi_{01tm}^j$ and $\pi_{11tm}^j = 1 - \pi_{10tm}^j$ for $j = h, b$. $L_k^j(\cdot)$, $k = 1, 2$, $j = h, b$, are logistic functions. The choice probabilities in (19) are given a logistic form to ensure that estimated probabilities remain between zero and one during iterations. As is common in the classification error literature, the probability of reporting a particular choice is a function of the current true (simulated) choice but is not directly affected by covariates (see Hausman et al., 1998). π_{01tm}^j and π_{10tm}^j, $j = h, b$, are the classification error rates.

Note that because there are M simulated participation and birth choices in every period t, for each unobserved type of individual, sequences of choice probabilities will be generated for each unobserved type. Specifically, we obtain the sequences $\{\{\pi_{jktm}^h\}_{t=\tau_i}^T\}_{m=1}^M$ and $\{\{\pi_{jktm}^b\}_{t=\tau_i}^T\}_{m=1}^M$ where j denotes the simulated choice h_{it}^m (b_{it}^m) and k denotes the reported choice h_{it}^* (b_{it}^*).

Assuming that classification error in reported participation status is independent of classification error in reported birth outcomes, an unbiased simulator of the type-specific likelihood contribution for each woman i can be formed by calculating the product of ($\pi_{jktm}^h \times \pi_{jktm}^b$) over time, and then averaging these products over the M replications. More formally,

$$\widehat{P}(h_i^*, b_i^* \| A_{1i}, A_{2i}, \theta)$$
$$= \frac{1}{M} \sum_{m=1}^M \prod_{t=\tau_i}^T \left(\left\{ \sum_{j=0}^1 \sum_{k=0}^1 \pi_{jktm}^h I[h_{it}^m = j, h_{it}^* = k] \right\} \left\{ \sum_{j=0}^1 \sum_{k=0}^1 \pi_{jktm}^b I[b_{it}^m = j, b_{it}^* = k] \right\} \right) \tag{20}$$

where h_i^* and b_i^* are vectors containing woman i's reported participation and birth outcomes over the sample period, and θ is the vector of model parameters. Construction of the unconditional likelihood contribution for each woman i entails weighting the conditional likelihood contributions in (20) by the mixing distribution in (17).[15]

Because the sample log-likelihood generated by this procedure is not everywhere smooth in θ, the objective function is maximized using a non-gradient based optimization routine. However, standard errors are computed off of a sample log-likelihood that is smooth in $\widehat{\theta}$, where the smoothing is accomplished by generalizing (20) to include importance sampling weights (see Keane and Sauer, 2005 for details).

5.3. Identification

The main identification issues that arise in this context are (i) separate identification of state dependence, permanent unobserved heterogeneity and AR(1) serial correlation, and (ii) separate identification of classification error rates from all other parameters of the model. We will briefly discuss these two issues in turn.

For purposes of illustration, consider the data on reported participation outcomes only. The strong persistence in participation status, as demonstrated in Tables 2 and 3, could be solely due to serial correlation in the error term (either in the form of random effects or an AR(1) error component). If this were the case, then conditional on the current X (a time-varying covariate), lagged X's would not help to determine current participation status. However, if true state dependence is present, then conditional on the current X, lagged X's will help to determine the current employment outcome, assuming that there is no direct effect of lagged X's. This is the variation in the data which distinguishes state dependence from serial correlation.

Although one can nonparametrically identify separate state dependence and serial correlation effects, it is not possible to distinguish state dependence from different forms of serial correlation (e.g., random effects or AR(1) transitory errors) without parametric assumptions (see Chamberlain, 1984). Moreover, if serial correlation in the errors is incorrectly modeled, one may obtain a spurious state dependence effects. Misspecification of the degree of state dependence can also lead to incorrect inferences about serial correlation. For these reasons, it is important to include different forms of state dependence and serial correlation in the model, as we do in Eq. (16). Note that identification of the nonparametric discrete mixing distribution follows from the analysis in Heckman and Singer (1984).

The conditions required for separate identification of classification error rates from the other parameters of the model are analyzed in a static discrete choice context by Hausman et al. (1998). They show that full identification (as opposed to semi-parametric identification up to scale) requires a nonlinear true choice probability, as well as a monotonicity condition. The corresponding monotonicity condition in our case is $\pi_{01tm}^h + \pi_{10tm}^h < 1$, which says that the sum of participation classification error rates should not exceed one. This condition ensures that the probability a woman reports

---

[15] In estimation, the restriction $\pi_{jktm}^h = \pi_{jktm}^b$ was imposed because it was difficult to separately identify participation and birth classification error rates.

