Assessing the Impact of the Maternity Capital Policy in Russia Using a Dynamic Stochastic Model of Fertility and Employment

Fabián Slonimczyk* Anna Yurko*

October 3, 2012

Abstract

With declining population and fertility rates below replacement levels, Russia is currently facing a demographic crisis. Starting in 2007, the federal government has pursued an ambitious pro-natalist policy. Women who give birth to at least two children are entitled to "maternity capital" assistance. In this paper we estimate a structural dynamic programming model of fertility and labor force participation in order to evaluate the effectiveness of the policy.

JEL classification: J13, C61

Key words: fertility; female labor supply; structural estimation; Russia

*International College of Economics and Finance. Higher School of Economics, Moscow, Russia; `fslonimczyk@hse.ru` and `ayurko@hse.ru`.

1 Introduction

For several decades now, economists have theorized fertility decisions as a special case of consumers' utility maximization problem.[1] Children produce certain satisfactions and have a net cost, and couples have to decide on the optimal number of children. A more recent development involves the recourse by a number of governments to the use of direct financial incentives in an attempt to revert declining fertility rates. While the details are different in each case, Australia, France, Germany, Canada (the province of Quebec), and Spain have all offered "baby bonuses" to couples.

Russia is among the countries with very low fertility rates: its total fertility rate (TFR) over the period 2001–2005 was only 1.3.[2] In order to encourage women to have more children, the State Duma (Russian Parliament) passed a law in December of 2006 establishing new measures of government support for families with children, commonly known as the maternity capital (MC) program. According to the law, starting in January 2007 women that give birth to or adopt a second or consecutive child are entitled to special financial assistance. The program is scheduled to expire by the end of 2016.[3]

MC assistance comes in the form of a certificate that entitles its holder to receive funds in the amount of approximately $\$11,000$ at any time after the child reaches the age of three.[4] The money can be used for a limited number of purposes. Specifically, parents can receive these funds if they intend to spend them on: 1) acquiring housing, 2) paying for children's education, or 3) investing in the mother's retirement fund. Women can apply for MC funds only once in their lifetimes.

Through the end of 2011, the Russian government has issued over three million MC certificates.[5] At the approximate value of $\$11,000$ per certificate, total liabilities due to the MC program are growing at a rate above \$7 billion per annum, or 2.4% of total federal government expenditures in 2011. In comparison, the fraction of the federal budget dedicated to education was 4.85%. Fortunately for public finances, parents are in no rush to claim and spend the money: out of the issued certificates only 26% have been claimed so far, most of them (98.1%) used on acquiring and improving

[1] See Becker (1960) for an early formulation. Hotz et al. (1997) and Arroyo and Zhang (1997) review the literature.

[2] The TFR is defined as the total number of children born to the average woman over her lifetime. It is computed as the sum of the current age-specific fertility rates. Population size is steady when the TFR is around 2.1.

[3] Currently, there is discussion over whether to extend the program until the end of 2025.

[4] The amount in Russian rubles is revised annually to adjust for inflation.

[5] Source: Ministry of Healthcare and Social Development, Russia, `http://www.minzdravsoc.ru/health/child/154`.

housing conditions.

How effective is this policy in increasing fertility? In 2006, Gary Becker wrote in his blog on the expected effect of the proposed MC policy: "I would guess that Russian fertility would increase by about 10–20 percent from current levels, or from the present total fertility rate of 1.28 to perhaps as high as 1.55." As of 2009, Russia's TFR was 1.54. It seems that Becker's prediction has been correct and the policy results in more births.

Predictably, the government attributes the higher birth rates to its policies, specifically to the MC program. Russian demographers are more skeptical, however, noting that the TFR has been increasing since 2000 at approximately constant rates and that TFR and other aggregate measures of fertility are very unreliable indicators of actual fertility behavior (Zakharov, 2012).

There are some previous studies that investigate the effect of financial incentives on fertility. For example, Dickert-Conlin and Chandra (1999) estimate that increasing the tax benefit of having a child by \$500 raises the probability of having the child in the last week of December by 26.9 percent. Similarly, using three substantial changes in tax policy in France, Chen (2011) finds mixed evidence that fertility responds to positive and negative changes in tax incentives. Gans and Leigh (2009) find that in Australia over 1000 births were "moved" so as to ensure that their parents were eligible for the Baby Bonus, with about one quarter being moved by more than one week. Finally, Milligan (2005) finds that the introduction of a pronatalist transfer policy in the Canadian province of Quebec had a strong effect on fertility.

In order to investigate whether the MC program has been successful in increasing fertility rates, in this paper we estimate a dynamic stochastic discrete choice model of fertility and employment. We then use the estimates of the structural parameters to predict the effect of the policy. The model we estimate builds on previous dynamic fertility models such as Wolpin (1984), Francesconi (2002) and Todd and Wolpin (2006). The decision horizon for each woman begins at age 22, after schooling is completed, and ends at the retirement age of 55.[6] At each age, a woman decides whether to work or not and whether to have a child or not, so as to maximize the expected discounted present value of remaining lifetime utility. The birth decision can only be made during the fertile period, which is assumed to end at age 40.

The woman's utility at every age depends on her current period's decisions, the number of children she already has, her consumption, work

[6] The purpose of the MC policy is to encourage women to have more than one child. While many women start having children before the age of 22, the majority does not have more than one child before that age. In fact, in our sample we do not observe any women younger than 22 with more than one child.

experience, and schooling. Her consumption is the difference between her income, which consists of her wages if she chooses to work and other income of her household (including, possibly, a partner's income), and the expenses of raising children and working outside the home if she works. The woman's earnings are endogenous and stochastic, and depend on her work experience and schooling. The utility function is specified so as to allow for both psychic costs and benefits of working and having children.

