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

Fabián Slonimczyk† Anna Yurko†

April 25, 2013

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

*We would like to thank Irina Denisova, Vladimir Gimpelson, Hartmut Lehmann, Sergey Roshin, Judith Shapiro and participants at the IZA-HSE workshop in Moscow and the Royal Economic Society conference for helpful comments. All remaining errors are our own.

†International College of Economics and Finance. NRU 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 2011, Russia's TFR was 1.58. 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.

Assessing the effect of the MC policy is challenging for two reasons. First, because in principle the policy is of universal application it is not easy to delineate reliable treatment-control distinctions. We exploit the fact that the program targets second and consecutive births to obtain a differences-in-differences estimate of the effect of the policy. However, even after controlling for a large number of observable characteristics, it is not possible to rule out the fact that the decision to give birth to a first child is fundamentally different from the decision to have two or more children. In particular, arguably the former decision is less sensitive to monetary incentives than the latter, so an improvement in economic conditions around the time the policy was introduced could lead to the false impression of a positive effect on fertility. A second challenge is that the fertility decisions of the women affected by the program will only be fully observable after they complete their fertile period. Without further assumptions it is not possible to distinguish an increase in completed fertility from a shift in the timing of births.

In order to investigate whether the MC program has been successful in increasing fertility rates while addressing these challenges, 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 analyze 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), and explicitly accounts for the differential costs and benefits of first and consecutive births. 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 45. Because women in the model are forward looking and rational, we are able to distinguish increases in long run fertility from shifts in the timing of births.

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

Our preliminary findings estimating the model show that the MC policy has had almost no effect in increasing fertility. The main driver of this result is that, if the MC policy were effective, it should lead to an increase of birth hazards at all parities. Intuitively, at the margin some women should choose to give birth to a first child because of the higher option value of a second birth later on. In contrast to this expectation, the data shows no change in the proportion of women giving birth to a first child. We cannot rule out, however, that a different specification of the model might yield predictions that are both consistent with the observed birth hazards and a more significant effect of the policy.

[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.

The paper is structured as follows. The next section presents diverse evidence on the effectiveness of the policy, including results from a 2008 poll, data on aggregate fertility rates from different sources, and results for before-after and difference-in-differences experiments. Section 3 introduces the structural model and the estimation method. Section 4 provides details on the estimating sample and section 5 presents estimation results. Finally, section 6 presents simulation results and concludes.

2 Fertility in Russia

This section presents diverse evidence on the effect of the MC program on fertility. After introducing the data sources, we look at poll evidence on the effect of the program. Next, we present time series evidence on birth rates and the total fertility rate based on official registry data. Thirdly, we use individual data from a representative sample to obtain before-after and differences-in-differences estimates. The final sub-section summarizes the findings and motivates the need for structural modeling and estimation of the fertility decision.

2.1 Data Sources

Official aggregate data on birth rates and total fertility rates is available from the Russian statistical agency's (Rosstat) website for the years 2000–2011. Information on earlier years comes from the Human Fertility Database (HFD).[7] These data are collected from official registries, so it is the best available source on fertility behavior in Russia. However, it is only available at aggregate level and does not provide any information other than the mother's age and birth order.

The main data source for this study is the Russian Longitudinal Monitoring Survey (RLMS), a household panel survey based on the first national probability sample drawn in the Russian Federation.[8] In a typical round, 10,000 individuals in 4,000 households are interviewed. These individuals reside in 32 oblast (regions) and 7 federal districts of the Russian Federation. A series of questions about the household (the "family questionnaire") are answered by one household member selected as the reference person. In turn, each adult in the household is interviewed individually (the "adult questionnaire"), providing information on labor market participation, experience, schooling and earnings.

[7]The Human Fertility Database. Max Planck Institute for Demographic Research and Vienna Institute of Demography. Available at www.humanfertility.org.

[8]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.

Table 1 – Poll Results on MC Policy

A. **Information and Beliefs about MC**	Males	Females
Knows “something” about MC	47.7% (4,823)	64.8% (6,693)
And gives correct answer to:		
... amount of MC benefit	93.9% (2,185)	94.9% (4,128)
... eligibility rule	91.5% (2,168)	93.3% (4,180)
... waiting time to receive benefit	87.8% (2,001)	90.3% (3,954)
... whether MC can be received in cash	90.1% (1,971)	91.2% (3,836)
... whether MC benefits can be split into separate uses	37.9% (1,822)	36.8% (3,504)
Has drawn up MC certificate (eligible women only)	-	50.8% (134)
Believes eligible women will be able to get the money from government	66.7% (1,914)	67.6% (3,598)
B. **Influence of MC Policy**	Males	Females
MC influenced the # of children they want	2.8% (3,821)	4.2% (3,755)
Average # of Children Desired		
... if claims not to have been influenced	1.6 (3,516)	1.6 (3,472)
... if claims to have been influenced	1.9 (102)	1.8 (147)
Influenced Decision to (eligible women)		
... give birth/get pregnant/adopt	-	5.6% (339)
... give birth and desired # of children	-	3.0% (303)
... give birth but not the desired # of children	-	3.3% (303)
C. **Benefit Use**	Males	Females
Planning to spend MC benefits in		
... improvement of family’s living conditions	47.6% (1,992)	46.4% (3,842)
... children’s education	48.4% (1,992)	49.4% (3,842)
... cumulative part of mother’s pension	4.0% (1,992)	4.2% (3,842)
Subsample of eligible women:		
... improvement of family’s living conditions	-	62.7% (324)
... children’s education	-	34.3% (324)
... cumulative part of mother’s pension	-	3.1% (324)

Note: Questions on MC policy asked to all adults in RLMS sample. Number of observations in parentheses.

We use the family roster to create a fertility history for each woman in the panel. Specifically, we record a birth every time a new child appears in the household roster. For households interviewed for the first time, we record a birth if the child is less than one year old. Because interviews are conducted between October and December, fertility measures using the RLMS do not exactly correspond to a calendar year. Below (figures 1 and 2) we compare fertility measures using the RLMS and official sources. As can be expected from a representative sample, the RLMS data is noisier but follows the official statistics quite closely.

2.2 Poll Data

The 2008 round of the RLMS included a series of questions on the MC policy that were administered as part of the adult questionnaire. Table 1 presents a summary of the poll results. We present results related not only to the program's effectiveness but also to the degree to which the public was informed about it and believed that the promised benefits would actually be provided to eligible applicants. These other aspects of the program's implementation will play a role in our modeling strategy below.

First, respondents were asked whether they knew anything about the policy.[9] Almost two years after the program was introduced, roughly two thirds of women and half the men answered affirmatively. Panel A also presents the percent of correct answers to a series of more detailed questions regarding the program (asked only to those who claimed to know something about it). In general, respondents (specially women) seem to have a good grasp of the basic program rules. Another item asked eligible women[10] whether they had done the paperwork to obtain their MC certificate. Roughly half of them answered affirmatively. Finally, about one third of respondents did not believe government would follow through with the promised transfer to eligible mothers. What we infer from these responses is that knowledge about the MC policy is less than universal and that some informed individuals do not believe the program benefits will materialize.

Panel B summarizes responses to questions on the effectiveness of the program. Only about 3% of males and 4% of females claimed that MC had led them to reconsider the number of children they would like to have. The difference in the number of desired children between those who claimed to have been influenced and those who claimed not to have been influenced is 0.2 among women and 0.3 among men. The questionnaire also included a number of items that were asked to eligible women only. Only 5.6% of them answered that the MC program had influenced their decision to have one more child. Responses to subsequent questions imply that in many cases the program only influenced the chosen timing of the birth and not the total number of desired children. If taken at face value, these responses would suggest a negligible effect of the program on fertility.

MC benefits can be given a limited number of uses (only a small fraction can be received in cash). Panel C presents responses to a series of questions on the planned use of program benefits. Among the overall population, children's education and housing improvements are equally favored options. Very few respondents indicated that they would use the benefits to make

[9]The exact wording to these questions (both in Russian and in English translation) is available at www.cpc.unc.edu/projects/rlms-hse/data/questionnaires.