Table 4
Selected SML point estimates and standard errors

	France		Italy		Spain	
	Work (1)	Birth (2)	Work (3)	Birth (4)	Work (5)	Birth (6)
y_{it}^m	−.3623	.0818	−.3782	.0854	−.4556	.0637
	(.0515)	(.0792)	(.0590)	(.0493)	(.0667)	(.1032)
$ychild_{it} < = 3$	−.7366	−1.2369	−.3982	−1.2533	−.5711	−1.2188
	(.1176)	(.1525)	(.1994)	(.2090)	(.1754)	(.2433)
$ychild_{it} > 3$	.0007	−.4963	−.4077	−.5734	−.1195	−.5768
	(.0962)	(.1240)	(.1418)	(.1516)	(.1244)	(.1842)
$E_{i\tau_i}^s$	.7814	.1514	2.0501	.3800	1.1410	.0899
	(.0783)	(.1191)	(.0855)	(.1133)	(.1246)	(.2156)
$E_{i\tau_i}^t$	.9280	.6013	1.3016	.6311	1.2949	.6163
	(.0750)	(.1075)	(.1513)	(.2066)	(.1483)	(.2545)
H_{it}	.0777	−.0474	.2377	−.0872	.2987	−.2325
	(.0424)	(.0466)	(.0444)	(.0492)	(.0520)	(.0713)
$h_{i,t-1}$	1.7499	-.0357	2.1426	.1359	1.8681	.0336
	(.0835)	(.1373)	(.1170)	(.1571)	(.1390)	(.2140)
A_{1i}	1.6432	1.1187	2.0312	1.2444	1.6137	1.0539
	(.0993)	(.1282)	(.1327)	(.2357)	(.1897)	(.3098)
A_{2i}	−1.3122	1.0689	−1.4669	1.0359	−1.3456	1.3129
	(.1172)	(.1409)	(.1383)	(.2354)	(.2062)	(.3158)
$\varepsilon_{i,t-1}^f$	.5933	–	.7778	–	.7361	–
	(.0341)		(.0127)		(.0237)	
$\Pr(A_{1i}), \Pr(A_{2i})$	(.5382,.2237)		(.3300, .5444)		(.4753,.3818)	
π_{01}, π_{10}	(.0648,.0210)		(.0472,.0167)		(.0669,.0241)	
Log-likelihood	−4092.48		−2652.73		−2474.69	
N	993		830		713	

Note: All specifications also include a quadratic in age, year and region dummies

participating in the labor market is increasing in the true probability, say $F(X_{it}'\beta)$, which in turn is increasing in $X_{it}'\beta$. A further discussion on the identification of classification error rates in dynamic models can be found in Keane and Sauer (2006). Keane and Sauer (2006) also demonstrate that large biases in the relative importance of state dependence and unobserved heterogeneity can result if even a small amount of classification error in the data is ignored.

6. Estimation results

6.1. Point estimates and standard errors

Table 4 reports selected SML estimates of the decision rule coefficients in (16), and their standard errors. The first two columns of the table present the results for the French $H_{it}^*(S_{it})$ and $B_{it}^*(S_{it})$ functions. Column (1) shows negative effects of transitory nonlabor income and very young children on labor market participation. Both effects are precisely estimated. More education has a significant positive effect on labor market participation. The state dependence effect associated with lagged participation status has a very strong positive impact, while the state dependence effect arising from accumulated human capital during the sample period is not strong nor precisely measured. The estimated coefficient on $\varepsilon_{i,t-1}^f$ suggests nonnegligible dynamics arising from unobserved productivity shocks. [16]

The coefficients on A_{1i} and A_{2i} indicate a strong and precisely measured influence of permanent unobserved heterogeneity. Type 1 women are more likely to participate in the labor market relative to type 0 women, and type 2 women are less likely to participate in the labor market relative to type 0 women.

The estimates in Column (2) show that the presence of a child in the house (regardless of the age of the child) significantly decreases the propensity to give birth. The youngest child being less than three has a relatively stronger impact. Having achieved a tertiary level of education in France significantly increases the propensity to give birth. There are also important unobserved heterogeneity effects. Type 1 women are more likely to give birth (and work) than type 0 women. Type 2 women are more likely to give birth (but less likely to work) than type 0 women.

The estimated type probabilities for France, shown at the bottom of Columns (1) and (2), indicate that, on average, types 1 and 2 women account for three-quarters of the population, with type 1 individuals being the majority, 53%. Nonlabor income and education have significant effects on a woman's unobserved type (not shown in the table). Nonlabor income

---

[16] Because of a lack of regional relocation in the data, r_{it} in (16) appears in estimation as a set of exogenous region dummies. The time effect t appears as unrestricted year dummies. $E_{i\tau_i}$ is discretized into $E_{i\tau_i}^s$ and $E_{i\tau_i}^t$, representing secondary and tertiary education levels, respectively. N_{it} is discretized into $ychild_{it} \leqslant 3$ and $ychild_{it} > 3$, representing the youngest child being three years of age or less, and the youngest child being older than three.

