CHILDREN OF A (POLICY) REVOLUTION: THE INTRODUCTION OF UNIVERSAL CHILD CARE AND ITS EFFECT ON FERTILITY

Stefan Bauernschuster
University of Passau, CESifo

Timo Hener
Ifo Institute

Helmut Rainer
University of Munich, Ifo Institute

Abstract
What role does affordable and widely available public child care play for fertility? We exploit a major German reform generating large temporal and spatial variation in child care coverage for children under the age of three. Our precise and robust estimates on birth register data reveal that increases in public child care have significant positive effects on fertility. The fertility effects are more pronounced at the intensive than at the extensive margin, and are not driven by changes in the timing of births or selective migration. Our findings inform policy makers concerned about low fertility by suggesting that universal early child care holds the promise of being an effective means of increasing birth rates. (JEL: J13)

The editor in charge of this paper was Fabrizio Zilibotti.

Acknowledgments: We would like to thank the editor, Fabrizio Zilibotti, and four anonymous referees for constructive comments and helpful suggestions that improved the paper. Seminar participants at the Ifo Institute, Statistics Norway, the University of Constance, the University of Freiburg, the University of Linz, the University of Munich, the University of Passau provided provided useful suggestions. We also thank participants at the 2013 North American Summer Meeting of the Econometric Society, the 2013 conference of the European Association of Labour Economists, the 2013 conference of the European Economic Association, the 2013 conference of the German Economic Association, the 2013 meeting of the Royal Economic Society, the 2014 E.ON Stipendienfonds and Research Council of Norway workshop on Family Policy, the 2014 CESifo Area conference on Employment and Social Protection, and the 2014 conference of the Work, Pensions and Labour Economics Study Group (WPEG).

E-mail: Stefan.Bauernschuster@uni-passau.de (Bauernschuster); hener@ifo.de (Hener); rainer@ifo.de (Rainer)

1. Introduction

Fertility levels in many developed countries are no longer sufficient to assure the long-term replacement of the population, with predictable economic consequences such as financial difficulties in health care and pensions systems. This raises the question of whether policy makers as agents of the public interest should care about fertility outcomes. Most traditional theories of fertility (see e.g. Leibenstein, 1957; Becker, 1960) assume that the benefits of having an extra child accrue entirely to the parents, and therefore imply that public intervention in individuals' fertility choices cannot be justified other than for equitative purposes. But suppose that the benefit of an extra child does not entirely go to the parents—that children are to some extent a public good and consequently involve positive social externalities. For example, it is well understood that social security schemes in which the active generation pays the pensions of the retired generation socialize part of the benefits of a child, mainly through a growing tax base (see e.g. Cigno, 1993; Folbre, 1994). What we are confronting then is a generalized prisoner's dilemma: fertility choices at the individual level are only based on the direct utility that a couple gets from its offspring, neglecting the fact that progeny benefits all in society. Since this disjunction between private interest and public good implies an insufficient number of children, it has been used by economists as well as demographers to advocate a reexamination of existing public policies and their appropriate redesign in a pronatalist direction (see e.g. Demeny, 1986; Sinn, 2004).

On a theoretical level, there is considerable work exploring approaches to public policy reform that could internalize the positive social externalities associated with offspring (see e.g. Groezen et al., 2003; Fenge and Meier, 2005). On a practical level, some governments have recently responded to concerns about low fertility by moving demographic considerations to the top of their political agenda and implementing reforms intended to induce people to have more children (see e.g. Rindfuss et al., 2010; Takayama and Werding, 2011). Chief among these have been efforts to expand public child care. However, there is relatively little empirical research on the fertility effects of such policies. Moreover, existing studies are plagued by identification problems due to the limited magnitude of the available policy variation. As a result, it is still open for debate whether public child care provision is an effective way to increase fertility rates.

The aim of this paper is to provide new empirical evidence on the relevance of public child care for fertility. To overcome problems of endogeneity, we draw upon a major German reform from the mid-2000s, which led to a large scale staggered expansion of public child care for children under the age of three. Germany has long been known for its low fertility, and one of the explicitly stated goals of the reform was to induce couples to have more children by making them less costly in terms of income and career opportunities. In essence, the reform included a commitment by the federal government to move from having almost no child care slots available for children under the age of three to having slots available for all children in this age group. However, although the federal government initiated the reform, local authorities were responsible for the expansion of public child care. This immediately generated large variation in child care coverage, both across time and between West Germany's 325 counties, which we exploit using difference-in-differences techniques. Our analysis draws upon birth registration records, which cover all births in West Germany—on average 580,000 annually. We match the information from the birth registers with administrative data on child care coverage at the county-level. The data allows us to examine the effects of public child care both at the extensive margin of fertility (i.e. entry into parenthood) and at the intensive margin (i.e., the number of parents' offspring). Age-specific birth rates allow us to examine possible effects on the timing of births within cohorts. In addition, we are able to ask whether public child care expansion has an effect on babies' health outcomes at birth, which might be expected if such an intervention leads to a change in the composition of parents (Dehejia and Lleras-Muney, 2004).