Current decisions affect the future: the decision to work increases her work experience and the decision to have a child increases the future number of children she needs to raise. The model is solved by backwards induction for each element of the state space at every age. The structural parameters of the model are estimated using individual level data on choices and earnings via the simulated maximum likelihood method.

We find that the MC policy has had almost no effect in increasing the number of births. Our results indicate that women in Russia are sensitive to economic incentives, so a well-designed pro-natalist policy should be effective. Our preliminary conclusion is that the design of the MC policy, in particular the fact that it can only be used for specific purposes, deems it ineffective.

The paper is structured as follows. Section 2 presents the model and the estimation method. Section 3 describes the data and section 4 presents estimation results. Finally, section 5 presents simulation results and concludes.

2 The Model

This section presents a dynamic stochastic model of fertility and labor force participation. We consider a woman making decisions among discrete alternatives at each point in time so as to maximize the present value of expected lifetime utility. The model focuses on two decisions. First, at each age t the woman decides whether to participate in the job market ($l_t = 1$).[7] Second, women in fertile age can choose to give birth ($n_t = 1$). To simplify matters, we assume fertility is a deterministic process over which women have full control.[8] We index the four mutually exclusive alternatives facing women by j:

[7] Part-time work is relatively rare in Russia. Only 3.3% of employed women in our sample work 20 or less hours per week. For this reason, we do not allow for separate full- and part-time participation decisions.

[8] Other studies, such as Hotz and Miller (1988), specify complex stochastic functions that make the probability of a birth depend on, among other factors, birth control intensity and the age of the mother.

$$
j = \begin{cases} 1 & \text{if no birth and no work} \\ 2 & \text{if no birth and work} \\ 3 & \text{if birth and no work} \\ 4 & \text{if birth and work} \end{cases}
$$

We let the decision process start at age 22, set the end of the fertile period at age 40, and end the program at the official retirement age of 55.[9] The starting point is convenient since a vast majority of Russian women finish their education by age 22.[10] Moreover, while some women become mothers at a younger age, second births —the focus of the MC policy— occur after our starting age in over 99% of cases. Formally, the woman's objective function can be written

$$
\mathrm{E}\left[\sum_{t=22}^{54} \rho^{t-22} U_t(c_t, l_t, n_t, X_{t-1}, N_t, B_t, S, m_t)\right]
$$

where ρ is the subjective rate of discount and the expectation is taken over the stochastic components of utility and earnings.[11] Women derive utility from consumption of a composite good (c_t), giving birth, and (disutility) from working. Utility is not intertemporally separable since labor market experience (X_{t-1}), the total number of children (N_t), and the presence of (one or more) children less than 3 years old (B_t) —all of them results from past decisions— are assumed to affect current tastes. Finally, utility is affected by the woman's education (S) and marital status (m_t). While education remains constant over time, marital status is assumed to evolve following a first-order markovian process whose underlying parameters are allowed to change as the woman ages.[12] The specific functional form for the utility function is

$$
\begin{aligned}
U_t = & c_t + \alpha_1 l_t + (\alpha_2 + \epsilon_t^n)\, n_t + \alpha_3 I_{N_t=1} + \alpha_4 I_{N_t=2} + \alpha_5 I_{N_t>2} \\
& + \beta_1 c_t l_t + \beta_2 c_t n_t + \beta_3 l_t n_t \\
& + (\delta_1 n_t + \delta_2 l_t + \delta_3 I_{N_t=1} + \delta_4 I_{N_t=2} + \delta_5 I_{N_t>2} + \delta_6 l_t n_t)\, m_t
\end{aligned}
$$

[9] Only in a handful of cases do women in our data set give birth after age 40. The estimation process ignores any fertility decision after the cutoff age. The last decision period is $t = 54$.

[10] According to the RLMS, only 0.5% of women 22 and older are students.

[11] Technically, the expectations operator should be time subscripted because the starting marital status affects future outcomes.

[12] Specifically, we allow transition probabilities to differ between women in different age intervals. The transition matrices are estimated outside the model (see table A.1 in the appendix for the estimated transition probabilities).

$$+ \big(\gamma_1 X_{t-1} + \gamma_2 S_1 + \gamma_3 S_2 + \gamma_4 S_3 + \gamma_5 S_4 + \gamma_6 I_{N_t=1} + \gamma_7 I_{N_t=2} + \gamma_8 I_{N_t>2} + \gamma_9 B_t\big) l_t \tag{1}$$

Instantaneous utility is linear and additive in consumption. Giving birth has both a deterministic (α_2) and a stochastic (ϵ_t^n) effect on utility. Note that I_x is an indicator function equal to 1 if statement x is true and zero otherwise. Work and births affect the marginal utility of consumption and births affect the marginal disutility of work. Marital status does not enter utility directly but modifies the effect of births, employment, and children. Finally, the disutility of work depends on previous work experience (habit formation), highest education completed[13], the number of children, and the presence of a small child.

The model does not permit either savings or borrowing. Consumption each period must equal total income minus the costs associated with work, giving birth, and rearing children. Formally, the budget constraint is written:

$$\begin{aligned} c_t =& y_t^f l_t + y_t^o + (\phi_1 + \phi_2 H) MC n_t K \\ & - b_1 l_t - b_2 n_t - b_3 I_{N_t=1} - b_4 I_{N_t=2} - b_5 I_{N_t>2} \end{aligned} \tag{2}$$

The linearity in consumption of the utility function means that the parameters corresponding to these monetary costs (b_s) cannot be separately identified from the "psychic" benefits. Therefore, we set the former parameters to zero and interpret the latter as benefits net of cost.