Table 5
Selected marginal effects and relative importance decomposition

	France		Italy		Spain	
	Work (1)	Birth (2)	Work (3)	Birth (4)	Work (5)	Birth (6)
Marginal effects						
y_{it}^m	−.023	−.000	−.015	.006	−.022	.002
$ychild_{it} \leqslant 3$	−.056	−.154	−.014	−.103	−.027	−.100
$ychild_{it} > 3$	.002	−.078	−.014	−.057	−.004	−.064
$E_{i\tau}^s$	.049	.004	.059	.016	.043	.007
$E_{i\tau}^t$	.055	.057	.047	.022	.053	.022
A_{1i}	.127	.058	.063	.035	.073	.026
A_{2i}	−.146	.055	−.056	.024	−.043	.043
H_{it}	.001	−.003	.002	−.003	.009	−.004
$h_{i,t} - 1$	.656	−.005	.816	.004	.779	−.006
Relative importance decomposition						
$\bar{R}_1^2\,(X\text{'s})$	.060	.117	.128	.075	.136	.064
$\bar{R}_2^2\,(X\text{'s},fe)$	.063	.133	.143	.084	.142	.072
$\bar{R}_3^2\,(X\text{'s},fe, \text{AR}(1))$	.159	.133	.260	.084	.274	.072
$\bar{R}_4^2\,(X\text{'s},fe, A\text{'s}))$	.529	.143	.478	.087	.392	.082
$\bar{R}_5^2\,(X\text{'s},fe, \text{AR}(1), A\text{'s})$	.623	.143	.596	.087	.523	.082
$\bar{R}_6^2(X\text{' s},fe, h_{i,t-1})$	.773	.133	.887	.084	.850	.077
$\bar{R}_7^2(X\text{' s},fe, h_{i,t-1}, H_{it})$	.774	.133	.887	.084	.850	.077
$\bar{R}_8^2(X\text{' s},fe, \text{AR}(1), A\text{' s}, h_{i,t-1}, H_{it})$	.821	.144	.900	.088	.870	.085

Note: The *X*'s include nonlabor income, children dummies, age, age-squared, and education dummies. The fixed effects (*fe*) are year and region dummies.

and secondary education increase the probability of being types 1 and 2 relative to type 0, while tertiary education decreases the types 1 and 2 probabilities. The significance of these variables implies that nonlabor income and education are indeed endogenous, and that it was important to account for this in estimation (albeit indirectly). It is also important to note that the main influence of nonlabor income on birth outcomes is indirect; it is through the type probabilities.

The estimated classification error rates for France are small in magnitude but are significantly different from zero. The probability of reporting participation in the labor market, when the true state is nonparticipation, is estimated to be .065. The probability of reporting nonparticipation, when the true state is participation, is .021.[17] The classification error rates for reported birth outcomes are constrained to be equal to the classification error rates for reported participation outcomes.

Columns (3)–(6) display the corresponding results for Italy and Spain, respectively. The point estimates are qualitatively similar to those reported in Columns (1) and (2) for France. Interestingly, the same pattern of coefficients on the type dummies are obtained in both the work and birth equations. Previous participation status is very important, as is serial correlation in productivity shocks. Note that the point estimate of the effect of lagged participation is strongest in magnitude in Italy and second strongest in Spain (marginal effects will be examined below).

In contrast to France, in both Italy and Spain accumulated actual work experience significantly decreases the propensity to give birth. There are also differences in the type proportions across countries, indicating substantially different distributions of permanent unobserved heterogeneity. Nonlabor income and education significantly affect the type probabilities in Italy and Spain in the same directions as in France. The classification error rates for Italy and Spain are small in magnitude as in France, but are also significantly different from zero.

6.2. Marginal effects and relative importance decomposition

The top panel of Table 5 reports selected marginal effects, corresponding to the point estimates reported in Table 4. The marginal effects are calculated in the following way. First, stochastic elements of the model are drawn from their estimated distributions. Second, participation and birth outcomes are simulated according to the decision rules. The third step is to run separate linear regressions of simulated participation and birth outcomes on the variables on the right-hand side of (16). The coefficients produced by these linear regressions are simulated marginal effects.

The marginal effects in Table 5 clearly illustrate the overriding importance of state dependence arising through lagged participation status. Having participated in the labor market in the previous period increases the probability of participating in the current period by 66 percentage points in France, 82 percentage points in Italy, and 78 percentage

---

[17] Keane and Sauer (2006) obtain classification error rates for reported labor market participation that are similar in magnitude.

points in Spain. The marginal effects for permanent unobserved heterogeneity (the type dummies) are much smaller than the marginal effects of lagged participation status, but are generally second in importance.

The marginal effects in the birth equations indicate that age (not shown) and the presence of children are the strongest determinants of births in each country. The effects of nonlabor income and education are weak. Permanent unobserved heterogeneity is much more important than previous participation status or actual accumulated work experience.