[10]For the purpose of the poll, eligible women are those who had given birth to a second or consecutive child since 2007.

Figure 1 – Birth Rates for Women Ages 15–49

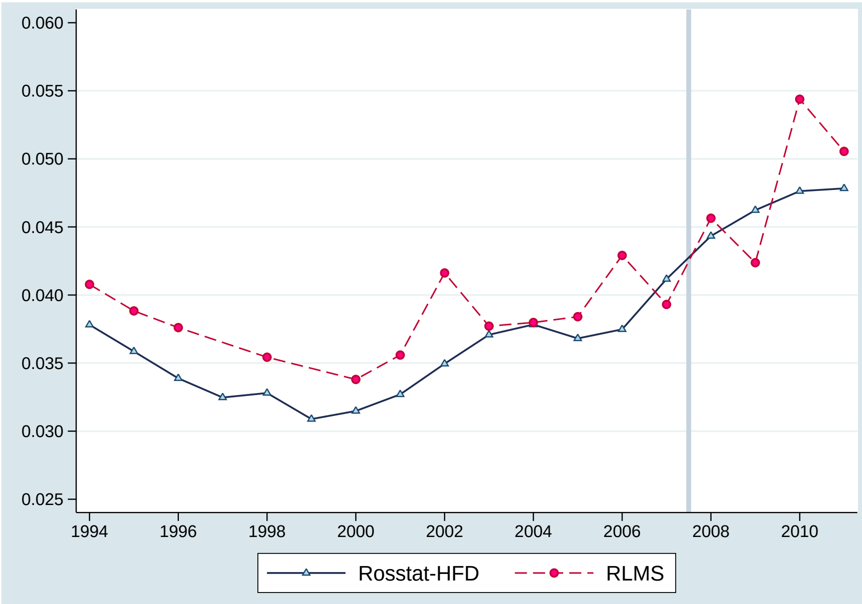

Notes: The data for the Rosstat-HFD series is from the Human Fertility Database for 1994–1999 and from Rosstat for 2000–2011.

an extra contribution to the mother's pension fund. Among eligible women there is a stronger preference for housing over the other options. In any case, these responses are quite different from the official statistics on program benefit use (98% housing improvement). This shows that poll results need to be interpreted cautiously as it is not rare for planned and realized behavior to differ.

2.3 Aggregate Trends

How does actual fertility behavior compare to people's opinions about it? Figure 1 shows official statistics on birth rates (BR) for women ages 15 to 49. The Rosstat-HFD series shows a declining trend during the late 90s. After a quick bounce back in the early 2000s, births per woman stabilized around a level of 3.7% in the years before the MC program was introduced. Starting in 2007, the BR has increased steadily and reached the highest level in the period under analysis. Note, however, that this maximum is still a low point in historical perspective (see figure A.1 in the appendix).

The BR is highly sensitive to changes in the age composition of the female population. An alternative indicator of overall fertility is the total fertility rate (TFR), which is the sum of the age-specific birth rates at a point in time. The TFR is independent of relative cohort sizes and measures

Figure 2 – Total Fertility Rate

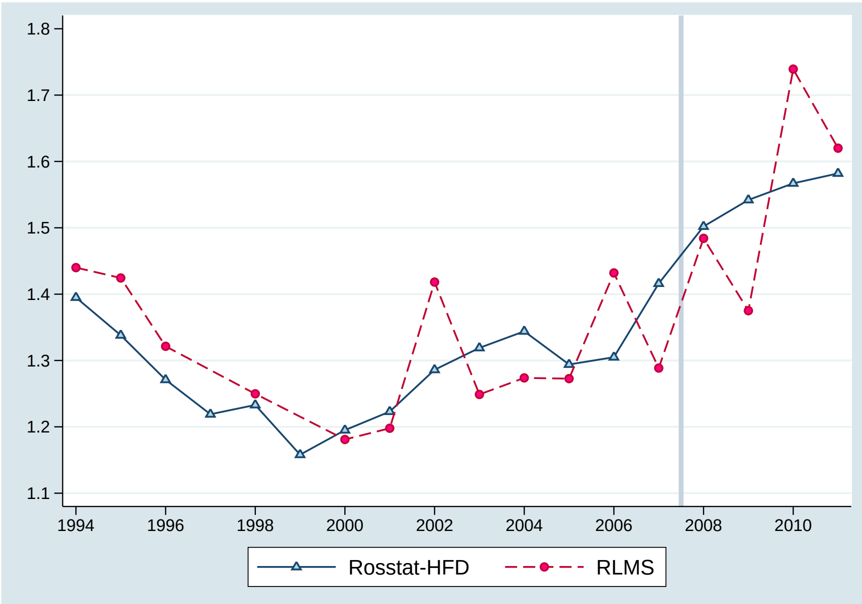

Notes: The data source is the same as in figure 1. The TFR is the sum of all age-specific fertility rates at a point in time.

fertility in an easy to interpret metric (number of children per woman). Its main shortcoming is that it correctly measures completed fertility only under the strong assumption that previous cohort's age-specific birth rates can be extrapolated to women who have not completed their fertile period. According to the official statistics, the TFR in Russia followed a path similar to the birth rate (see figure 2).

It is far from clear whether the recent increases in the BR and the TFR are related to the implementation of the MC program. First, the ascending trend in fertility seems to have begun before the program was in place. Second, even the more recent increases might be due to other factors. In particular, it could be the case that the MC policy encouraged a shift in the timing of births without actually affecting total number of children desired.[11] However, the aggregate time series evidence has given some hope to the supporters of the MC policy.

Figure 3 shows birth rates by birth order using the RLMS data.[12] An item in the adult questionnaire that asks for the number of children the person has can be used to determine birth order but is only available start-

[11] In demographers' jargon, the MC policy could have affected the "tempo" of births without changing their "quantum".

[12] Birth rates by birth order are also available from the HFD. They closely follow the RLMS series and are omitted to keep the figure uncluttered.

Figure 3 – Birth Rates By Birth Order for Women Ages 15-49

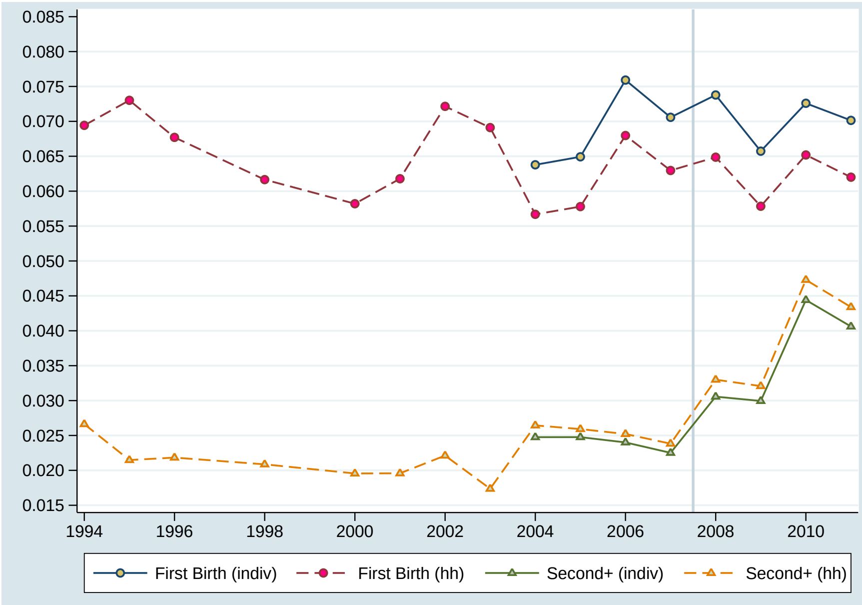

Notes: The data source is the RLMS. The series based on the individual questionnaire (indiv) use an item asking for the number of children the respondent has. The series based on the family questionnaire (hh) use the number of sons and daughters living in the household. "Second+" refers to second and consecutive birth orders.

ing in 2004. The family questionnaire provides information on the number of sons and daughters for all years on the condition that they live in the household. The discrepancy between the two series is caused by births by relatively older women with children outside the household.[13]

Regardless of which source is used to determine birth order, it is clear that the rate of first births was fairly constant over the period and that the increase in the birth rate in recent years is due exclusively to second and consecutive births. Since the MC policy explicitly targeted women with at least one child, the evidence on birth rates of different orders can in principle be counted as in favor of a positive effect of the program on fertility.