Women receive labor income y_t^f when employed and income from other household members y_t^o, including the spouse's income when married. In addition, eligible women ($MC = 1$) receive maternity capital assistance in the amount K if they give birth.[14] Because assistance can only be obtained three years after the birth and must be used for specific purposes, we estimate two parameters (ϕ_s) that convert assistance dollars into a monetary equivalent consumption value.[15] The state variable H is an indicator of whether a household member owns the residence. We expect ϕ_2 to be positive, reflecting the fact that home owners are in a better position to "cash in" the MC assistance.[16]

The woman's non-labor income depends on her characteristics and marital status. Women are assumed to form expectations according to

[13] S_1 through S_4 correspond to secondary school, vocational school, technical school, and university respectively.

[14] We set $K = 108,557$, the average real value (in rubles of year 2000) of MC assistance over the period 2007–2010.

[15] Keane and Wolpin (2010) use the same procedure when evaluating welfare participation in the U.S.

[16] We assume that home ownership is a state variable that remains constant over time.

$$\overline{\log y^o_t} = c_0 + c_1 m_t + c_2 t + c_3 t^2 + c_4 S_1 + c_5 S_2 + c_6 S_3 + c_7 S_4 \quad (3)$$

Equation 3 does not depend on current or future decisions so it is estimated outside the model.[17] Note that non-labor income depends on the random state m_t, so women use the transition probabilities in table A.1 to estimate the expected value.

The earnings offer function depends on the woman's accumulated human capital as follows:

$$\log y^f_t = a_0 + a_1 X_{t-1} + a_2 X^2_{t-1} + a_3 S_1 + a_4 S_2 + a_5 S_3 + a_6 S_4 + \epsilon^y_t \quad (4)$$

The shock ϵ^y_t captures variation in earnings that is independent of the decision process. The two shocks $(\epsilon^n_t, \epsilon^y_t)$ are jointly normally distributed with zero mean, finite variance, and non-zero contemporaneous covariance. The shocks are assumed to be serially independent, so past realizations do not provide information on future shocks.

The model allows for unobserved individual heterogeneity in the following parameters: utility of giving birth (α_2,δ_1), utility associated with having children (α_3–α_5,δ_3–δ_5), and baseline earnings (a_0). Heterogeneity is introduced as a set of unobservable types, with each type having its own associated set of parameters. The proportion of women corresponding to each type is estimated jointly with the model parameters as explained below.

In addition to the shocks and the realization of the marital status process, the state variables informing employment and fertility decisions include the history of choices up to age t. Let the state space be denoted by $\Omega_t = (N_{t-1}, X_{t-1}, n_{t-1}, n_{t-2}, n_{t-3}, S_1, \ldots, S_4, H, \overline{y^o_t}, m_t, \epsilon^n_t, \epsilon^y_t)$.[18] The value function $V(\Omega_t)$ is the maximal expected present value of the remaining lifetime utility given the state at age t.[19] Because the alternatives facing the woman are discrete, the value function can be written as the maximum over alternative-specific value functions:

$$V(\Omega_t) = \max_{j \in J_t} [V_j(\Omega_t)]$$

where $J_t = \{1, \ldots, 4\}$ for $t = 22, \ldots, 40$ and $J_t = \{1, 2\}$ for $t = 41, \ldots, 54$. The alternative-specific value functions obey the Bellman equation:

[17] See table A.2 in the appendix for the estimated coefficients.

[18] The presence of a child under 3 years old is defined as $B_t = I(n_{t-1} n_{t-2} n_{t-3} = 1)$.

[19] Technically, because this is a finite horizon problem, the value function should be time subscripted. We omit it to simplify notation (the time subscript would always be the same as that of the state space).

$$V_j(\Omega_t) = U_{j,t} + \rho \mathrm{E}_t \left[V(\Omega_{t+1}) \mid \Omega_t, j \in J_t \right] \qquad \text{for } t < 54$$
$$= U_{j,54} \qquad \text{for } t = 54$$

Finally, the pre-determined state variables evolve according to

$$N_t = N_{t-1} + n_t$$
$$X_{t-1} = X_{t-2} + l_{t-1}$$

2.1 Model Solution and Estimation

The solution to the finite-horizon dynamic programming problem can be found using backward recursion, which in turn enters into the estimation of the structural parameters.

A woman in her last period only needs to evaluate two alternatives. The alternative utility levels depend on the pre-determined part of the state space (Ω_t^d) and the particular realization of the random components.[20] Therefore, the last period's decision can be seen as a static random utility model. Given data on actual decisions of 54 year old women, their earnings, and the observable components of the state space, it would be straightforward to obtain parameter estimates using maximum likelihood methods.

The extension to a dynamic setting is better understood by first considering 53 year old women. While still facing two alternatives, women of this age need to consider the effect of their choices on the next period's prospects. For example, evaluating the alternative "work" involves the following steps: 1) compute the flow utility corresponding to the alternative "work" at age 53; 2) Update the state space for age 54 (e.g. add one year of experience); 3) Given the new state, the fact that she will act optimally at age 54 allows the use of the value functions for age 54 (this the recursive step); 4) With these inputs it is possible to calculate the age 53 value of working.

These steps need to be repeated for the alternative "not work". At this point, the decision at age 53 only depends on the (unobservable to the researcher) shock ϵ_t^y.[21] Given data for 53 year old women, the solution to the dynamic program makes it possible to estimate the parameter values that maximize the likelihood of observed behavior. The same logic applies to younger women.[22]

[20] Marital status is included in Ω_t^d.

[21] Only women in fertile age are affected by ϵ_t^n.

[22] The solution for women 40 years old and younger is more computationally demanding since it involves the doubling of the decision tree that must be considered.