An additional way to measure the relative importance of the factors determining labor market participation and birth outcomes is to add different sets of variables to the linear regressions on the simulated data, and examine the changes in the adjusted R-squared. The results of this exercise are reported in the bottom panel of Table 5. In what follows, we will discuss the results for the participation equation only.

The base specification for the analysis contains nonlabor income, age of youngest child dummies, age, age-squared and education dummies (referred to as the X's). The adjusted R-squared, $\overline{R}_1^2$, for the base specification is quite low; it is .06 for France, .13 for Italy and .14 for Spain. Adding region and year dummies to the base specification (referred to as the fixed effects (fe's) yields a $\overline{R}_2^2$ which is only slightly higher. Adding the simulated serially correlated productivity shock (referred to as AR(1)) has a more substantial effect. $\overline{R}_3^2$ reaches .16 for France, .26 for Italy, and .27 for Spain. Adding permanent unobserved heterogeneity (referred to as the A's) instead of the AR(1) shock yields a much more substantial rise. $\overline{R}_4^2$ increases to .53 for France, .48 for Italy and .39 for Spain. Adding the A's and the AR(1) shock simultaneously reveals that permanent unobserved heterogeneity is relatively more important than the AR(1) shock.

In order to assess the relative importance of state dependence, the next specification adds lagged participation status to the X's and the fe's. Notice that $\overline{R}_6^2$ jumps dramatically to .77 for France, .89 for Italy and .85 for Spain. This again illustrates the overriding importance of first-order state dependence relative to permanent unobserved heterogeneity and AR(1) serial correlation. The next experiment adds accumulated experience during the sample period to lagged participation status, the X's and the fe's. There is virtually no change in the adjusted R-squared. In the final row of the panel, all of the variables appearing in the approximate decision rules are included in the regressions. $\overline{R}_8^2$ increases by relatively little over $\overline{R}_6^2$ in all three countries.

The results of our relative importance decomposition exercise illustrate that first-order state dependence (adjustment costs) is by far the most important factor in explaining female labor supply decisions in all three countries. This is true even after controlling for higher order state dependence (accumulated human capital). There is also a clear ordering in the adjustment cost effect across countries. First-order state dependence is relatively more important in Italy and Spain than in France. We will further discuss this cross-country ranking in adjustment cost effects after examining model fit.

6.3. Model fit

The reliability of our results rests to a certain extent on the ability of the estimated decision rules to fit the raw data. We assess model fit by comparing descriptive statistics produced from simulations of the model, using the estimated decision rules, to the descriptive statistics displayed earlier. Fig. 3 compares the annual participation rates across countries in the raw data (as in Fig. 1) to predicted participation rates. As can be seen clearly in the graph, the predicted rates track the actual rates very well. The order in the level of participation rates across countries is reproduced, as are the trends over time within each country.

Fig. 4 compares the annual birth rates in the raw data (as in Fig. 2) to predicted annual birth rates. Again, the order in the level of birth rates across countries is reproduced as are the time/age trends in each country. However, the fit is relatively
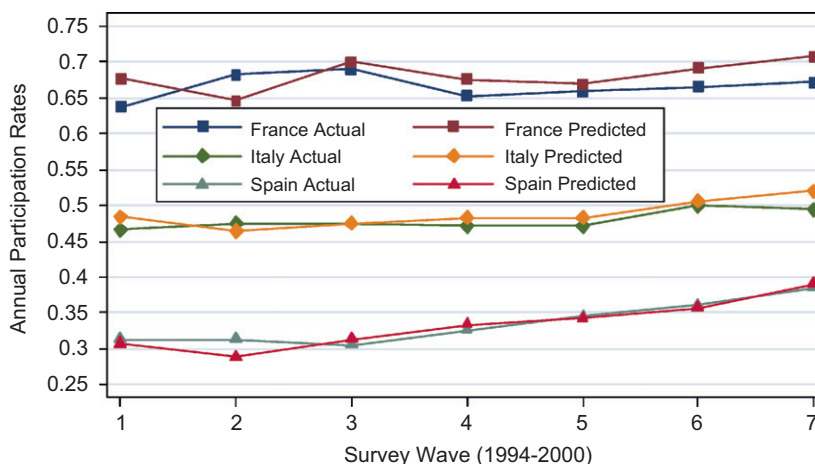

Fig. 3. Annual participation rates by country: actual vs. predicted (1994–2000). *Note*: Survey wave 1 corresponds to the year 1994 and survey wave 7 corresponds to the year 2000.