2.4 Before-After and Difference-in-Differences Estimates

It is possible to use the individual level data from the RLMS to obtain more formal tests of the hypothesis that the MC policy increased fertility, while also controlling for a number of observable characteristics. Specifically, we estimate the following equations:

[13] Figure A.2 in the appendix shows the same series restricted to younger women.

$$
\begin{aligned}
birth_{it} =& \alpha_0 + \alpha_1 post_t + f_1(age_{it}) + \boldsymbol{X}\boldsymbol{\gamma_1} + \varepsilon_1 && (1)\\
birth_{it} =& \beta_0 + \beta_1 post_t + \beta_2 MCelig_{it} + \beta_3 \left(MCelig_{it} \times post_t\right)\\
&+ f_2(age_{it}) + \boldsymbol{X}\boldsymbol{\gamma_2} + \varepsilon_2 && (2)
\end{aligned}
$$

where *birth* is an indicator of whether woman i gave birth in time t and *post* is an indicator equal to one for the years 2008–2011 and zero otherwise. Note that while the policy was put in place in January 2007, a large majority of the births observed in that year's RLMS interview correspond to pregnancies from 2006.[14] The variable *MCelig* equals one for women with one or more children and who have not given birth to a second or consecutive child after 2007. Correspondingly, it is zero in two cases: 1) for women without children, and 2) for women with two or more children after the birth of a second or consecutive child in the program period.[15]

We estimate equation (1) for MC-eligible women only. Under the strong assumption that no other important unobservable determinant of fertility changed at the same time as the MC policy was implemented, the before-after (BA) comparison summarized by the α_1 parameter identifies the effect of the program.

An alternative strategy is to apply difference-in-differences (DID). The β_3 parameter in equation (2) identifies the causal effect under the well-known "common trends" assumption, namely that any time-varying unobservables have the same effect on treated and non-treated women.[16] Note that this estimate is a lower-bound since the MC program could also incentivize first births by increasing the option-value of a second birth later on. We further discuss these assumptions in the next sub-section.

Columns numbered 1 through 7 in table 2 present OLS estimates of the treatment effects for different specifications of the age function ($f(\cdot)$) and the set of control variables ($\boldsymbol{X}$).[17] Standard errors are robust and clustered at the individual level, as recommended by Bertrand et al. (2004).

The BA and DID estimates are very similar and robust to the inclusion of a wide variety of controls. We have applied special care to controlling for age of the mother, which is obviously an important factor determining

[14] Using the child's month of birth, we constructed a quarterly birth rate series (available upon request from authors). We found no evidence of an increase in births of any order in the last quarter of 2007 (in fact, the birth rate in that quarter is substantially lower than in quarters 1–3).

[15] Since we do not have information on program take-up, the best we can do is look at the effect of eligibility. This is equivalent to intent-to-treat analysis in clinical trials.

[16] This is only approximately equivalent to the assumption that unobservables affect first births and higher birth orders equally since our treatment indicator also depends on whether the woman gave birth after the program was in place.

[17] Table A.1 in the appendix presents the full regression results for specification 7. Other estimation results are available from the authors.

fertility behavior. We experimented with a quadratic functional form, as well as with the inclusion of age-group dummies, and a linear and a cubic spline. We have also included marital status, a number of demographic and human capital characteristics, household composition, and year and location dummies.[18] Finally, in some specifications we included controls for the age of the youngest child, as well as a series of indicators for the availability and cost of child care in the locality where the woman lives.

The estimated effect on the birth probability always falls in the range 1.6–2.4%, which is consistent with figure 3. We can obtain a back-of-the-envelope estimate of the expected increase in the number of children as follows. The life period we analyze is 34 years long. Our estimates indicate that the average treated woman's birth probability increased by approximately 2% per annum. This gives an expected increase of $0.02 \times 34 \approx 0.68$ children over the life period considered. Since in the post-reform period roughly 58% of women were treated, the implication is that the program led to an estimated overall increase in fertility of about 0.4 children. This is consistent with the increase in TFR shown in figure 2.

In sum, these estimates tell us that the increase in birth rates observed after 2007 is not an artifact of changes in age composition or in any of the other observable determinants of fertility. It also confirms that the increase in second and consecutive births is statistically significant.

As a final robustness check, we obtained BA and DID estimates based on a nearest-neighbor matching technique. Specifically, we calculated

$$\widehat{BA}_{match} = \frac{1}{n_{T\cap P}} \sum_{i \in T\cap P} \left[birth_i - birth^{nn(i)}_{T\cap\overline{P}} \right] \tag{3}$$

$$\begin{aligned} \widehat{DID}_{match} =& \frac{1}{n_{T\cap P}} \sum_{i \in T\cap P} \left[birth_i - birth^{nn(i)}_{\overline{T}\cap P} \right] \\ & - \frac{1}{n_{T\cap\overline{P}}} \sum_{i \in T\cap\overline{P}} \left[birth_i - birth^{nn(i)}_{\overline{T}\cap\overline{P}} \right] \end{aligned} \tag{4}$$

where T is the set of MC-eligible women, P is the set of women observed in the program years, and n_x is the number of individuals in set x. The BA matching estimator in equation (3) compares births by MC-eligible women post 2007 with their matched counterparts in the pre 2008 period (the set $T \cap \overline{P}$). The match is chosen by the nearest neighbor criterion applied to the propensity score metric.[19] Matching is done without replacement and ties are resolved by randomization. We impose the common support condition.[20] The propensity score is estimated via a logit model with *post*

[18] Table 2 contains the full list of control variables.

[19] We experimented with other matching techniques and obtained similar results.

[20] The region of common support is the subset of values of the controls that are observed

as the dependent variable and age plus all the baseline controls in the right-hand-side (see table 2 for a list; year dummies are excluded).

The DID estimator in equation (4) is based on double matching (see Smith and Todd, 2005). MC-eligible women before and after the program was effective are matched to ineligible women. The matching criterion is again nearest neighbor. The propensity score is estimated on the P and $\overline{P}$ sub-samples by logit models with MC-eligibility as the dependent variable and age and all baseline controls (including year) in the right-hand-side.

Matching-based estimates are presented in the last column of table 2. These estimates are statistically undistinguishable from the ones obtained via regression. The OLS estimator imposes a restrictive linear-additive form to the control function. In addition, OLS does not restrict estimation to the region of common support, so an unknown level of extrapolation goes undetected. Given that the matching estimates are very close to the regression estimates, we can rule out that the findings in this sub-section are due to the specifics of OLS regression. However, it should be emphasized that matching is no silver bullet. In particular, the identification assumptions behind the BA and DID strategies are fundamentally the same in all cases.

2.5 Discussion

In this section, we have presented diverse evidence on the effectiveness of the MC policy. First, poll results show negligible effects of the program. Second, aggregate time series evidence based on registry data show a significant increase both in the BR and in the TFR. The increase in births is explained exclusively by second and higher birth orders. Finally, BA and DID estimates show that the increase in fertility is not due to changes in any of a large number of control variables.

It is clear that evidence based on actual births carries more weight than poll responses that at best reflect the respondents' plans. However, we think the poll results are a reason for concern. In particular, a significant fraction of MC-eligible respondents indicated that the program influenced their decision to give birth but not the total number of desired children. Moreover, previous studies have found the rescheduling of births to be a common response to pronatalist policies (eg. Dickert-Conlin and Chandra, 1999, Gans and Leigh, 2009). Neither the BA nor the DID strategy is robust to unobservable changes in the "tempo" of births.

The BA and DID results have a causal interpretation only under strong assumptions. Although we are able to control for age, marital status, and a number of other observable determinants of fertility, there are important unobserved factors that could be behind the increase in birth rates post 2007.

in both groups under comparison.