We present six classes of results. *First*, we find consistent and robust evidence of a substantial positive effect of public child care expansion on fertility. To be concrete, our estimates suggest that a 10 percentage point increase in public child care coverage increases the number of births per 1,000 women by 1.2, or roughly 2.8% of the baseline birth rate. Under the strict assumption of linearity, this result implies that an increase in public child care coverage by 30 percentage points—as ultimately achieved by the reform under consideration—leads the average woman to have roughly 0.12 more children. Given that the total fertility rate in Germany has been hovering around 1.4 for decades, this effect appears to be quantitatively important. *Second*, we provide evidence that the increase in fertility brought about by public child care expansion is not driven by births brought forward in time, which suggests a positive effect on completed fertility. *Third*, we find that the effects of the public child care expansion on fertility are stronger at the intensive than at the extensive margin: a 10 percentage point increase in child care coverage increases the incidence of second and third births by 4% and 7%, respectively. *Fourth*, there is no evidence that children born in response to increases in public child care have inferior health outcomes at birth such as a lower birth weight or a lower birth length. *Fifth*, we show that the positive fertility effects are accompanied by increases in female employment. *Sixth*, a simple cost-benefit analysis suggests that the fertility effect of a given amount of public spending on child care exceeds the effect of increasing spending on child benefits by a factor of five. A battery of robustness checks, which amongst others deal with the common trend assumption, time-varying regional heterogeneity, selective migration or the timing of fertility responses, corroborate our results.

Taken together, our findings contribute to ongoing academic and public debates on family policies and low fertility in developed countries. In particular, our analysis provides evidence suggesting that policies that facilitate the combination of parenthood and employment hold the promise of being an effective way to positively influence birth rates where these rates are considered to be too low.

The remainder of the paper is organized as follows. Section 2 provides a review of the related literature and discusses how our study contributes to it. Section 3 provides the institutional setting. Section 4 outlines the empirical strategy. Section 5 describes the data. In Section 6, we present the results, a battery of robustness checks, and a simple cost-benefit analysis. Section 7 concludes.

2. Related Literature

Ever since the seminal works of Leibenstein (1957) and Becker (1960), economists have taken an interest in how government policies influence fertility. From a theoretical perspective, the impact of child care availability on fertility is as follows. At the extensive margin, better access to child care lowers the opportunity cost of having a first child, and so encourages women to enter into motherhood. Consider next the fertility responses at the intensive margin. Increases in child care availability are likely to allow mothers to return to work sooner after the birth of a first child or to choose more high powered careers. On one side, this generates a positive income effect which works to increase the likelihood of a second or higher-order birth. However, at the same time, it raises the opportunity cost of having an additional child, generating a substitution effect which works in the opposite direction. A low price of public child care makes it more likely that the positive income effect dominates the negative substitution effect (Ermisch, 2003; Apps and Rees, 2004). Since public child care is heavily subsidized in Germany, we would expect that increases in child care availability also encourage second and higher-order births.

From an empirical perspective, there are quite a few impressive studies on the impact of *financial incentives* on fertility. Milligan (2005) provides evidence that the introduction of a pronatalist cash transfer policy in the Canadian province of Quebec had a positive effect on fertility, especially among women with high family income. Cohen et al. (2013), using Israeli data and variation in Israel's child subsidy, finds strong positive effects on fertility among women in the lower range of the income distribution. Raute (2014) exploits changes in financial incentives arising from a reform in parental leave benefits in Germany and finds strong effects on fertility, driven mainly by highly educated women.[1]

Much less can be said for our knowledge of the impact of *child care provision* on fertility.[2] The closest antecedent to this study is an interesting paper by Rindfuss et al. (2010), which documents a positive link between child care availability and fertility in Norway. Our analysis offers several major innovations on this study. First and most importantly, we implement a quasi-experimental strategy while Rindfuss et al. (2010) use the results from a discrete-time hazard model to simulate the effect of different child care availability scenarios on fertility patterns. Second, and related

1. Milligan (2005) provides a comprehensive survey of the earlier literature on financial incentives and fertility until the early 2000s.

2. Most previous studies which analyze child care provision have considered as major outcomes female labor supply (see, e.g., Gelbach, 2002; Baker et al., 2008; Lefebvre and Merrigan, 2008; Cascio, 2009; Havnes and Mogstad, 2011a) and/or child development (see, e.g., NICHD – Early Child Care Research Network, 2003, 2004; Baker et al., 2008; Havnes and Mogstad, 2011b).

to the first point, we exploit a policy reform which led to a large scale expansion of child care over a short time horizon. In contrast, Rindfuss et al. (2010) exploit the growth in child care slots in Norway from 1973 to 1998, which leaves the identification vulnerable to endogeneity bias. In particular, given the 25-year study period, it is difficult to separate supply-side shocks to child care availability from fertility-driven spikes in the demand for child care. Third, Norway is characterized by several macro-level factors that make it in many ways a unique country in Europe: its egalitarian gender ideology, its social democratic political economy, and its oil wealth. It is important to understand whether the positive link between child care availability and fertility in Norway is mainly due to the presence of these differentiating macro-level factors. By focusing on Germany, we are able to assess wether a country with a comparatively sharp gender differentiation and a traditional focus on the male breadwinner model can expect a substantial fertility increase by adopting Norwegian-style child-care policies. There is also an earlier literature that examines the effect on fertility of child care provision (Del Boca, 2002; Hank and Kreyenfeld, 2003; Hank et al., 2004; Del Boca et al., 2009). Overall, these studies cover periods with limited policy variation and yield inconclusive results.