Fig. 4. Annual birth rates by country: actual vs. predicted (1994–2000). *Note*: Survey wave 1 corresponds to the year 1994 and survey wave 7 corresponds to the year 2000.
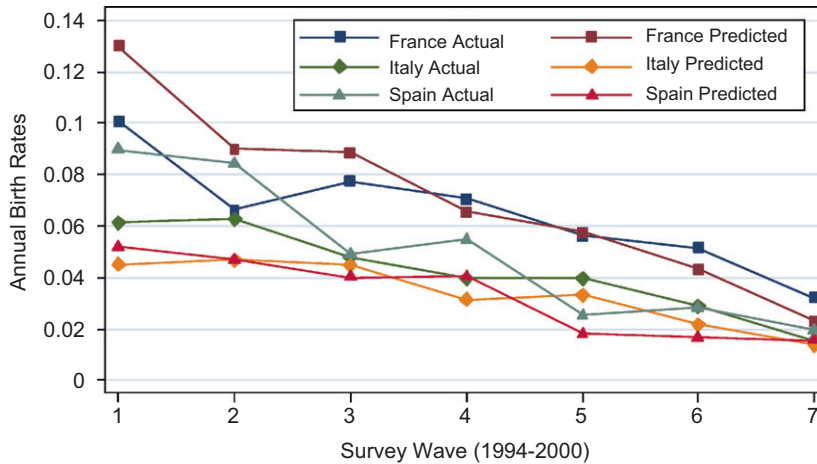
Table 6
Distribution of panel years worked by country: actual vs. predicted (column percentages)

Years worked	France		Italy		Spain	
	Actual	Predicted	Actual	Predicted	Actual	Predicted
0	.180	.195	.395	.423	.492	.562
1	.045	.039	.053	.033	.086	.044
2	.039	.033	.042	.035	.060	.030
3	.050	.041	.037	.024	.048	.028
4	.058	.035	.032	.030	.036	.035
5	.068	.061	.034	.027	.027	.037
6	.099	.034	.037	.010	.032	.013
7	.459	.561	.369	.418	.219	.251

less good than in Fig. 3. Fig. 4 indicates that the model slightly under-predicts birth rates in each country. Very low birth rates is probably an important reason why it is more difficult to fit these rates, as opposed to rates of labor force participation.

Table 6 displays the distribution of years worked during the sample period (as in Table 3) and the corresponding predicted distribution. Notice that the predicted distribution reproduces the two modal points at zero and seven years worked. The predicted distribution also reproduces the correct ordering of modal points across countries. The proportion working zero years in France and Italy is captured accurately, but is slightly over-estimated in Spain. The proportion working all seven years is captured well in Spain and Italy, but is a bit too high in France.

Table 7 compares actual and predicted transition matrices. In France, the diagonal rates are a bit too high. However, the ordering of transition frequencies is exactly reproduced. In Italy, the fit is better. The ordering of transition frequencies is again exactly reproduced. In Spain, as opposed to France and Italy, the model predicts too many transitions out of employment and hence the employment to employment transition rate is too high. Note that since the model over-predicts employment persistence in each country, state dependence effects might also be over-estimated. However, because we focus on relative state dependence effects across countries, over-estimation is less of a concern.

Overall, the fit of the model to the data is quite good. This suggests that the estimated approximate decision rules are reasonable proxies for the exact decision rules, and reliable benchmarks for the policy simulations we perform below.

7. Discussion

The estimated decision rules reveal that state dependence, arising from the cost of adjusting participation status, is stronger in Italy and Spain (where labor market participation rates are relatively lower) compared to France. Using Eq. (5) as motivation, in the next subsection we note the connection between aggregate proxies for labor market flexibility and child care availability and the cross-country pattern of first-order state dependence effects estimated in the model.

In the second subsection, we use the estimated approximate decision rules to perform additional simulations. We quantify the effect of the institutional environment on participation and fertility choices by examining how these choices would change if women in one country were to face the estimated decision rule parameters of another country.

Table 7
Employment transitions by country: actual vs. predicted (row percentages)

France

Actual				Predicted		
	Work in t				**Work in t**	
Work in $t-1$	0	1		Work in $t-1$	0	1
0	.839	.161		0	.907	.093
1	.073	.927		1	.035	.965

Italy

Actual				Predicted		
	Work in t				**Work in t**	
Work in $t-1$	0	1		Work in $t-1$	0	1
0	.932	.068		0	.966	.034
1	.064	.936		1	.025	.975

Spain

Actual				Predicted		
	Work in t				**Work in t**	
Work in $t-1$	0	1		Work in $t-1$	0	1
0	.928	.072		0	.963	.037
1	.112	.888		1	.033	.967

Table 8
Employment protection and child care

Country	Employment protection index (1)	Part-time work (2)	% Child care (<3) (3)	Child care opening hours (4)	Child benefits (% GDP) (5)
France	1.4	31.0	39	10	2.8
Italy	1.5	17.4	7	8	.9
Spain	1.4	17.2	5	6	.4

Note: The Employment Protection Index is derived from Blanchard and Wolfers (2000). The data on public child care and child care benefits are drawn from the OECD, Eurostat and Fondazione Innocenti Statistics (2000).