For the BA strategy, the timing of birth or any other change in unmeasured conditions that led to increasing fertility around the time the program was put in place could explain the increases in the BR. To cite an example, in 2007 maternity leave and child care benefits were raised. While the amount of the rise was not significant enough to influence fertility behavior, the BA comparison incorrectly assigns any response to this or other policies to the MC program.

The DID strategy can be expected to be more robust in this respect, since unobservable factors affecting second and consecutive births should in principle also affect first births. However, it can also be argued that the decision to give birth to a first child is less sensitive to economic factors than the decision to have two or more children. In other words, the common trends assumption might not apply in this case.

Our response to this methodological challenges it to model the fertility decision in a dynamic setting that allows for differential costs and benefits for births of different orders, as well as tradeoffs between family and work life. The main advantage of this strategy is that the model explicitly takes into account women's optimal timing of births when providing estimates of the long run changes in fertility.

Table 2 – Before-After and Diff-in-Diff Estimates: women 15–49

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Before/After Change	0.016*** (0.002)	0.017** (0.008)	0.017** (0.008)	0.017** (0.008)	0.017** (0.008)	0.019** (0.008)	0.017* (0.009)	0.016*** (0.003)
Difference in Differences	0.018*** (0.004)	0.019*** (0.004)	0.021*** (0.004)	0.022*** (0.004)	0.022*** (0.004)	0.024*** (0.004)	0.023*** (0.005)	0.024*** (0.005)
Age Control	NO	Quadratic	Groups[e]	Linear Spline[e]	← Cubic Spline[e] →			Propensity Score NN Matching
Baseline Controls[a]	NO	YES	YES	YES	YES	YES	YES	
Other Controls	NO	NO	NO	NO	NO	YES[b]	YES[b,c]	NO
Observations[d]	36,659/60,051	33,801/55,535	33,798/55,531	33,801/55,535	33,801/55,535	33,801/55,535	32,755/53,463	33,803/55,535
R-squared[d]	0.002/0.009	0.030/0.063	0.031/0.067	0.032/0.070	0.031/0.069	0.034/0.074	0.036/0.076	

Notes: The dependent variable is *birth*. The before/after regression is restricted to MC-eligible women. Robust standard errors clustered at individual level; matching s.e. are bootstrapped. [a]Baseline controls (in parenthesis the # of categories for discrete controls): Number of children (5), Education (5), Experience, Russian national, Russian Born, Married, Parents in HH (3), Grand Parents in HH (3), Siblings in HH (5), Year (16), Location (39), Urban location. Other controls: [b]Age Group of Youngest Child (5); [c]Local State Nursery, Local Private Nursery, Local State Pre-school, Local Private Pre-school, Local median amount paid for hourly childcare, Local fraction that receive free care. [d]Figures corresponds to before-after and diff in diff respectively. [e]Age groups are 15–20, 20–25, . . . , 45–50. The same intervals are used to choose the knots of the splines.

3 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 ($l_t = 1$).[21] 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.[22] We index the four mutually exclusive alternatives facing women by j:

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

We let the decision process start at age 22, set the end of the fertile period at age 45, and end the program at the official retirement age of 55.[23] The starting point is convenient since a vast majority of Russian women finish their education by age 22.[24] 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

$$\mathrm{E}\left[\sum_{t=22}^{54} \rho^{t-22} U_t(c_t, l_t, l_{t-1}, n_t, n_{t-1}, n_{t-2}, X_{t-1}, N_t, B_t, S, m_t, p_t)\right]$$

where ρ is the subjective rate of discount and the expectation is taken over the stochastic components of utility and earnings.[25] Women derive utility

[21] 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. We emphasize that from here on t stands for the woman's age, and not calendar time.

[22] 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. In our model, the utility of births has a random component. These two modeling strategies are not separately identifiable.

[23] The estimation process ignores any fertility decision after the cutoff age. The last decision period is $t = 54$.

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

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

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), the age of the youngest child (B_t), as well as lagged employment and births —all of them results from past decisions— are assumed to affect current tastes. Finally, utility is affected by the woman's education (S), marital status (m_t), and whether she cohabits with parents (p_t). While education remains constant over time, marital status and cohabitation with parents are assumed to evolve following a first-order markovian process whose underlying parameters are allowed to change as the woman ages.[26] The specific functional form for the utility function is

$$\begin{aligned}U_t =&c_t + \alpha_1 l_t + (\alpha_2 + \epsilon_t^n)\, n_t + \alpha_3 I_{N_t=1} + \alpha_4 I_{N_t=2} + \alpha_5 I_{N_t>2} + \beta_1 c_t l_t \\ &+ n_t\left[\beta_2 c_t + \beta_3 l_t + \beta_4 l_{t-1} + \beta_5 n_{t-1} + \beta_6 n_{t-2} + I_{N_t>1}\left(\beta_7 + \beta_8 \epsilon_t^n\right)\right] \\ &+ m_t\left[\delta_2 n_t + \delta_3 I_{N_t=1} + \delta_4 I_{N_t=2} + \delta_5 I_{N_t>2} + l_t\left(\delta_1 + \delta_6 I_{B_t<4} + \delta_7 I_{B_t\geq 4}\right)\right] \\ &+ l_t\big[\gamma_1 X_{t-1} + \gamma_2 S_1 + \gamma_3 S_2 + \gamma_4 S_3 + \gamma_5 S_4 \\ &+ \gamma_6 I_{N_t=1} + \gamma_7 I_{N_t=2} + \gamma_8 I_{N_t>2} + \gamma_9 I_{B_t<4} + \gamma_{10} I_{B_t\geq 4}\big] \\ &+ p_t\left[\mu_2 n_t + \mu_3 I_{N_t=1} + \mu_4 I_{N_t=2} + \mu_5 I_{N_t>2} + l_t\left(\mu_1 + \mu_6 I_{B_t<4} + \mu_7 I_{B_t\geq 4}\right)\right]\end{aligned} \tag{5}$$

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. Cohabitation with parents enters in a similar fashion. The disutility of work depends on previous work experience (habit formation), highest education completed[27], and the number of children. Finally, we let the work interactions depend on the presence of a small child (3 years old or less) or an older 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:

[26]Specifically, we allow transition probabilities to differ between women in different age intervals. The transition matrices are estimated outside the model (see table A.2 in the appendix for the estimated transition probabilities).

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

$$c_t = y_t^f l_t + y_t^h m_t + y_t^o + \phi MC n_t K$$
$$- b_1 l_t - b_2 n_t - b_3 I_{N_t=1} - b_4 I_{N_t=2} - b_5 I_{N_t>2} \tag{6}$$

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, income from their partners y_t^h when cohabiting, and income from other household members y_t^o. In addition, eligible women ($MC = 1$) receive maternity capital assistance in the amount K if they give birth.[28] Because assistance can only be obtained three years after the birth and must be used for specific purposes, we estimate a parameter (ϕ) that converts assistance dollars into a monetary equivalent consumption value.[29]

The woman's income from other household members depends on her characteristics. Women are assumed to form expectations according to

$$\overline{\log y_t^h} = c_0 + c_1 t + c_2 t^2 + c_3 S_1 + c_4 S_2 + c_5 S_3 + c_6 S_4 + c_7 G \tag{7}$$
$$\overline{\log y_t^o} = d_0 + d_1 m_t + d_2 p_t + d_3 t + d_4 t^2 + d_5 S_1 + d_6 S_2 + d_7 S_3 + d_8 S_4 + d_9 G \tag{8}$$

where G indicates urban residence. Equations (7) and (8) do not depend on current or future decisions so they are estimated outside the model.[30] Note that non-labor income depends on the random states (m_t, p_t), so women use the transition probabilities in table A.2 to estimate the expected value.

Women receive a job offer with a probability (π_t) that depends on whether they were employed the previous period and whether they reside in an urban area. Formally,

$$\pi_t = \frac{\exp(\psi_t)}{1 + \exp(\psi_t)} \tag{9}$$
$$\psi_t = z_0 + z_1 l_{t-1} + z_2 G$$

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

[28] We set $K = 365,698$, the average real value (in rubles of year 2011) of MC assistance over the period 2007–2011.