Our study also makes contact with a paper by Mörk et al. (2013), which exploits the exogenous variation in *child care costs* caused by a Swedish child care fee reform. The results suggest that the reduction in child care costs increased the number of first births, but only seemed to affect the timing of second births. With our focus on the introduction of universal, highly-subsidized child care in Germany, the treatment we exploit differs quite markedly from that in Mörk et al. (2013). In addition, the Swedish child care fee reform took place in a context in which child care enrolment was already almost universal and the labor force participation of mothers very high. In contrast, the German reform we exploit took place at a time when child care for young children was virtually non-existent and the labor force participation of mothers relatively low. Thus, the margins of adjustment are likely to be very different. Finally, Björklund (2006) shows that the economic incentives created by Sweden's family policy package from the 1960s to around 1980 had a strong impact on fertility. However, since the policy package included a mix of financial and in-kind support for families, the study does not shed light on the effect of child care availability on fertility.

3. Background and Context

Ever since the 1970s, Germany has been among the twenty countries with the lowest fertility rates worldwide (Population Reference Bureau, 2007). Historically, fertility rates in Germany were increasing during the 1950s and early 1960s from just above 2.0 to 2.5, but they then dramatically decreased in the late 1960s and early 1970s to a level of 1.5 in 1974. During the last four decades, fertility stayed constant at a very low level of roughly 1.4.[3] Germany's population reached a maximum shortly after the turn of the millennium and has started to decline thereafter as a result of sustained very low fertility rates (Dorbritz, 2008). As a consequence, the German government now approaches demographic issues in an official way and wants to encourage higher fertility through policy interventions, which it refrained from doing after World War II because anything resembling pronatalism was discredited for historical reasons (Takayama and Werding, 2011).

A key initiative in this regard was the introduction of universal public care for *children under the age of three*. In 1996, the German government had already enacted legislation that granted *children aged three to six* the right to a place in a public kindergarten. By the early 2000s, this reform had led to full provision of half-day public child care for children in that age group in West Germany (Bauernschuster and Schlotter, 2015). However, up until the mid-2000s, public child care for younger children—i.e., those under the age of three—was virtually non-existent or at least severely rationed in West Germany. For example, in a survey conducted in 2005, 35 percent of West German mothers with under three year olds stated a demand for a child care slot (Bien et al., 2006), while only roughly 5 child care slots per 100 children in this age group were available.[4,5] At the same time, virtually no private market for child care had emerged.[6] Prompted by the severe rationing of public child care, the German government implemented a set of public child care reforms during the period 2005-2008, with the explicit intention to increase fertility levels:

- At the beginning of 2005, a federal law (*"Tagesbetreuungsausbaugesetz"*) became effective which included the commitment to create 230,000 additional child care

3. Data from the World Bank (2009) depict a fertility rate of 1.38 for Germany in 2008. Thus, Germany lies well below the EU-27 average of 1.60 and close to Poland (1.39), Portugal (1.37), Hungary (1.35), or Japan (1.34). Fertility rates in the US (2.10), France (2.00), Norway (1.96), or Sweden (1.91) are substantially higher.

4. Wrohlich (2008) estimates that more than 50 percent of West German mothers with children aged 0-3 were queuing for a child care slot in the mid-2000s, suggesting that the excess demand for child care was even more severe.

5. The situation was completely different for early child care in East Germany. Throughout the history of the former German Democratic Republic, the East German government strongly supported the use of public daycare for children of all ages. The East German child care system survived the German reunification, with more than one-half of all East German children under the age of three and almost all East German children between three and six attending a child care center in the mid-1990s. At the turn of the millennium, parents in East Germany demanded fewer child care slots for children than were available (Hank et al., 2001).

6. One reason for the lack of a private market are the strict regulations (set at the state-level) faced by child care providers. As pointed out by Felfe and Lalive (2012), these regulations concern dimensions such as opening hours, group sizes, staff-child rations, but also qualifications of the staff before being allowed to work in the sector.

slots for under three year old children by 2010 in West Germany. The specific aim was a child care coverage rate of 17% by 2010 in West Germany.[7]

- In 2007, a summit (called *"Krippengipfel"*) of the three federal levels—i.e., federal state, "Länder", local authorities—agreed upon increasing the child care coverage rate for under three year olds to 35% by 2013.
- At the end of 2008, the federal law to promote children (*"Kinderförderungsgesetz"*) established the legal claim to a child care slot for all preschool children aged one and above by 2013.

In the run-up to the reforms, the three federal levels agreed that each level bears a share of the expansion costs.[8] The child care slots that were then created were heavily subsidized. In 2006, the total operating costs of child care for under three year olds amounted to €14.1 billion, with roughly 79% of these covered by public subsidies, 14% by parents and 7% by private organizations. Parental fees are regressive in family size and progressive in family income and range from 0 to 600 Euros per month. Thus, the initiative led to low-cost care for young children across all German counties.