7.1. Labor market flexibility and child care services

As a proxy for labor market flexibility, consider the employment protection index. The employment protection index ranks countries on the basis of employment protection legislation (EPL)—on the basis of regulations governing individual dismissals and hiring of workers (e.g., severance pay and advance notice). Theoretical models indicate that employment should be more stable when EPL is stricter. Given a constant cyclical wage pattern, higher firing costs stabilize employment in downturns but also deter employers from hiring in upturns. Since stricter EPL leads to less turnover, and an overall lower supply of jobs, it should be associated with higher job search costs and stronger state dependence. Column (1) of Table 8 ranks France, Italy and Spain in terms of EPL. Consistent with Eq. (5), Italy has the highest index score as well as the greatest degree of first-order state dependence. On the other hand, there are no differences in the index score between Spain and France.

Column (2) compares the three countries in terms of the proportion of workers employed in part-time jobs. Assuming that the underlying demand for part-time work is the same in all three countries, the higher the proportion of workers employed in part-time jobs, the greater is the supply of these jobs and the more flexible is the labor market. According to this measure, sharper differences between the countries arise. Italy and Spain are similar, but France has a much higher proportion of workers employed in part-time work. This second proxy is consistent with the cross-country pattern of adjustment costs effects.

If we likewise assume that the underlying demand for child care services is the same in each country, then the percentage of children in child care will proxy for the supply of child care services and the level of child care search costs. Column (3) reports the percentage of children less than three years of age in child care in France, Italy and Spain. Consistent with Eq. (5), Italy and Spain have much lower percentages relative to France. Column (4) compares across countries the average opening hours of child care (for children less than three years of age). France has a greater availability of child care

Table 9
Simulated effect of institutional environment on women's work and birth decisions

	Less than HS educ. % point change in rate of		More than HS educ. % point change in rate of	
	Participation	Birth	Participation	Birth
French parameters				
Italian women	17.5	.3	3.8	−.2
Spanish women	29.4	.1	21.9	2.1
Italian parameters				
French women	−17.8	−.4	−.5	−.4
Spanish women	12.0	−.3	21.7	2.3
Spanish parameters				
French women	−30.0	−.1	−24.0	−2.1
Italian women	−12.4	.2	−24.3	−1.3

Note: The numbers in parentheses are the simulated percentage point changes assuming no permanent unobserved heterogeneity or transitory serial correlation.

services on this measure as well. Column (5) compares public child benefits as a percentage of GDP. The French percentage far exceeds the percentages in Italy and Spain.[18]

Overall, the data in Table 8 indicate that in Italy and Spain, relative to France, part-time work is in scarce supply and child care services are more limited in the number of weekly hours available (see also Del Boca, 2002). This raises job search and child care costs, or participation adjustment costs, in Italy and Spain relative to France, and plausibly explains the cross-country order of state dependence effects found in estimation. Facing more substantial adjustment costs, women in Italy and Spain that decide to bear a child less often withdraw from the labor market, or more often fail to re-enter. Moreover, women who remain employed tend to have full-time work commitments, which is not compatible with having many children, so overall fertility in Italy and Spain is relatively lower than in France (see also Boeri et al., 2005).

7.2. Measuring the effect of the institutional environment

In Table 9, we report the results of a simulation exercise which quantifies the overall influence of the institutional/social policy environment on labor force participation and birth rates. In the simulation, predicted participation and birth outcomes are generated for each women in one country, using the SML estimates of the decision rules in an alternative country. The results of the counterfactual exercise are partial equilibrium only, in the sense that the background characteristics (e.g., nonlabor income and education) are assumed to remain the same after changing the institutional environment. Moreover, we hold constant the country-specific distribution of unobserved heterogeneity in the simulations, so that changes in work and fertility choices will be driven mostly by differential first-order state dependence effects.

In the top panel of Table 9, Italian and Spanish women face the decision rule parameters of the French. The results indicate that if Italian women, who have not completed secondary education, were to make work and birth decisions in the relatively more flexible French environment, they would increase their average participation rate over the sample period by 17.5 percentage points. However, their average birth rate would increase by only .3 percentage points. Among Italian women who have completed secondary education, the participation rate would increase by relatively less, 3.8 percentage points, and the birth rate would decrease by a small amount (.2 percentage points). The main payoff to a more flexible environment for Italian women would be a higher participation rate among the less educated.

If Spanish women were to face the more flexible French institutional environment, the participation rate of less educated women would increase by a very large 29.4 percentage points. More educated women would also increase their participation rate by a substantial amount (21.9 percentage points). The increase in the birth rate among less educated Spanish women would be a negligible .1 percentage points, but more educated Spanish women would increase their birth rate by 2.1 percentage points. In contrast to Italian women, both less educated and more highly educated Spanish women would increase their labor market participation rates were they to face the French institutional environment, and more educated Spanish women would increase their birth rate.