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

[30] See table A.3 in the appendix for the estimated coefficients.

$$\log y_t^f = a_0 + a_1 X_{t-1} + a_2 X_{t-1}^2 + a_3 S_1 + a_4 S_2 + a_5 S_3 + a_6 S_4 + a_7 G + \epsilon_t^y \tag{10}$$

The shock ϵ_t^y captures variation in earnings that is independent of the decision process.[31] The two shocks $(\epsilon_t^n, \epsilon_t^y)$ 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_2), utility associated with having children (α_3–α_5,δ_3–δ_5), the baseline job offer probability (z_0), the baseline earnings (a_0), and the standard deviation of the birth shock (σ^n). 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 and parental cohabitation 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 Ω_t.[32]. The value function $V(\Omega_t)$ is the maximal expected present value of the remaining lifetime utility given the state at age t.[33] 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 $t = 22, \ldots, 45$ and $J_t = \{1, 2\}$ for $t = 46, \ldots, 54$. The alternative-specific value functions obey the Bellman equation:

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

Finally, the pre-determined state variables evolve according to

[31] Table A.4 presents OLS estimates of the earnings regression and logit estimates of the employment probability. We use these coefficients as starting values in the ML search.

[32] $\Omega_t = (l_{t-1}, N_{t-1}, X_{t-1}, n_{t-1}, n_{t-2}, B_t, S_1, \ldots, S_4, G, \overline{y_t^h}, \overline{y_t^o}, m_t, p_t, MC, \epsilon_t^n, \epsilon_t^y)$

[33] 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).

$$
\begin{aligned}
N_t &= N_{t-1} + n_t \\
X_{t-1} &= X_{t-2} + l_{t-1} \\
B_t &= \begin{cases} 1 & \text{if } n_{t-1} = 1 \\ B_{t-1} + 1 & \text{otherwise} \end{cases}
\end{aligned}
$$

3.1 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.[34] 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.[35] 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.[36]

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

[34] Marital status and cohabitation with parents are included in Ω_t^d.

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

[36] The solution for women 45 years old and younger is more computationally demanding since it involves the doubling of the decision tree that must be considered.

$$\Pr(d_{i,t} \mid \Omega_t^d) = \Pr\left(j = \arg\max_k V_k(\Omega_t)\right) \qquad \text{for } j = 1,3$$

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

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

$$\Pr(d_{i,22}, \ldots, d_{i,54} \mid \Omega_{22}^d) = \prod_{t=22}^{54} \Pr(d_{i,t} \mid \Omega_t^d) \tag{11}$$

In turn, the likelihood for a sample of women is simply the product of (11) over the N women in the sample. In order to generate the probabilities in the right hand side of (11), we solve the dynamic program for 30 simulations of the random shocks and use a kernel smoothing function (McFadden, 1989). 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

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

where $\boldsymbol{\theta}$ is the vector of parameters, including the errors variance-covariance matrix and the type proportions (κ_h).[37]

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

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

[37]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.[38]

4 Estimating Sample

In this section we provide details on how we processed the RLMS data in order to estimate the structural model.

4.1 Variable Definitions

Employment The RLMS contains information on a main job and a secondary job.[39] 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.[40] In subsequent rounds we let experience evolve in a way that is consistent with the observed employment history.

Births and number of children As already mentioned, whether a woman has given birth during the year preceding interview is determined on the basis of the household roster. 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.

[38]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.

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

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

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 or partner, 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 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 2011 using the Russian CPI.

4.2 Sample Selection and Descriptive Statistics

Our sample is composed of women between 22 and 54 years of age. The model requires accurate information on the number of children each woman has regardless of her age. These data are available for the years 2004–2011. After deleting observations with missing values in the relevant variables, our unbalanced panel comprises 9,321 individuals and a total of 31,911 person-year observations. Table 3 has descriptive statistics.

In our model, women's fertile period ends exogenously at age 45. Over 74% of individuals enter our sample before crossing this threshold. Women in the sample exhibit wide variation in initial labor market experience and education attainment.

Women in our sample have completed fertility rates significantly below the replacement rate. For example women over 40 have on average 1.72 children. Low fertility rates occur despite the fact that Russia has one of the highest marriage rates in the world.

Table 4 focuses on the relationship between marital status, family size, and the decision variables (employment and fertility). 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

Table 3 – Descriptive Statistics

	Mean	Std Dev
Individuals (9,321 observations)		
Years in sample	3.4	2.5
Age in first period	35.6	10.3
Experience in first period	13.0	10.3
Less than Secondary Educ	6.3%	
Secondary Educ Complete	17.1%	
Vocational School Complete	18.1%	
Technical School Complete	27.1%	
University Degree or above	31.4%	
Person-year (31,911 observations)		
Age	37.6	9.7
Married	67.5%	
Cohabits with Parents	28.2%	
Birth	0.036	
No Children	18.2%	
One Child	40.9%	
Two Children	32.6%	
Three Children	6.6%	
Four+ Children	1.7%	
Employed	73.4%	
Experience	14.4	10.0
Labor Income	8,861	10,840
Spouse's Income	10,533	16,067
Other Income	10,847	30,039
MC Eligible (2008–2011)	52.4%	

Notes: Income variables in 2011 rubles.

Table 4 – Employment and Births by Marital Status and Number of Children

# of	Single/Not Cohab.			Married/Cohab.		
children	% Empl.	% birth	Obs	% Empl.	% birth	Obs
0	69.1	2.33	3,598	71.6	16.68	2,751
1	80.1	1.02	4,125	75.3	4.72	8,827
2	79.6	0.33	2,146	74.8	1.20	7,910
3	65.9	1.01	398	57.5	1.41	1,653
4+	48.1	1.85	106	34.2	3.88	396
Total	75.3	1.34	10,373	72.5	4.69	21,538

Note: Number of children does not include recent born.

Table 5 – Employment and Births: rural vs. urban

# of	Rural Location			Urban Location		
children	% Empl.	% birth	Obs	% Empl.	% birth	Obs
0	55.4	8.21	1,254	73.8	8.64	5,095
1	69.1	5.17	2,573	78.7	3.14	10,379
2	67.4	1.05	2,964	79.4	1.00	7,092
3	47.7	1.72	987	69.9	0.96	1,046
4+	35.7	4.87	308	39.2	1.42	212
Total	62.5	3.71	8,087	77.1	3.57	23,824

Note: Number of children does not include recent born.

of one or two children are more likely to be employed than women without children. Only after the third birth does participation decline significantly. Predictably, the probability of a birth is always higher for married women (inclusive of cohabitation). The table shows that the birth counting process is non-linear. It is highest for women without children. It then decreases monotonically for women with one and two children but picks up again for women with three and more children.

Women in rural areas have lower participation rates and higher birth rates (table 5). In particular, birth rates for high birth orders are much lower in cities.

Finally, table 6 shows the evolution of women's choices over their life-cycle. 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 about 60% and increases 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 much later in life.

Table 6 – Choice Distribution

Age Group	Non-employed No Birth	Non-employed Birth	Employed No Birth	Employed Birth	Total
22–24	1,105	136	1,823	138	3,202
	34.51	4.25	56.93	4.31	100.0
25–27	825	115	2,088	173	3,201
	25.77	3.59	65.23	5.41	100
28–30	695	90	2,030	152	2,967
	23.42	3.03	68.42	5.12	100
31–33	731	62	2,116	104	3,013
	24.26	2.06	70.23	3.45	100
34–36	623	29	2,226	71	2,949
	21.13	0.98	75.48	2.41	100
37–39	600	18	2,142	39	2,799
	21.44	0.64	76.53	1.39	100
40–44	919	9	3,254	13	4,195
	21.91	0.21	77.57	0.31	100
45–49	1,022	0	3,573	1	4,596
	22.24	0	77.74	0.02	100
50–54	1,507	0	3,482	0	4,989
	30.21	0	69.79	0	100
Total	8,027	459	22,734	691	31,911
	25.15	1.44	71.24	2.17	100

Note: Number of observations and percentages.

5 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 7.