Since public child care for under three year olds was virtually non-existent in West Germany before 2005, all counties in each state had to substantially expand public child care in order to be able to fulfill legal claims to a child care slot for all preschool children aged one and above by 2013.[9] In order to assess the expansion brought about by these initiatives, Figure 1 provides two maps which illustrate the child care coverage rate for West Germany's 325 counties in 2002 and 2009, respectively. In 2002, we observe that the child care coverage rate was consistently below 5% across virtually all West German counties. In succeeding years, the child care coverage rate more than quintupled to reach an average of 15% in 2009, with roughly a third of all slots being full-time slots (Destatis, 2010). However, it is also evident from the map that the counties differ distinctly in the magnitude of public child care expansion. In 2009, the public child care coverage rates vary from 3.7% to 35.9%. The percentage point increases in child care coverage from 2002 to 2009 range from a minimum of 3 percentage points to a maximum of 27 percentage points. Also note that there is considerable variation in this expansion across counties even within the same state.

7. A draft of the law was forwarded to the federal assembly ("*Bundesrat*") on August 13, 2004, followed by the publication of the official government draft on September 6, 2004. The law was passed on December 18, 2004, and became effective on January 1, 2005.

8. The reform did not crowd out public spending on other major family-related benefits. In particular, the funds allocated to child benefits, child allowances and paid maternity leave were unaffected by the reform.

9. For regional authorities, the only binding aspect of the federal laws was the introduction of the legal claim to a child care slot. The relevant deadline for this law was August 1, 2013. The federal government did not penalize regional authorities for violating the legal claim set out in the law. However, parents could sue municipalities if they were unable to offer child care slots for their eligible children. Court rulings to date have established that regional authorities have to (i) cover the costs if parents have to resort to private child care arrangements due to a lack of public child care slots, and (ii) pay foregone earnings in cases where the lack of public child care slots acts as an impediment to parental employment.

FIGURE 1. Public child care coverage in West German counties in 2002 and 2009

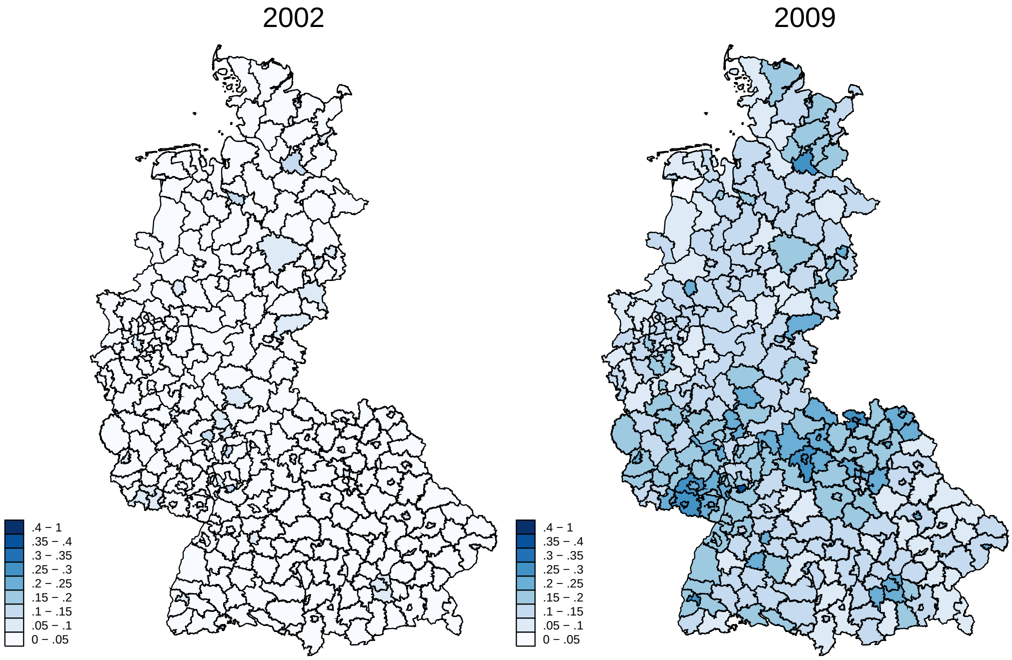

Notes: The left panel shows child care coverage in West German counties in 2002, the right panel shows child care coverage in West German counties in 2009. East German counties are shaded in gray.

Indeed, two thirds of the variation in child care coverage are attributable to variation within states, while one third is attributable to differences between states.

Where does this variation come from? The process of opening up new child care slots involved many complex and intertwined decisions of municipality level, county level and state level authorities, respectively. On the one hand, authorities at the municipality and county level had the responsibility to assess the local demand for child care, with demographic and economic factors such as current cohort sizes and labor market conditions entering the projections. On the other hand, authorities at the state level had to approve proposals to set up new child care centers which were submitted by non-profit organizations. This administrative process was prone to problems that varied substantially across counties and that could not be influenced by local authorities (see e.g. Huesken, 2011). Amongst them are varying routines and knowledge about the complicated funding system (with subsidies coming from the federal state, the state and the municipality), shortages in construction ground, various regulations for building child care centers, shortages in qualified child care workers, serious delays in approval or final rejections of applications due to non-compliance with regulations. As result, the growth of child care slots differed at the county level not only due to some well defined predictors of local child care demand, but also due to shocks to the local supply of new child care slots emanating from lengthy and intricate administrative processes and rules (see e.g. Felfe and Lalive, 2012). The latter component is arguably orthogonal to expected changes in fertility rates, and provides the basis for our identification strategy.