Letting $d_{i,t}$ denote the combination of the choice and earnings (i.e. $d_{i,t} = j$ for $j = 1, 3$ and $d_{i,t} = (j, y_t^f)$ for $j = 2, 4$) for woman i at age t, we have

$$
\begin{aligned}
\Pr(d_{i,t} \mid \Omega_t^d) =& \Pr\left(j = \arg\max_k V_k(\Omega_t)\right) && \text{for } j = 1, 3 \\
\Pr(d_{i,t} \mid \Omega_t^d) =& \Pr\left(j = \arg\max_k V_k(\Omega_t)\right) \\
& \times \Pr\left(y_t^f \mid j = \arg\max_k V_k(\Omega_t)\right) && \text{for } j = 2, 4
\end{aligned}
$$

Given the serial independence of the shocks, the joint probability of a sequence of choices is

$$\Pr(d_{i,22}, \ldots, d_{i,54} \mid \Omega_{22}^d) = \prod_{t=22}^{54} \Pr(d_{i,t} \mid \Omega_t^d) \tag{5}$$

In turn, the likelihood for a sample of women is simply the product of (5) over the N women in the sample. In order to generate the probabilities in the right hand side of (5), we solve the dynamic program for 20 simulations of the random shocks. Thus, the estimation program involves two loops: the first loop iterates over parameter values, while the second loop —for given parameter values— solves the model using backward recursion and obtains via simulation the likelihood of observing the actual choices in the data. The procedure stops when the likelihood of the sample data is maximized.

The introduction of unobservable types into the model modifies the objective likelihood function as follows

$$\mathrm{L}_i(\boldsymbol{\theta}) = \sum_{h=1}^{H} \mu_h \prod_{t=22}^{54} \Pr(d_{i,t} \mid \Omega_t^d, \text{type} = h)$$

where $\boldsymbol{\theta}$ is the vector of parameters, including the errors variance-covariance matrix and the type proportions (μ_h).[23]

It is standard in this setting to assume earnings are measured with error. Let observed earnings, $\widetilde{y}_t^f$, be given by

$$
\begin{aligned}
\log \widetilde{y}_t^f &= \log y_t^f + u_t^f \\
u_t^f &\sim \mathrm{N}(0, \sigma_u^2)
\end{aligned}
$$

[23]The only non-estimated parameter is the time discounting rate, ρ, which we set to 0.95.

where u_t^f is measurement error, which is assumed to be uncorrelated with other shocks and also over time. The rationale for including measurement error in the estimation step is twofold. First, it is reasonable to believe that earnings are not reported accurately. Second, an extra error component is necessary to prevent a degenerate likelihood due to outliers. Technically, this could happen in situations when the woman is observed working but her earnings are too low to justify her decision given the parameter values and the realized earnings shock.[24]

3 Data Description

Our data comes from the Russian Longitudinal Monitoring Survey (RLMS), a household panel survey based on the first national probability sample drawn in the Russian Federation.[25] We use data from rounds XIII–XIX covering the period 2004–2010. In a typical round, 10,000 individuals in 4,000 households are interviewed. These individuals reside in 32 oblasts (regions) and 7 federal districts of the Russian Federation. A series of questions about the household (referred to as the "family questionnaire") are answered by one household member selected as the reference person. We use the family roster to create a fertility history for each woman in the panel. In turn, each adult in the household is interviewed individually (the "adult questionnaire"), providing information on labor market participation, experience, schooling and earnings.

3.1 Variable Definitions

Employment The RLMS contains information on a main job and a secondary job.[26] A woman is considered employed if she usually works 10 or more hours per week at all jobs.

Experience The adult questionnaire includes an item regarding past labor market experience. We construct our experience variable as follows. First, we use the RLMS data to determine previous experience in the first round the individual is interviewed.[27] In subsequent rounds we let expe-

[24]Alternatively, one could include a random disturbance to the disutility of work. However, it is harder to justify the assumption of zero correlation, both with other shocks and serially.

[25]The RLMS is conducted by the Higher School of Economics and the "Demoscope" team in Russia, together with Carolina Population Center, University of North Carolina at Chapel Hill.

[26]In addition, there are a series of items regarding irregular informal activities. We do not consider irregular activities in determining employment status.

[27]In cases when the response is missing, we use data from other rounds to impute a value.

rience evolve in a way that is consistent with the observed employment history.

Births Whether a woman has given birth during the year preceding interview is determined on the basis of the household roster. This task would be straightforward if the exact date of birth for each household member were known. Unfortunately, for privacy reasons the RLMS censors the day and month of birth. Because some households are interviewed as early in the year as October, this creates a problem for determining to which period children born in the year preceding the interview should be assigned. We proceed as follows. If these children do not appear in the previous round's roster, then we assume they were born in the period between interviews. We are thus unable to determine births for women not observed in consecutive years.

Because birth are quite unfrequent events, it could be the case that even a large representative sample like the RLMS might produce biased estimates of fertility rates. The Russian statistical agency —Rosstat— selectively publishes yearly data on the number of births based on State registries, as well as official estimates of the female population based on the census. In figure 1 we compare birth rates obtained from the RLMS and official data. While the sample data are noisier, they seem to offer a reasonable approximation to the birth rate based on official statistics.

Number of children The procedure to create our number of children variable is analogous to the one applied for labor market experience. First, we use an item from the adult questionnaire to determine the number of children in the first round the woman is observed. We then let the variable evolve in a manner consistent with her birth history.

Marital Status We consider a woman as married when there is a cohabiting spouse in the household roster. While information on marital status is also available from the individual questionnaire, the emphasis on cohabitation better represents the opportunity set confronting the woman.