- α_1, the disutility of work, is negative as expected. In addition, working implies giving up around 3% of consumption (this suggests consumption and leisure are complements). Note that working married women do not experience significantly lower utility (δ_1 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.

- Children 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 low return to on the job experience (two percent, similar to the OLS estimates is table A.4)
- The multipliers associated with MC policy are around 1% regardless of the woman's type.
- Obtaining standard errors is computationally time consuming. They will be provided at a later stage.

Figures 4 and 5 show the ability of the model to reproduce key aspects of the data.

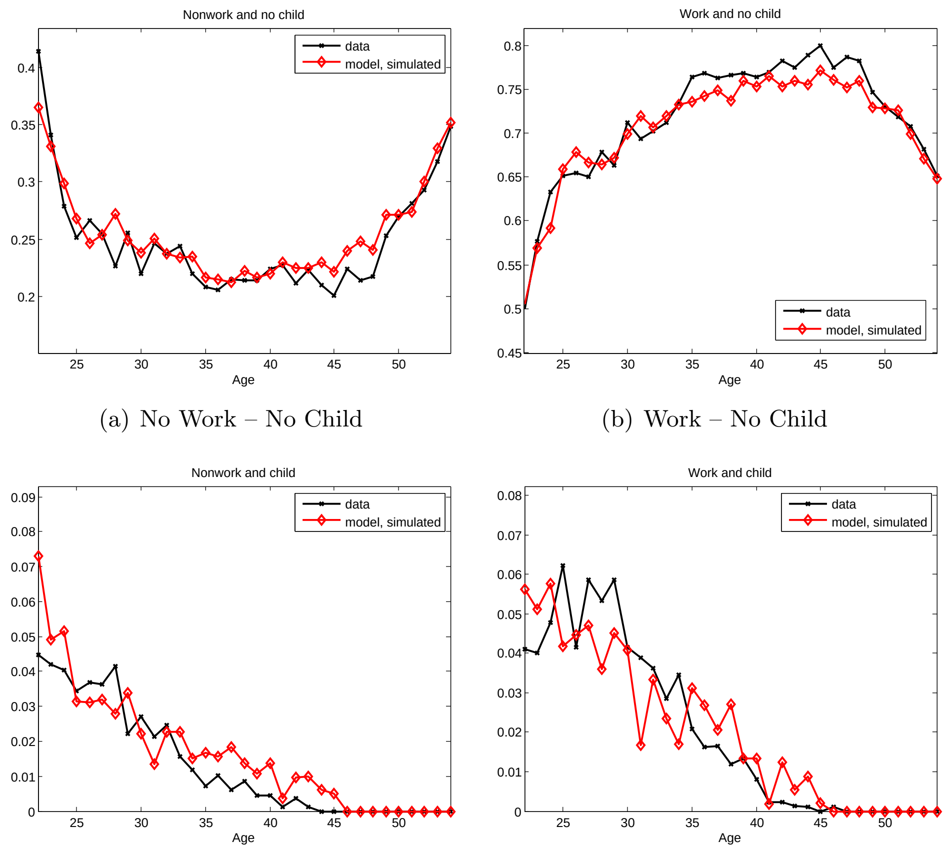

(a) No Work – No Child
(b) Work – No Child
(c) No Work – Child
(d) Work – Child

Figure 4 – Model Fit for Mutually Exclusive Choices

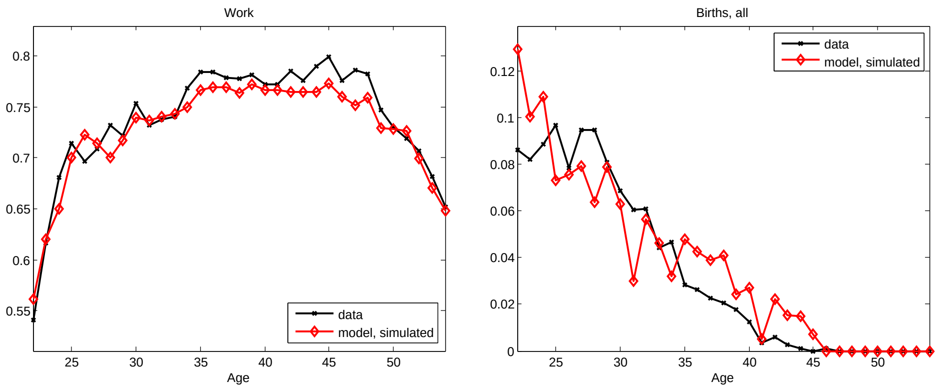

(a) LF Participation (b) Total Births

Figure 5 – Model Fit for LF Participation and Total Births

Table 8 shows transition probabilities among the mutually exclusive choices for women ages 22–45 and 46–54 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 9 shows that ability of the model to fit behavior [by unobservable type forthcoming...]. 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.

Table 7 – Maximum Likelihood Estimates

Utility Function			
α_1	-7,645.874	β_1	-0.033
α_2 $(type = 1)$	-133,524.835	β_2	-0.042
α_2 $(type = 2)$	-118,074.870	β_3	444.281
α_2 $(type = 3)$	-125,019.418	β_4	-1597.326
α_3 $(type = 1)$	5447.009	β_5	-4178.833
α_3 $(type = 2)$	6236.358	β_6	-925.088
α_3 $(type = 3)$	5,562.133	β_7	-6720.012
α_4 $(type = 1)$	7,732.085	β_8	1.080
α_4 $(type = 2)$	9,208.776	γ_1	201.288
α_4 $(type = 3)$	8,215.993	γ_2	-232.071
α_5 $(type = 1)$	12,571.802	γ_3	-34.670
α_5 $(type = 2)$	10,709.053	γ_4	-598.013
α_5 $(type = 3)$	11,499.079	γ_5	-6,419.642
δ_1	-193.904	γ_6	-3,675.894
δ_2 $(type = 1)$	19,518.675	γ_7	-4,448.491
δ_2 $(type = 2)$	21,822.785	γ_8	-7,427.210
δ_2 $(type = 3)$	21,229.702	γ_9	-1,087.545
δ_3 $(type = 1)$	4,070.991	γ_{10}	499.685
δ_3 $(type = 2)$	4,784.797	μ_1	0.000
δ_3 $(type = 3)$	4,621.392	μ_2	493.950
δ_4 $(type = 1)$	5,534.788	μ_3	540.138
δ_4 $(type = 2)$	5,938.933	μ_4	606.865
δ_4 $(type = 3)$	5,976.404	μ_5	795.341
δ_5 $(type = 1)$	7,929.693	μ_6	1495.582
δ_5 $(type = 2)$	7,647.865	μ_7	426.565
δ_5 $(type = 3)$	7,958.517		
δ_6	209.752		
δ_7	744.689		
Labor Income		*Error Structure*	
a_0 $(type = 1)$	11.666	$\sigma_n(type = 1)$	25175.426
a_0 $(type = 2)$	10.202	$\sigma_n(type = 2)$	24202.447
a_0 $(type = 3)$	10.891	$\sigma_n(type = 3)$	23186.271
a_1	0.023	σ_y	0.283
a_2	0.000	σ_u	0.338
a_3	0.076	$\rho_{n,y}$	-0.251
a_4	0.016	*MC Policy*	
a_5	0.095	$\phi(type = 1)$	0.009
a_6	0.466	$\phi(type = 2)$	0.010
a_7	0.390	$\phi(type = 3)$	0.010
$z_0(type = 1)$	-1.628	*Type Proportions*	
$z_0(type = 2)$	-2.890	κ_1	0.2659
$z_0(type = 3)$	-1.164	κ_2	0.3753
z_1	3.785	κ_3	0.3589
z_2	0.400	$\log \mathrm{L}$	-38,978

Table 8 – Transition Probabilities: data vs. model

	Ages 22–45				Ages 46–54	
	no birth no work	no birth work	birth no work	birth work	no birth no work	no birth work
no birth no work	0.6635	0.2936	0.0342	0.0086	0.7818	0.2182
	0.5956	0.3347	0.0536	0.0161	0.7323	0.2677
no birth work	0.0819	0.8768	0.0077	0.0336	0.0851	0.9149
	0.1085	0.8516	0.0087	0.0311	0.1271	0.8729
birth no work	0.6770	0.2422	0.0745	0.0062		
	0.7029	0.2508	0.0374	0.0089		
birth work	0.1237	0.8281	0.0063	0.0419		
	0.2251	0.7346	0.0145	0.0258		

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

Table 9 – Data versus Model

	Births (per 1,000)	Participation Rate
All	36.04	73.41%
	38.78	72.50%

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

6 Simulations, preliminary conclusions, and model reformulation

Having estimated the parameters of the model, it is possible to compare the fertility behavior of women under different scenarios. In particular, we are interested in comparing behavior in a counterfactual in which the MC policy is not present. For this purpose, we simulate the lifetime decisions of a group of 22 year old women with the same initial characteristics as those found in our estimating sample. The simulations results are averages over 100 different draws from the joint distribution of the shocks. With the MC policy present, we find that on average women give birth to 1.752 children and work 21.85 years. In contrast, a simulation setting the ϕ parameter to zero yields 1.697 children and 21.92 years of work. In sum, the model predicts that the MC policy has a very small effect on the fertility and labor force participation decisions.