4. Empirical Strategy

To identify the effect of public child care on fertility, we exploit spatial and temporal variation in child care coverage using difference-in-differences (DID) techniques. We start with a basic DID model that employs a dichotomous treatment group variable, and then generalize this approach using a continuous treatment variable. The basic DID model is used to provide some first intuitive graphical insights. The more detailed results are based on the generalized model.

For the basic DID model, we need to generate a dichotomous treatment group variable. To this end, we order all West German counties by the absolute size of the increase in public child care coverage from 2002 to 2009.[10] We define counties whose increase in child care coverage was above the median as the treatment group. Accordingly, counties whose increase was below the median constitute the control group. By choosing this approach, we follow Havnes and Mogstad (2011a) who use a similar identification strategy for analyzing the effects of universal child care on maternal employment in Norway. It is important to note that this approach defines treatment and control groups (or more and less intensely treated groups) based on variation in the outcome of a set of policy initiatives, i.e., how much child care expanded in a county over the sample period. This is not the same as an approach that exploits variation coming from differences in formal rules or policies across counties. Thus, the validity of our empirical design depends on the variation in child care growth being orthogonal to expected changes in fertility rates. In what follows, we will investigate wether this key identifying assumption holds.

We estimate the basic DID model for the period from 1998 to 2010. Thus, we make use of data from several pre- and post-reform years. In its basic form, the model can be written in the following way:

$$y_{ct+1} = \beta_c + \gamma_t + \sum_{t=1998}^{2003} \delta_t (D_c \times \gamma_t) + \sum_{t=2005}^{2009} \delta_t (D_c \times \gamma_t) + \varepsilon_{ct+1} \tag{1}$$

y_{ct+1} is the number of births by 1,000 women aged 15 to 44 living in county c in year $t+1$. The outcome variable is measured in $t+1$ because there are at least 9

10. We choose the year 2002 as the baseline for two reasons: First, since public child care for toddlers was not a major political issue until the year 2005, the year 2002 is certainly a year which is unaffected by any political decisions aimed at expanding public child care. Second, the year 2002 is the last year where administrative data on child care coverage at the county level is available before it became a political issue in 2005. Data exists for the years 1998, 2002 and 2006-2009.

months from the decision to have a child to the actual birth. We will later allow the outcome variable to react to the treatment within an (empirically validated) time frame of up to 22 months. β_c are county fixed effects and thus capture time-invariant regional heterogeneity. γ_t are year fixed effects capturing year-specific differences in birth rates that are common to the treatment and control group. D_c is the treatment group indicator for county c, which is unity for counties with above median increase in public child care coverage, and zero for counties with below median increase. The coefficients δ on the interactions of the treatment group indicator and the year fixed effects identify deviations from a common trend. If the child care expansion had positive effects on birth rates, we should observe δ being zero in the pre-treatment years and increasing in size in the years after 2005. Standard errors ε are clustered at the county level. This basic DID model identifies intention-to-treat effects (ITT). In order to make statements about the average treatment effects on the treated (ATT), we have to rescale the reduced form estimates by the emerging difference in child care coverage between treatment and control group counties (see, e.g., Baker et al., 2008; Havnes and Mogstad, 2011b).

A straightforward alternative to this rescaling procedure is to estimate a more generalized DID model that uses the local child care coverage rate as a continuous treatment variable. In this generalized model, we are able to exploit the full variation in local child care coverage. Furthermore, we are able to avoid questions regarding the exact definition of treatment and control group. The regression equation can be written as follows:

$$y_{ct+1} = \eta_c + \mu_t + \rho d_{ct} + X'_{ct}\lambda + \zeta_{ct+1} \tag{2}$$

where η_c is a county fixed effect for county c and thus captures time-invariant heterogeneity between counties, μ_t is a fixed effect for year t, and X'_{ct} is a vector of covariates of county c that vary over time t. The key variable of interest, d_{ct}, is now continuous and represents the child care coverage rate of county c in year t. Accordingly, ρ captures the treatment effect of the child care coverage expansion on fertility. As before, y_{ct+1} is the outcome variable measuring the number of births per 1,000 women aged 15 to 44 living in county c in period $t+1$. Standard errors ζ are allowed to correlate within counties over time. Note that by using this specification, we restrict the marginal effects of expansions in public child care to be constant. Furthermore, in contrast to the basic DID specification, we only use years t for which we actually observe public child care coverage in the data.

The key identifying assumption of any DID model is the common trend assumption. In our case this means that, conditional on county fixed effects and the set of time-varying covariates, there are no unobserved characteristics of a county that vary over time and are correlated with public child care expansion and future changes in fertility. Note that the expansion of public child care need not be orthogonal to county characteristics since we control for county fixed effects. It is nevertheless informative to investigate differences in pre-reform characteristics between counties that expanded quickly and counties that expanded slowly. As we will show below, the treatment and control group counties in the basic DID specification are very similar in their pre-reform characteristics.