Labor and Other Income The RLMS contains information on the previous month's after-tax earnings for each job, as well as an item on overall after-tax income. Our labor income variable adds earnings from the main and the second job. Individuals who work less than 10 hours per week are imputed zero labor income. Women receive other income from three sources: a) income in excess of labor income, b) income from the spouse, and c) some fraction of income from other household members. The first source is calculated as the difference between total after-tax income and our labor

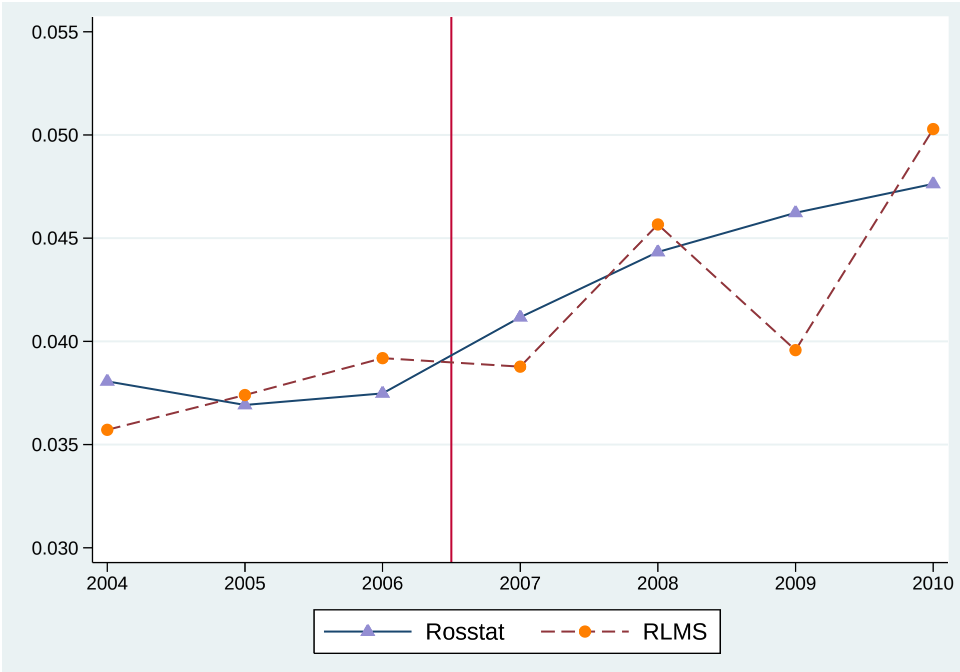

Figure 1 – Birth Rates for Women Ages 15-49

income variable. The second is obtained from the spouse's answer to the RLMS individual questionnaire. In order to estimate the third component, we proceed as follows. From the household interview, we obtain total after-tax family income. From this amount we subtract the woman's income and (if present) the spouse's income. Finally, we assume that the woman receives a fraction of this income that is proportional to the size of her nuclear family (herself, her spouse, and children living in the household) relative to overall household size. All nominal amounts are converted to rubles from year 2000 using the Russian CPI.

3.2 Sample Selection and Descriptive Statistics

Our sample is composed of women between 22 and 54 years of age. We only retain women who are observed at least 3 times during the period of analysis. After deleting observations with missing values in the relevant variables, our unbalanced panel comprises 2,031 individuals and a total of 12,117 person-year observations. Table 1 contains descriptive statistics.

In our model, women's fertile period ends exogenously at age 40. Over 60% of individuals enter our sample before crossing this threshold. Women in the sample exhibit wide variation in initial labor market experience and education attainment. The residence ownership indicator is based on an item

Table 1 – Descriptive Statistics

	Mean	Std Dev
Individuals (2031 observations)		
Years in sample	6	1.2
Age in 1st period	36	9.2
Experience in 1st period	13	10.0
Residence Owner	0.75	
Less than Secondary Educ	0.05	
Secondary Educ Complete	0.19	
Vocational School Complete	0.23	
Technical School Complete	0.31	
University Degree or above	0.22	
Person-year (12,117 observations)		
Age	38.7	9.1
Number of Children	1.4	0.9
Experience	15.2	10.1
Labor Income	2,446	2,846
Other Income	5,909	11,857
Married	0.69	
Birth	0.02	
Employed	0.72	
MC Eligible (2007–2010)	0.81	

in the family questionnaire asking whether any of the household members owns the residence. The high relative prevalence of owners vis a vis tenants reflects the successful privatization policy of residencies during the early transition period.

Women in our sample have completed fertility rates significantly below the replacement rate. For example women over 40 have on average 1.8 children. Low fertility rates occur despite the fact that Russia has one of the highest marriage rates in the world. Table 2 focuses on the relationship between cohabitation, family size, and employment. Russian women have a long history of strong attachment to the labor market. High participation rates prevail both among married and unmarried women. Moreover, mothers of one or two children are more likely to be employed than women without children. Only after the third birth does participation seem to decline somewhat.

Finally, table 3 shows the evolution of women's choices over their lifecycle. Unsurprisingly, births are concentrated in the 20s and become less and less frequent after age 30. Employment rates follow a pattern that contrasts and complements the fertility cycle. Participation in the labor market starts at less than 60% and increases slowly during the 20s. The pick employment rate is reached only in the mid-30s and remains high until the late 40s. While our model restricts the planning horizon to the official retirement age at 55, a very significant fraction of Russian women work until

Table 2 – Employment by Marital Status and Number of Children

Number of	Unmarried		Married		All	
Children	Obs.	% Employed	Obs.	% Employed	Obs.	% Employed
0	1,108	66.0	649	64.4	1,757	65.4
1	1,640	78.2	3,281	76.9	4,921	77.3
2	856	80.7	3,362	74.5	4,218	75.8
3	128	62.5	803	53.1	931	54.4
4+	25	48.0	265	31.3	290	32.8
Total	3,757	74.4	8,360	71.2	12,117	72.2

much later in life.

4 Estimation Results

In this section we describe our parameter estimates and evaluate how well the model's predictions fit the sample data. At this stage we consider these findings preliminary.

Parameter estimates are presented in table 4.