On one hand, this result is not too surprising considering the very low estimated value of the ϕ parameter. On the other hand, it is natural to wonder how the model's long run prediction can be reconciled with the evidence presented in section 2. The BA and DID regressions estimated an increase in the birth rate for second and consecutive children of about 2%. In order to better understand the predictions of the model, we conducted a different type of simulation.

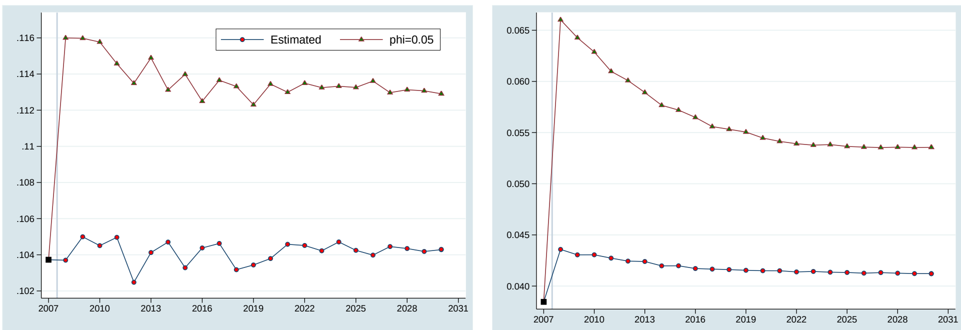

(a) First Birth

(b) Second+ Births

Figure 6 – The Effect of the MC Program in Calendar Time

First, we take a group of women with the same distribution of age and other characteristics as the one found in 2007 and simulate their decisions in the absence of the MC program. Next, we introduce the MC program and simulate their decisions for 23 periods. In order to keep the age composition constant, we replenish the sample with new 22 year old women every year. This simulation is meant to capture the short run dynamics predicted by the model right after the policy is launched.

Figure 6 presents the simulation results for two scenarios. In the first, the ϕ parameter takes the value estimated by the model. In the second, we set $\phi = 0.05$ for every woman type. Panel (a) shows that the model predicts no change in the birth rate for first children, while panel (b) shows a predicted immediate increase in second births of about half a percentage point. While the model misses in the level of these changes, it gets the story right qualitatively by predicting an increase only in second and consecutive births. In contrast, setting ϕ equal to a higher (but still quite low) number leads to predicted behavior that is completely at odds with the data. With $\phi = 0.05$, the higher option-value of a second birth leads to a sharp increase in first births. Another striking difference between the these simulations and the data applies to all scenarios. The increase in second births observed in figure 3 is gradual. In contrast, for both simulated values of ϕ the model predicts instantaneous jumps in the fertility rates followed by rapid convergence to a lower long run level.

The intuition behind the behavior predicted by the model (both the increase in first births and the jump in the birth rates) follows straightforwardly from the underlying assumptions.

6.1 Reformulating the model

We have estimated a modified version of the model that better mimics the observed behavior. We introduced two changes in our baseline specification:

1. Women believe that, with certain probability, the MC policy will not be available to them in the following period

2. Not all women are informed about the MC program benefits. Learning occurs following a process with constant hazard. There is no unlearning

The first modification is intended to break the tight link between first and consecutive births that follows necessarily from the assumption of rational expectations. The second modification is meant to match the gradual reaction to the introduction of the program. Both these assumptions are consistent with the poll evidence presented above.

Estimation of this modified version is currently under way. Preliminary result: $\phi = 0.0886$. Figure 7 shows the calendar time simulation.

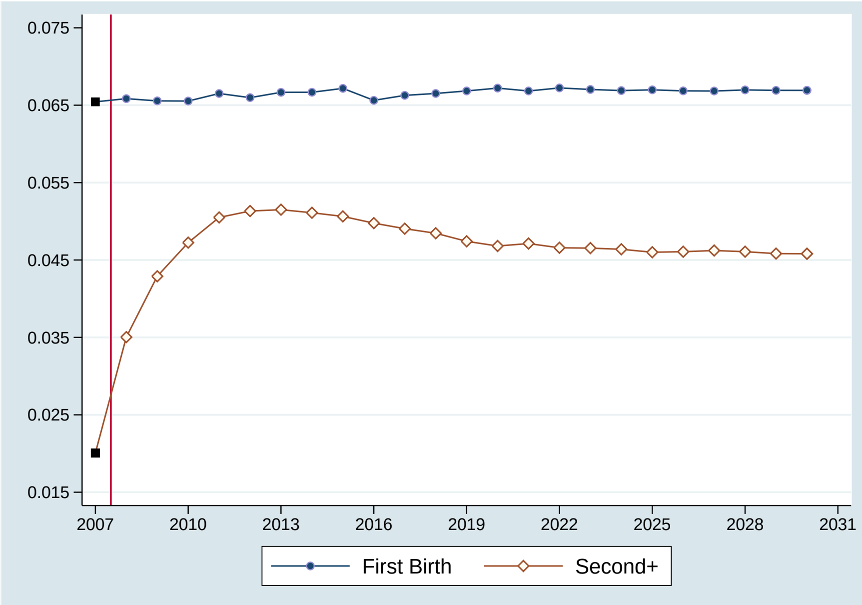

Figure 7 – The MC Program in Calendar Time. Alternative Specification.

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.

Bertrand, M.; Duflo, E., and Mullainathan, S. How Much Should We Trust Differences-in-Differences Estimates? *Quarterly Journal of Economics*, 119(1):249–275, 2004.

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.

McFadden, D. A Method Of Simulated Moments For Estimation Of Discrete Response Models Without Numerical Integration. *Econometrica*, 57(5): 995–1026, Sep 1989.

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

Smith, J. and Todd, P. Does Matching Overcome LaLonde's Critique of Nonexperimental Estimators? *Journal of Econometrics*, 125(1):305–353, 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: additional figures and tables

Figure A.1 – Birth Rate and TFR in the Long Run

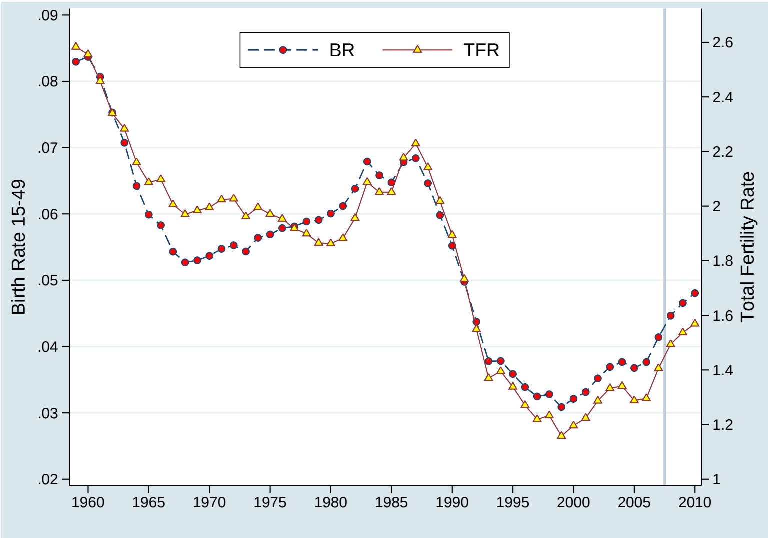

Notes: The data source is the Human Fertility Database. For the period 1994–2010, it coincides with the Rosstat-HFD series in figures 1 and 2.