Despite this, it might be the case that time-varying factors which are correlated with fertility evolve differently in strong expansion counties as compared to weak expansion counties and thus bias the estimates. To address this concern, the basic DID model investigates whether treatment and control group counties follow the same fertility trend in the pre-treatment period. This test is commonly considered a useful check for the validity of the common trend assumption. However, even if the trends are very similar prior to the treatment, this does not safeguard us against the possibility that they deviate from each other in the post-reform years for reasons other than the expansion of public child care. Thus, to further investigate the robustness of our results with respect to time-varying county characteristics, we will run the regressions both without any covariates and with a rich set of county-specific time-varying covariates.

Let us briefly outline the basic intuition behind our choice of time-varying covariates. One potential confounding factor could be changes in predicted fertility at the local level. As we have argued above, virtually all West German counties had to massively expand public child care in response to the federal child care initiatives. However, it is conceivable that the pace of expansion is affected by changes in predicted fertility within a county over time. In other words, although predicted child care demand exceeds current supply in virtually all counties, the political pressure to quickly increase child care supply might be higher in counties where predicted increase in fertility (or closely related, future increase in child care demand) is relatively high, which would bias the estimates upwards. To minimize this problem, we control for time-varying local socio-demographic factors which are typically used to predict fertility, and therefore also might be relevant for local authorities, in all regressions. In particular, we include extremely detailed information on a county's population age structure to capture local demographics. These are year of age specific population shares of females aged 15 to 44 and of the whole population aged 45 and above. Furthermore, we control for population density to account for regional agglomeration tendencies, and the male employment rate to capture local labor market conditions. By controlling for GDP per capita, we capture a county's prosperity and mitigate effects of the financial crisis that may affect counties differentially. Moreover, we directly control for local political attitudes by including the conservative vote share as an additional covariate.[11] We also control for the share of highly educated women up until age 44 to account for the labor market attachment of women in fertile age. In extended regressions, we also include municipalities' gross revenue and debt to capture time-varying differences in local public finance, which may be important in

11. This variable might also be a good proxy for local cultural attitudes about women and the family. This is because all major parties other than the conservative party are associated with more liberal family policies. We also experimented with including vote shares for all parties separately and found the results to be unaffected.

the decisions to expand child care. Furthermore, we include the number of newly built dwellings which may be geared towards attracting families.

Controlling for these county characteristics makes sure that we only exploit the component of child care growth that can neither be explained by time-invariant differences between counties nor by a rich set of potential predictors of local child care demand and fertility. We argue that the remaining variation in the pace of the mid- to late-2000s child care expansion was due to shocks to the local supply of new child care slots emanating from lengthy and intricate administrative processes and rules. Accordingly, conditional child care expansion should be exogenous to future changes in fertility.

5. Data

We use administrative data from the Statistical Offices of the German Länder (*Statistische Landesämter*) on public child care for children under the age of three. These data are available for the years 1998, 2002, 2006, 2007, 2008, and 2009. The number of public child care slots are reported to the authorities in the first half of March in every year.[12] Combining these data with detailed administrative data on the counties' population structure, we build the key variable of interest, public child care coverage, defined as public child care slots (measured in March of year t) over the population of children less than three years old (measured in December of year t-1), which we also simply refer to as *child care* in the following. Table 1 shows that child care coverage averages 7.7% over the whole period of observation and varies widely from 0 to 35.9%. The average coverage rate across West Germany's 325 counties was very low in 1998 (1.6%) and 2002 (2.2%). The modest increase from 1998 to 2002 is mainly due to a decrease in births rather than an increase in child care slots (DJI, 2005). Yet, there is already some variation across counties with some reporting no child care for under three year olds at all and others reporting coverage rates up to 13.0%. After 2005, the reform takes effect and the rise in coverage rates accelerates. In 2006, the rates reach 7.3% on average. The minimum value is lifted above zero and the maximum value up to 23.3%. Until 2009, the average coverage rate is doubled to a value of 14.2%. Note that while the whole distribution of child care coverage shifted to the right, we do not observe a convergence process between counties. Instead, the standard deviation of coverage rates steadily increases from 1998 to 2009. Closer inspection of the data reveals that patterns of child care expansion are very heterogenous across counties. Some counties expand very slowly and others very fast. Some counties gradually increase child care over time, some start off strong but come to a halt, and again others are delayed by a couple of years and later increase

12. To be precise, the number of slots are reported from 1998 to 2002 whereas from 2006 to 2009 the number of children attending child care are reported. However, due to extremely severe rationing of public child care for under three year olds, the number of children attending child care should be very close to the number of available slots.

Table 1. Child care coverage over time.

		Child care coverage				
Year	N	Mean	Median	S.D.	Min	Max
1998	325	0.016	0.009	0.020	0.000	0.118
2002	325	0.022	0.014	0.023	0.000	0.130
2006	325	0.073	0.068	0.038	0.010	0.233
2007	325	0.094	0.085	0.044	0.022	0.289
2008	325	0.117	0.109	0.048	0.033	0.352
2009	325	0.142	0.135	0.050	0.037	0.359
Total	1950	0.077	0.071	0.061	0.000	0.359

Note: The figures show mean child care coverage rates across West German counties as well as standard deviations, median, minimum, and maximum values. All information is provided for the years 1998, 2002, 2006, 2007, 2008, and 2009.

coverage steeply. Overall, we find many different types of expansion patterns (see Online Appendix Figures A1-A2).