- α_1, the disutility of work, is negative as expected. In addition, working implies giving up around 1% of consumption (this suggests consumption and leisure are complements). Note that working married women do not experience significantly lower utility (δ_2 is small)
- The disutility of giving birth is large in absolute value, while having children results in positive net benefits realized over the remaining lifetime. In other words, having children involves large short-term losses that have to be balanced with long term gains. For married women, the costs of giving birth are lower while the gains from having children are higher.
- Labor market experience, births, and children all increase the disutility of work. Relative to secondary school dropouts, women with a degree suffer from disutility levels that increase with education attainment. One possible explanation is that the value of leisure time is higher for highly educated women who tend to work more than others on average.
- We estimate a very low return to on the job experience (half a percent, compared to 1% in an OLS regression)
- The multipliers associated with MC policy are essentially zero

Table 3 – Choice Distribution

Age Group	Non-employed No Birth	Non-employed Birth	Employed No Birth	Employed Birth	Total
22–24	263	27	385	21	696
	37.8	3.9	55.3	3.0	100
25–27	322	22	641	21	1,006
	32.0	2.2	63.7	2.1	100
28–30	281	26	703	34	1,044
	26.9	2.5	67.3	3.3	100
31–33	311	19	842	28	1,200
	25.9	1.6	70.2	2.3	100
34–36	282	9	930	13	1,234
	22.9	0.7	75.4	1.1	100
37–39	273	6	888	10	1,177
	23.2	0.5	75.5	0.9	100
40–44	426	2	1,363	3	1,794
	23.8	0.1	75.9	0.2	100
45–49	507	0	1,605	0	2,112
	24.0	0	76.0	0	100
50–54	591	0	1,263	0	1,854
	31.9	0	68.1	0	100
Total	3,256	111	8,620	130	12,117
	26.87	0.92	71.14	1.07	100

Note: Number of observations and percentages.

- Obtaining standard errors is computationally time consuming. They will be provided at a later stage.

Figures 2 and 3 show the ability of the model to reproduce key aspects of the data.

Table 5 shows transition probabilities among the mutually exclusive choices for women ages 22–40 and 41–54 and compares them to model predictions obtained from 200 simulations. The overall fit seems reasonable, although without standard errors it is not possible to determine whether some discrepancies are statistically significant.

Table 6 shows that ability of the model to fit behavior by unobservable type. We use the likelihood function to assign a type to each woman in the data. The fit for the work decision is remarkably accurate, whereas the model performs relatively less well for births. Clearly, type one women specialize in work, type 2 specialize in child production, while the third type balances both activities. We estimate that type 1 is the most prevalent in Russia (43%) and type 2 the least prevalent (21%). Types' behavior can be rationalized by looking at the parameter estimates in table 4. Type 1 women receive the highest wage offers (a_0), seconded by type 3. Note that births are about equally costly to all types (α_2 and δ_1) but the gains associated

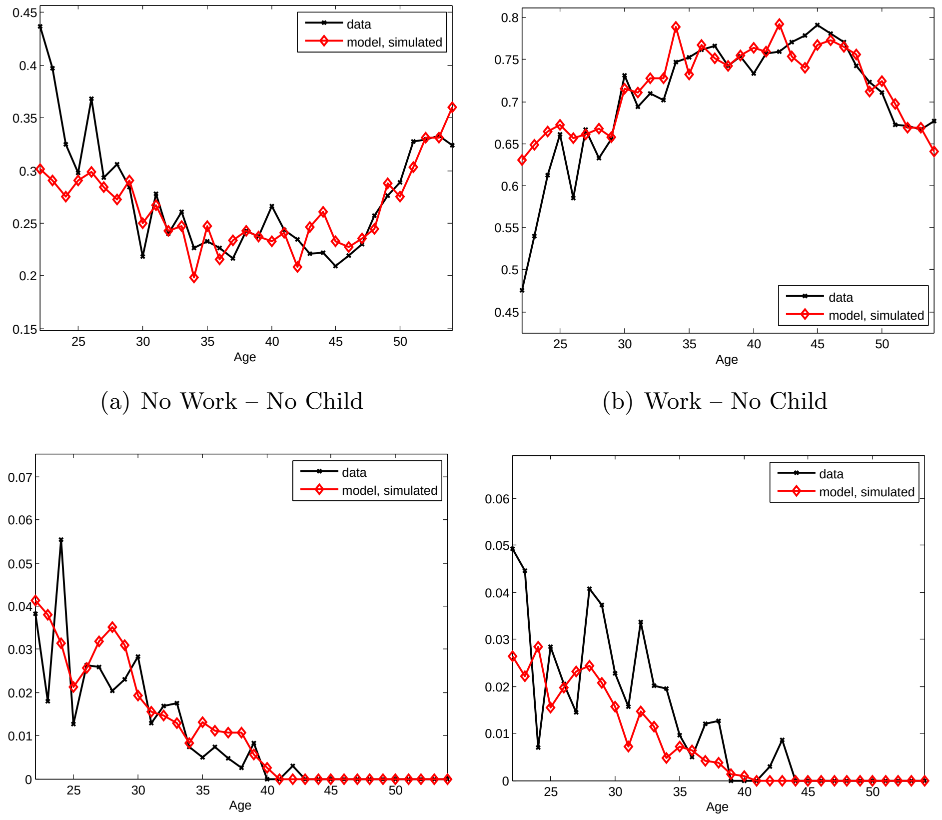

(a) No Work – No Child

(b) Work – No Child

(c) No Work – Child

(d) Work – Child

Figure 2 – Model Fit for Mutually Exclusive Choices

with children vary substantially. Indeed, it is necessary to compensate type ones and threes if they are to have children.