The two bottom panels of Table 7 perform the analogous experiments of having French and Spanish women face the Italian parameters, and French and Italian women face the Spanish parameters. The results are mostly symmetric. French women would decrease their participation rates in the Italian and Spanish environments, and Spanish women would benefit from the Italian environment in terms of participation outcomes. There is an especially large increase in the

---

[18] The extension of Allocation Parentale d' 'Education (APE) to births of parity 2 in 1994 is often cited as a cause of the recent growth in fertility (Laroque and Salaniè, 2005).

participation of more educated Spanish women when facing the Italian parameters. This is because the effect of secondary education on participation decisions is much stronger in Italy than in Spain.

Overall, the simulations suggest that institutions/policies that promoted more flexible labor markets and child care availability in Italy in Spain, as in France, would result in substantial convergence to the estimated French female labor force participation rate of 68%. With French parameters, Italian female participation rates rise from 53% to 63%, and Spanish female participation rates increase dramatically from 35% to 62%. Thus, our model indicates that Italy and Spain could meet EU female employment target levels, at least 60% by 2010, were French social policies to be adopted in those countries.[19]

Note that since we have limited the set of countries to only those with similar cultural characteristics (such as attitudes towards gender roles), these simulations provide a better approximation to the underlying influence of employment and social policies than if we had attempted to include a larger set of countries that differed more substantially in other unobserved dimensions. Nonetheless, the overall changes in participation and birth rates should be considered only as upper bound employment and social policy effects.

8. Conclusion

In this paper, we formulate a general dynamic utility maximization model of female labor market participation and fertility choices and estimate the approximate decision rules of the model using data from the ECHP on married women in Italy, Spain, and France. The main focus of the paper is on measuring the differential relative importance of state dependence and permanent unobserved heterogeneity in work and family preferences across countries. The estimated decision rules indicate that in each country first-order state dependence is the most important factor explaining female labor force participation rates. We examine the relationship between first-order state dependence and aggregate measures of social policies, and find that the order of state dependence effects is consistent with the cross-country ranking in the extent of labor market flexibility and the supply of child care services. This suggests important differences in institutional environments.

We also use the estimated decision rules to quantify the effects of the institutional environment by simulating counterfactual female participation and birth outcomes when women in one country face the decision rule parameters of a different country. The results of the simulation suggest that Italian and Spanish women would substantially increase their participation in the labor market were they to face the relatively more flexible French social policy environment. The convergence of Italian and Spanish participation rates to the higher estimated French rates is especially pronounced amongst less educated women. The counterfactual simulation indicates that female employment rates in Italy and Spain would reach EU target levels, at least 60% by 2010, were French social policies to be adopted in those countries.

One limitation of our study is that we were only able to provide indirect evidence on the effects of employment and social policies on female labor market participation and persistence. This is the only way to proceed when aggregate proxies for social policies do not vary sufficiently over individuals and over time. If better proxies for social policies were to become available, then direct evidence of policy effects might be obtainable by directly entering the proxies into the individual's approximate decision rules. This remains an area for future data collection and research.

An additional limitation of our study is that we estimate approximate decision rules rather than exact decision rules. The exact solution approach is computationally more intensive, but it more faithfully incorporates cross-equation and forward-looking restrictions implied by the dynamic decision problem. Thus, exact decision rules may look very different from approximate ones. The comparison of exact and approximate decision rules is another area for future research.

References

Angrist, J.D., Evans, W.N., 1998. Children and their parents' labor supply: evidence from exogenous variation in family size. American Economic Review 88 (3), 450–477.
Azmat, G., Güell, M., Manning, A., 2004. Gender gaps in unemployment rates in OECD countries. CEPR Discussion Paper No. 4307.
Blanchard, O., Wolfers, J., 2000. The role of shocks and institutions in the rise of european unemployment: The aggregate evidence. Economic Journal 110 (462), 1.
Boeri, T., Del Boca, D., Pissarides, C., 2005. Women at Work: An Economic Perspective. Oxford University Press, Oxford.
Browning, M., 1992. Children and household economic behavior. Journal of Economic Literature 30, 1434–1465.

---

[19] Approximately 70% of the increase in Italian and Spanish participation rates is due to differential state dependence effects. This was determined by additional simulations that equalize the state dependence parameters across countries.

Buchinsky, M., Gotlibovski, C., 2006. The residential location choices and labor market outcomes of immigrants: Estimating the impact of government intervention in the housing market. Unpublished manuscript.

Carrasco, R., 2001. Binary choice with binary endogenous regressors in panel data: Estimating the effect of fertility on female labor participation. Journal of Business & Economic Statistics 19 (4), 385–394.