Register data based on the universe of birth certificates of all 325 West German counties[13], covering roughly 580,000 births per year, are the basis of our fertility measures. We collapse the individual birth data on county-year cells for the period from 1998 to 2010. At the county level, the data is combined with the administrative data on the population structure to compute fertility outcome variables. The main dependent variable *births per 1,000 women* is calculated as the sum of births over 1,000 women in reproductive age, i.e., between 15 and 44 years. We will refer to this outcome variable as the *birth rate*. Additionally, the data also allow us to compute disaggregated age-specific birth rates, i.e., the number of births per 1,000 women of a specific age over 1,000 women of this specific age. The denominators of these fertility measures ensure that the results are not confounded by changes in the size of the relevant female population. For births to married mothers, we know whether the birth is the first, second, third, fourth or higher-order birth. This allows us to analyze the effects of child care both at the extensive and the intensive margin. Moreover, we draw on information on children's birth weight and birth length, which is available for all births in the data set.

We measure births in year t, i.e., in the same year as the independent variables, and in year $t+1$. The reason is that the main variable of interest, child care, is measured in the first half of March each year, while births (as the sum of births in a year) are measured on December 31. By allowing fertility to respond in t and in $t+1$, we allow for a conception and gestation lag of a maximum of 22 months. The suitability of this timing specification will later be validated empirically. We observe 44.2 births per 1,000 women of fertile age on average. There is a pronounced inverted u-shape relation between birth rates and age peaking at around 30. Since age-specific fertility rates differ substantially, changes in the composition within the population

13. Data is provided by the Statistical Offices of the German Länder (*Statistische Landesämter*).

of 15 to 44 year old women would affect the main outcome variable. Therefore, we include the share of women aged 15, 16, 17, 18, ..., 44 in all women aged 15-44 within a county in our regressions. This ensures that our results are not biased by compositional changes.

Additional data from the Statistical Offices of the German Länder (*Statistische Landesämter*) and the Federal Employment Agency (*Bundesagentur für Arbeit*) complement the county level panel data set. In particular, we use information on population density, GDP per capita, the male employment rate[14], year of age specific (45, 46,...,75+) population shares, vote shares of political parties[15], municipalities' debt and revenues as well as the number of newly built dwellings. In addition, we use data from the German Microcensus to compute the share of highly educated females in all females aged up to 44 years excluding pupils. Highly educated women are women with a tertiary education (ISCED codes 5 and 6). Finally, in robustness checks dealing with potential selective migration, we use female gross migration flows as a share of a county's total population. Descriptive statistics of all variables used in the analysis are given in Online Appendix Tables A1-A4. The Appendix also provides a description of the data merging process as well as background information on the child care data.

6. Results

6.1. Basic difference-in-differences results

We start by comparing pre-treatment characteristics of West German counties with an above-median increase in child care (treatment group) and a below-median increase in child care (control group), respectively. Table 2 depicts the means as well as the results of t-tests for differences in means for the two groups in the pre-treatment year 2002. The statistics show that the difference in child care coverage is very small (0.5 percentage points), whereas the birth rate is significantly lower in the treatment group. This means that it was counties with low fertility that more strongly expanded child care in response to the federal reform. Interestingly, treatment and control group are statistically indistinguishable along a number of fundamental determinants of local birth rates and public child care. These include county measures such as the population density, GDP per capita, the male employment rate, the conservative vote share, municipalities' gross revenue and debts, and the number of new dwellings. The only noticeable difference is the positive gap between treatment and control group

14. To be precise, the variable measures the number of employed males subject to social security contributions as a share of all males aged 15 to 64.

15. We use the vote share of the conservative party in the last general elections and interpolate the years between elections by county.

Table 2. Pre-treatment descriptives for treatment and control group

Variable	Mean		Mean-Diff.	T-test	
	Control	Treatment	(T-C)	t-stat	p-value
Child care					
Child care coverage	0.019	0.024	0.005	1.839	0.067
Dependent variables					
Birth rate (t)	45.740	43.614	-2.126	-5.602	0.000
Birth rate (t+1)	44.843	42.642	-2.201	-6.059	0.000
Control variables					
Population density	590.764	541.310	-49.454	-0.647	0.518
Employment rate (m)	0.604	0.599	-0.005	-0.880	0.380
GDP per capita (in 1,000)	25.465	26.867	1.402	1.270	0.205
Conservative vote share	0.438	0.448	0.010	0.700	0.484
Female high education share	0.130	0.165	0.035	5.862	0.000
Gov revenue	357.555	359.393	1.838	0.040	0.968
Gov debt	0.668	0.635	-0.033	-0.602	0.548
New dwellings	0.219	0.215	-0.005	-0.149	0.881
Share of women 15-19	0.138	0.132	-0.006	-3.756	0.000
Share of women 20-24	0.139	0.140	0.001	0.701	0.484
Share of women 25-29	0.136	0.136	0.001	0.403	0.687
Share of women 30-34	0.179	0.179	-0.000	-0.102	0.919
Share of women 35-39	0.210	0.213	0.003	2.484	0.013
Share of women 40-44	0.199	0.200	0.001	1.008	0.314
Population fraction 45-49	0.070	0.071	0.001	2.811	0.005
Population fraction 50-54	0.064	0.065	0.001	2.104	0.036
Population fraction 55-60	0.053	0.052	-0.000	-0.167	0.867
Population fraction 60-64	0.067	0.066	-0.000	-0.557	0.578
Population fraction 65-69	0.056	0.054	-0.001	-1.937	0.054
Population fraction 70-74	0.044	0.043	-0.001	-2.237	0.026
Population fraction 75+	0.073	0.069	-0.003	-2.388	0.018
Detailed population structure[a]					