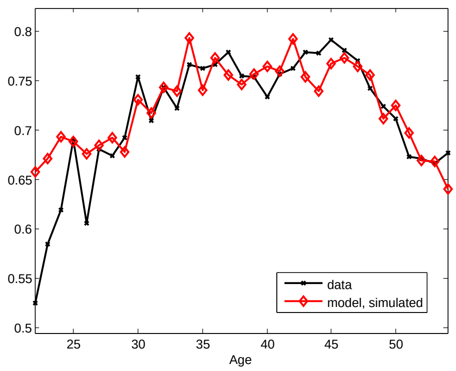

(a) LF Participation

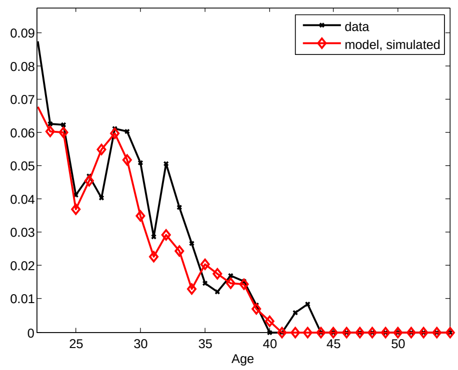

(b) Total Births

Figure 3 – Model Fit for LF Participation and Total Births

Table 4 – Maximum Likelihood Estimates

Parameter	Estimate	Parameter	Estimate
Utility function		*Earnings function*	
α_1	-3336.9907	a_0 $(type = 1)$	9.8628
α_2 $(type = 1)$	-33707.2965	a_0 $(type = 2)$	7.1211
α_2 $(type = 2)$	-33658.7978	a_0 $(type = 3)$	8.6398
α_2 $(type = 3)$	-33676.2642	a_1	0.0051
α_3 $(type = 1)$	2053.7001	a_2	$-7.3652e-05$
α_3 $(type = 2)$	1108.4243	a_3	0.2083
α_3 $(type = 3)$	1597.0186	a_4	0.5536
α_4 $(type = 1)$	2285.9512	a_5	0.7179
α_4 $(type = 2)$	1647.1113	a_6	1.1641
α_4 $(type = 3)$	2083.0012		
α_5 $(type = 1)$	3004.0086	*Error structure*	
α_5 $(type = 2)$	2545.9300	σ_n	7363.8444
α_5 $(type = 3)$	2723.5810	σ_y	0.7403
β_1	-0.0106	σ_u	0.3261
β_2	-0.0123	$\rho_{n,y}$	-0.2537
β_3	-315.4730		
δ_1 $(type = 1)$	5897.9532	*MC policy*	
δ_1 $(type = 2)$	5891.1060	ϕ_1 $(type = 1)$	$4.0782e-06$
δ_1 $(type = 3)$	6269.1885	ϕ_1 $(type = 2)$	$6.8281e-06$
δ_2	-52.8118	ϕ_1 $(type = 3)$	$4.1889e-06$
δ_3 $(type = 1)$	1167.5531	ϕ_2 $(type = 1)$	$3.4410e-06$
δ_3 $(type = 2)$	1004.3368	ϕ_2 $(type = 2)$	$6.2555e-06$
δ_3 $(type = 3)$	1098.9226	ϕ_2 $(type = 3)$	$4.8157e-06$
δ_4 $(type = 1)$	1879.4659		
δ_4 $(type = 2)$	1742.1353	*Type proportions*	
δ_4 $(type = 3)$	1844.7761	μ_1	0.4259
δ_5 $(type = 1)$	2384.4997	μ_2	0.2078
δ_5 $(type = 2)$	2392.6942	μ_3	0.3663
δ_5 $(type = 3)$	2340.1999		
δ_6	223.6071	$\log \mathrm{L}$	-15151.3962
γ_1	-66.0565		
γ_2	-64.5793		
γ_3	-9.5467		
γ_4	-167.6547		
γ_5	-1463.8074		
γ_6	-1174.3124		
γ_7	-851.3766		
γ_8	-2615.3805		
γ_9	-1337.2530		

Table 5 – Transition Probabilities: data vs. model

	Ages 22–40				Ages 41–54	
	no birth no work	no birth work	birth no work	birth work	no birth no work	no birth work
no birth no work	0.7082	0.2546	0.0325	0.0046	0.8044	0.1956
	0.6197	0.3316	0.0426	0.0061	0.6991	0.3009
no birth work	0.0754	0.8919	0.0086	0.0241	0.0747	0.9253
	0.1013	0.8767	0.0077	0.0144	0.1105	0.8895
birth no work	0.7273	0.2222	0.0303	0.0202		
	0.6580	0.2975	0.0371	0.0074		
birth work	0.1698	0.8113	0	0.0189		
	0.1992	0.7800	0.0080	0.0128		

Note: White cells contain actual transition probabilities. Gray cells contain model predictions based on 200 simulations.

Table 6 – Data versus Model: Analysis by Type

	Births (per 1,000)	Participation Rate
Type 1	13.7597	0.9802
	9.1318	0.9892
Type 2	34.0408	0.1005
	34.5695	0.1072
Type 3	19.0583	0.7722
	16.5762	0.7759
All	19.8894	0.7224
	17.1140	0.7289

Note: Gray cells contain model predictions based on 200 simulations.

5 Simulating alternative policy scenarios. Preliminary conclusions.

Having structurally estimated the parameters of the model, it is possible to address questions many of which would be out of reach for most other methodologies. Table 7 presents results from simulations in which we alter, one at a time, some of the important parameters of the model.