Chamberlain, G., 1984. Panel data. In: Griliches, Z., Intrilligator, M.D. (Eds.), Handbook of Econometrics, vol. 2. Elsevier, Amsterdam (Chapter 22).

Croda, E., Kyriazidou, E., 2004. Intertemporal labor force participation of married women in Germany: A panel data analysis. Unpublished manuscript.

Del Boca, D., 2002. The effect of child care and part time on participation and fertility of Italian women. Journal of Population Economics 14.

Eckstein, Z., Wolpin, K.I., 1989. Dynamic labor force participation of married women and endogenous work experience. Review of Economic Studies 56, 375–390.

Francesconi, M., 2002. A joint dynamic model of fertility and work of married women. Journal of Labor Economics 20, 336–388.

Hausman, J.A., Abrevaya, J., Scott-Morton, F.M., 1998. Misclassification of the dependent variable in a discrete-response setting. Journal of Econometrics 87, 239–269.

Heckman, J.J., 1981. The incidental parameters problem and the problem of initial conditions in estimating a discrete time–discrete data stochastic process and some Monte Carlo evidence. In: Manski, C.F., McFadden, D. (Eds.), Structural Analysis of Discrete Data with Econometric Applications. MIT Press, Cambridge, MA.

Heckman, J.J., MaCurdy, T., 1980. A life cycle model of female labour supply. Review of Economic Studies 48, 47–74.

Heckman, J.J., Singer, B., 1984. A method for minimizing the impact of distributional assumptions in econometric models for duration data. Econometrica 2, 271–320.

Heckman, J.J., Willis, R., 1977. A beta-logistic model for the analysis of sequential labor force participation by married women. Journal of Political Economy 85, 27–58.

Hotz, V.J., Miller, R.A., 1988. An empirical analysis of life cycle fertility and female labor supply. Econometrica 56 (1), 91–118.

Hotz, V.J., Kydland, F.E., Sedlacek, G.L., 1988. Intertemporal preferences and labor supply. Econometrica 56 (2), 335–360.

Hyslop, D.R., 1999. State dependence, serial correlation and heterogeneity in intertemporal labor force participation of married women. Econometrica 67 (6), 1255–1294.

Jakubson, G., 1988. The sensitivity of labor supply parameter estimates to unobserved individual effects: Fixed- and random-effects estimates in a nonlinear model using panel data. Journal of Labor Economics 6, 302–329.

Keane, M.P., Sauer, R.M., 2005. A computationally practical simulation estimation algorithm for dynamic panel data models with unobserved endogenous state variables. Unpublished manuscript.

Keane, M.P., Sauer, R.M., 2006. Classification error in dynamic discrete choice models: Implications for female labor supply behavior. Unpublished manuscript.

Keane, M.P., Wolpin, K.I., 1997. Introduction to the JBES special issue on structural estimation in applied microeconomics. Journal of Business & Economic Statistics 15 (2), 111–114.

Keane, M.P., Wolpin, K.I., 2001. The effect of parental transfers and borrowing constraints on educational attainment. International Economic Review 42, 1051–1103.

Keane, M.P., Wolpin, K.I., 2002. Estimating welfare effects consistent with forward-looking behavior: Part II: Empirical results. Journal of Human Resources 37 (3), 600–622.

Keane, M.P., Wolpin, K.I., 2006. The role of labor and marriage markets, preference heterogeneity and the welfare system in the life cycle decisions of black, hispanic and white women. Pier Working Paper no. 06-004.

Laroque, G., Salaniè, B., 2005. Does fertility respond to financial incentives? Unpublished manuscript.

Magnac, T., 2000. Subsidized training and youth unemployment: Distinguishing unobserved heterogeneity from state dependence in labour market histories. Economic Journal 110, 805–857.

Moffitt, R., 1984. Profiles of fertility, labour supply and wages of married women: A complete life cycle model. Review of Economic Studies 51 (2), 263–278.

Mroz, T.A., 1987. The sensitivity of an empirical model of married women's hours of work to economic and statistical assumptions. Econometrica 55 (4), 765–799.

Nakamura, A., Nakamura, M., 1985. Dynamic models of the labor force behavior of married women which can be estimated using limited amounts of information. Journal of Econometrics 27, 273–298.

Nicoletti C., Peracchi, F., 2004. Survey response and survey characteristics: Micro-level evidence from the European Community Household Panel. CHILD Working Paper.

Rosenzweig, M.R., Schultz, T.P., 1985. The demand for and supply of births: Fertility and its Life cycle consequences. American Economic Review 75 (5), 992–1015.

Rosenzweig, M.R., Wolpin, K.I., 1980. Life-cycle labor supply and fertility: Casual inferences from household models. Journal of Political Economy 88 (2), 328–348.

Weeks, M., Orme, C., 2004. The statistical relationship between bivariate and multinomial choice models. Working Paper No. 9912, University of Cambridge.