Figure A.2 – Birth Rates By Birth Order for Women Ages 15-39

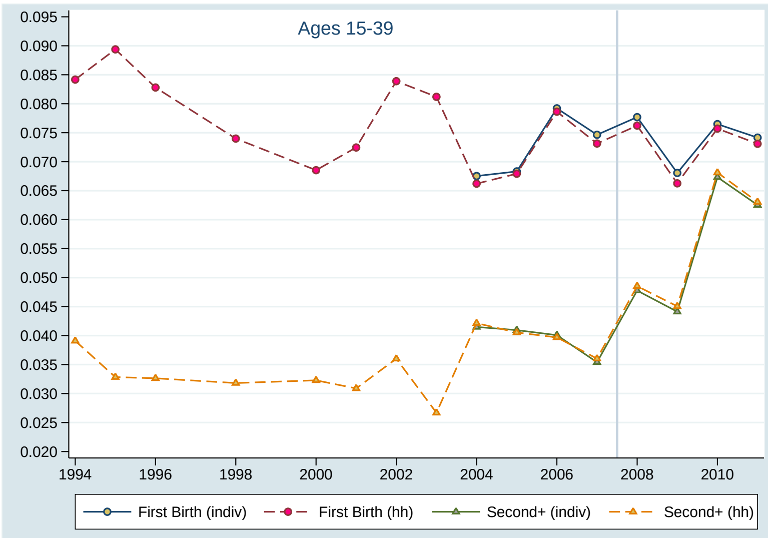

Notes: The data source is the RLMS. See notes to figure 3.

Table A.1 – Full Regression Results

	Difference-in-Differences	Before-After
post	-0.012	0.017*
	(0.008)	(0.009)
MCelig	-0.061***	
	(0.019)	
MCelig × *post*	0.023***	
	(0.005)	
Age Cubic Spline Coeff (6 knots)		
1	0.011***	0.000
	(0.001)	(0.003)
2	-0.099***	-0.004
	(0.009)	(0.019)
3	0.224***	-0.018
	(0.031)	(0.053)
4	-0.153***	0.058
	(0.045)	(0.060)
5	0.054	-0.030
	(0.039)	(0.043)
# of children (excluding new born)		
One	0.030	
	(0.019)	
Two	0.015	-0.016***
	(0.019)	(0.002)
Three	0.027	-0.010**
	(0.019)	(0.004)
Four or more	0.050**	0.006
	(0.022)	(0.012)
Education		
High School Completed	-0.015***	-0.019***
	(0.004)	(0.006)
Vocational Completed	-0.002	-0.016***
	(0.004)	(0.006)
Technical Completed	-0.008*	-0.019***
	(0.004)	(0.006)
University+ Completed	0.002	-0.014**
	(0.004)	(0.006)
Russian National	0.004*	-0.001
	(0.002)	(0.002)
Russian Born	-0.006*	-0.006*
	(0.003)	(0.003)
Urban Location	0.000	-0.004
	(0.004)	(0.004)
Married/Cohabiting	0.072***	0.022***
	(0.002)	(0.002)
One parent in HH	0.004*	0.003
	(0.003)	(0.003)
2+ parents in HH	-0.006	0.004
	(0.003)	(0.004)
One Grand-parent in HH	-0.006	0.003
	(0.005)	(0.009)
2+ Grandparents in HH	-0.019**	0.015
	(0.008)	(0.034)
One Sibling in HH	-0.022***	-0.008
	(0.003)	(0.005)
Two Siblings	-0.023***	0.004
	(0.006)	(0.017)
Three Siblings	-0.032***	0.011
	(0.009)	(0.032)
Four+ Siblings in HH	-0.023	0.050
	(0.016)	(0.074)
LM Experience	-0.000	0.000
	(0.000)	(0.000)
Age Group Youngest Child		
Less than 3	-0.075***	-0.025***
	(0.004)	(0.005)
Less than 7	-0.048***	-0.004
	(0.004)	(0.004)
Less than 12	-0.019***	0.009***
	(0.003)	(0.002)
Less than 18	-0.009***	0.004***
	(0.002)	(0.001)

continued in next page

continued from previous page

Child Care		
State Nursery	-0.012***	-0.014***
	(0.003)	(0.003)
Private Nursery	-0.007**	-0.004
	(0.004)	(0.004)
State Preschool	-0.010**	-0.006
	(0.005)	(0.005)
Private Preschool	0.007*	0.003
	(0.004)	(0.004)
Median Amount Paid for Childcare	-0.000	0.001
	(0.001)	(0.001)
Proportion with Free Care	0.012	0.007
	(0.020)	(0.022)
Observations	53,463	32,746
R-squared	0.076	0.036

Notes: Robust s.e. clustered at individual level. The cubic spline knots are 20, 25, . . . , 45. Education baseline is no degree. Year, location, and the constant are omitted.

Table A.2 – Evolution of Marital Status and Cohabitation with Parents

Origin		Destination		Age Group						
m_{t-1}	p_{t-1}	m_t	p_t	22–25	25–30	30–35	35–40	40–45	45–50	50–55
0	0	0	1	6.25	2.46	1.55	1.91	1.72	0.86	0.59
		1	0	15.91	23.16	14.95	8.78	6.41	3.93	2.84
		1	1	2.27	1.05	0.26	0.38	0	0.37	0
0	1	0	0	1.88	2.21	4.12	3.23	3.36	6.69	12.42
		1	0	5.52	5.41	2.26	2.08	1.83	1.34	0
		1	1	4.69	7.38	5.14	3.46	2.75	1.34	1.24
1	0	0	0	3.7	3.3	3.01	3.19	2.61	3.31	3.06
		0	1	1.92	0.86	0.24	0.39	0.11	0	0
		1	1	2.74	2.22	1.74	1.57	0.91	1.17	1.1
1	1	0	0	0.29	0.48	0.68	0.44	0.6	0.31	0
		0	1	3.76	6.5	2.88	3.08	2.71	2.19	2.15
		1	0	16.18	10.94	12.01	7.49	10.84	9.06	12.45

Table A.3 – OLS Estimates for Partner's Income and Other Household Income

	Eq. (7) $\log y_t^h$	Eq. (8) $\log y_t^o$
m_t		-0.6015***
		(0.0414)
p_t		1.1731***
		(0.0437)
t	0.0366***	-0.1071***
	(0.0059)	(0.0176)
t^2	-0.0006***	0.0019***
	(0.0001)	(0.0002)
S_1	0.1819***	-0.0926
	(0.0288)	(0.0876)
S_2	0.2367***	-0.1747**
	(0.0281)	(0.0862)
S_3	0.2746***	-0.2898***
	(0.0270)	(0.0833)
S_4	0.5568***	-0.1340
	(0.0272)	(0.0830)
G	0.4441***	0.0224
	(0.0140)	(0.0438)
Constant	10.9080***	12.1415***
	(0.1107)	(0.3271)
Observations	17,727	25,269
R-squared	0.129	0.056

Table A.4 – Labor Income and Employment Probability (logit)

	Eq. (10) $\log y_t^f$	Eq. (9) l_t
l_{t-1}		3.7192***
		(0.0387)
X_{t-1}	0.0223***	
	(0.0021)	
X_{t-1}^2	-0.0004***	
	(0.0001)	
S_1	0.0656**	
	(0.0284)	
S_2	0.0124	
	(0.0272)	
S_3	0.1076***	
	(0.0263)	
S_4	0.4911***	
	(0.0261)	
G	0.3775***	0.3715***
	(0.0124)	(0.0384)
Constant	10.9239***	-1.8919***
	(0.0285)	(0.0419)
Observations	22,320	31,911
R-squared	0.124	