Table 7 – Simulations

	Births (per 1,000)[a]	Participation Rate[a]	N avg.	X avg.
Baseline model	22.584	0.645	1.186	22.428
MC policy efficacy (ϕ_1)				
0.1	+16.367	−0.012	+0.594	−0.413
0.5	+21.055	−0.021	+1.007	−0.721
1	+15.565	−0.027	+1.025	−0.941
Net utility of birth (α_2)				
+5000	+14.434	−0.014	+0.524	−0.448
+10000	+23.836	−0.024	+0.896	−0.780
Net utility from children (α_3–α_5)				
+500 (per child)	+19.670	−0.025	+0.758	−0.833
+1000 (per child)	+28.461	−0.041	+1.193	−1.334
Mean earnings (a_0)				
+10%	−0.319	+0.000	−0.013	−0.002
+30%	−0.939	+0.008	−0.035	+0.275
Earnings, return to experience (a_1)				
+1 percentage point	−0.623	−0.014	−0.022	−0.490
+3 percentage points	−1.501	−0.009	−0.050	−0.313
Mean other income (c_0)				
+10%	−0.084	+0.000	−0.003	+0.003
+30%	−0.071	−0.002	−0.004	−0.075
Utility of working with baby (γ_9)				
+1000	+3.448	−0.003	+0.123	−0.096
+5000	+17.622	−0.011	+0.657	−0.280
College graduates				
+10%	−1.812	+0.063	−0.068	+2.054
+30%	−2.834	+0.091	−0.104	+2.994

Note: Effects of changes in parameters correspond to changes with respect to baseline.
[a] Births per 1,000 and participation rates are computed using the age distribution data from the RLMS.

- According to our estimates (table 4), the multipliers corresponding to the MC policy are essentially zero for all types and independently of whether the woman is a home owner. We simulate the model setting different values for the multipliers. The result is a monotonic increase in the average number of children each woman would bear over her lifetime. If the MC were fully effective ($\phi_1 = 1$), the overall fertility rate would be 2.22, just above the replacement rate. All the simulations also show that a more effective MC policy would reduce LF participation and hence accumulated experience over the lifetime. However, the latter effect is relatively mild.

- We experiment with reducing the disutility from giving birth. The effects are comparable to those of an effective MC policy. Similar results are obtained if we increase the utility per child.
- Increases in earnings, husband's income, and the returns to experience have small effects on fertility.
- Reducing the disutility of working with a small child induces more births.
- Increases in the fraction of college graduates raise labor force participation and accumulated experience but do not affect fertility substantially.

These simulation exercises lead us to the (preliminary) conclusion that, while the MC policy as currently applied is ineffective in increasing birth rates, the underlying rationale —that fertility behavior responds to economic incentives— is correct. What would seem to be necessary is a reformulation of the policy so that these incentives are actually perceived by economic actors. It should be kept in mind, however, that effectiveness of policy in achieving its stated aims is fundamentally different from the issue of efficiency. A reformulation of the MC policy might be effective but undesirable if it fails to raise attained levels of utility for the population. In other words, there may be other more cost-effective ways to achieve the same ends.

References

Arroyo, C. and Zhang, J. Dynamic Microeconomic Models of Fertility Choice: A survey. *Journal of Population Economics*, 10(1):23–65, 1997.

Becker, G. Grappling with Russia's Demographic Time Bomb. *The Becker-Posner Blog*, 2006. `http://www.becker-posner-blog.com`.

Becker, G. S. An Economic Analysis of Fertility. In *Demographic and Economic Change in Developed Countries*, Universities-National Bureau of Economic Research Conference Series 11, pages 209–231. NBER, Princeton, NJ, 1960.

Chen, D. Can Countries Reverse Fertility Decline? Evidence from Frances Marriage and Baby Bonuses, 1929–1981. *International Tax and Public Finance*, 18(3):253–272, 2011.

Dickert-Conlin, S. and Chandra, A. Taxes and the Timing of Births. *Journal of Political Economy*, 107(1):161–177, 1999.

Francesconi, M. A Joint Dynamic Model of Fertility and Work of Married Women. *Journal of Labor Economics*, 20(2):336–380, 2002.

Gans, J. and Leigh, A. Born on the First of July: An (un) Natural Experiment in Birth Timing. *Journal of Public Economics*, 93(1-2):246–263, 2009.

Hotz, J. V.; Klerman, J., and Willis, R. The Economics of Fertility in Developed Countries. In Rosenzweig, M. R. and Stark, O., editors, *Handbook of Population and Family Economics*, volume IA, pages 275–347. Elsevier Science, 1997.

Hotz, V. J. and Miller, R. A. An Empirical Analysis of Life Cycle Fertility and Female Labor Supply. *Econometrica*, 56(1):91–118, 1988.

Keane, M. and Wolpin, K. The Role Of Labor And Marriage Markets, Preference Heterogeneity, And The Welfare System In The Life Cycle Decisions Of Black, Hispanic, And White Women. *International Economic Review*, 51(3):851–892, 2010.

Milligan, K. Subsidizing the Stork: New evidence on Tax Incentives and Fertility. *Review of Economics and Statistics*, 87(3):539–555, 2005.

Todd, P. and Wolpin, K. 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. *The American Economic Review*, 96(5):1384–1417, 2006.

Wolpin, K. An Estimable Dynamic Stochastic Model of Fertility and Child Mortality. *The Journal of Political Economy*, 92:852–874, 1984.

Zakharov, S. Increases in Fertility in 2007-2010: Cause for Optimism? *Demoscope Weekly*, pages 495–496, Jan 23–Feb 5 2012. In Russian, `http://demoscope.ru/weekly/2012/0495/tema01.php`.

A Appendix

Table A.1 – Evolution of Marital Status

Age Group	Transition Probabilities	
	$\Pr(m_t = 1 \mid m_{t-1} = 0)$	$\Pr(m_t = 0 \mid m_{t-1} = 1)$
22–25	9.36	8.25
26–30	16.36	4.78
31–35	12.31	4.05
36–40	5.19	3.6
41–45	4.52	2.38
46–50	4.47	3.05
51–55	1.17	2.15

Table A.2 – Log Non-labor Income Regression

	Coefficient	Standard Error
Married	0.966	0.020
Age	-0.022	0.009
Age Squared	0.001	0.0003
Secondary School	0.169	0.042
Vocational School	0.136	0.041
Technical School	0.144	0.040
University	0.452	0.041
Constant	10.114	0.173
Observations	11,359	
R-squared	0.187	

Note: OLS regression estimated on person-year observations with positive non-labor income.