Note: The table shows means, differences in means and differences-in-means-tests for the control and the treatment group in 2002, the pre-treatment period. The last two columns depict results of T-tests for equality in means for each variable as t-statistics and p-values. Birth rates are defined as births per 1,000 women aged 15-44 years. Debt and income of municipalities are not reported for the federal city states Hamburg and Bremen (including Bremerhaven). Income information is missing in 2001 from all 15 Schleswig-Holstein counties. Income information in 2009 is not included due to fragmentary raw data. Revenue and debt figures are divided by 1,000,000 EUR and the number of new dwellings is divided by 1,000.

a. Tables of descriptive statistics for share of females and the population by years of age as used as control variables can be found in the Appendix.

in the share of highly educated females, which amounts to 3.5 percentage points.[16] Coming to pre-treatment differences in the age structure of the population, we have aggregated the age specific values to age groups to provide a short overview in Table 2.

16. Differences in female education are not able to explain a substantial share of the differences in child care expansion between counties, though. The difference in child care expansion between treatment and control group decreases from 6.8 to 6.5 percentage points once we condition on differences in female education.

FIGURE 2. Child care and fertility in treatment and control group

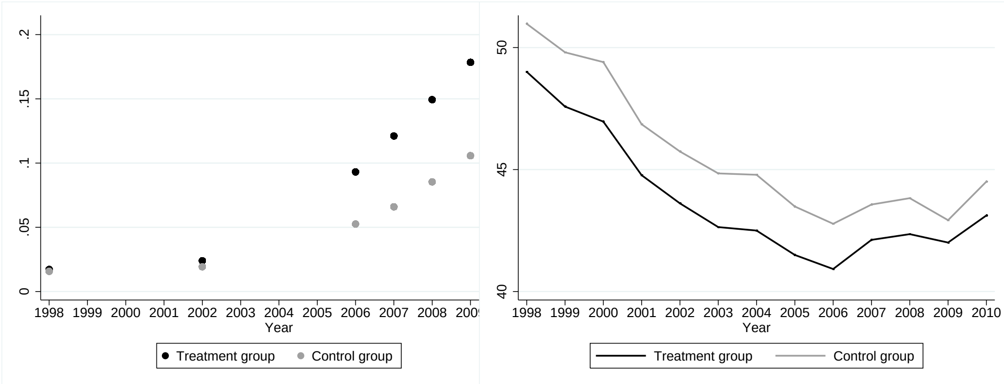

(i) Child care coverage (ii) Birth rates

Notes: The figures show averages of the treatment group (D=1) and the control group (D=0) over time. The treatment group consists of all counties with above median increase in child care coverage rates from 2002 until 2009, whereas the control group consists of counties with below median increase in child care coverage rates from 2002 until 2009. The figure in panel (i) depicts child care coverage over time, separately for treatment and control group. The figure in panel (ii) depicts birth rates over time, separately for treatment and control group.

We observe that treatment and control group counties only marginally differ from each other in the population's age structure. In the age range from 20 to 34, which is most relevant for fertility, differences are tiny and insignificant. In cases where the difference turns out to be significant, the absolute and relative sizes of the differences are very small.[17] Thus, the general picture we get from Table 2 supports our argument that the expansion of child care largely happened in an unpredictable and unsystematic way. Nevertheless, we will use county fixed effects in all our regressions and show that the inclusion of all observable time-varying county characteristics as controls does not affect the results.

Figure 2 shows how child care and birth rates evolved over time, both for treatment and control counties. As is evident from panel (i), child care for under three year olds was virtually non-existent in 1998, with coverage rates of not even 2% for both groups of counties. From 1998 to 2002, we observe hardly any dynamics; the slight rise in coverage rates is due to a decrease in the number of births rather than an increase in the number of child care slots. The federal child care initiatives took effect in 2005. Unfortunately, there are no administrative data available for child care in the period from 2003 to 2005. By 2006, coverage rates had increased to 5.3% in the control group and to 9.2% in the treatment group. In 2010, coverage was already 20.8% in treatment counties and 13.3% in control counties. Thus, while treatment and control group counties started out from the same low level of child care in 2002,

17. The exact differences in means of the detailed age specific variables used in the regressions can be found in Online Appendix Tables A5-A6.

Figure 3. Child care and fertility: Difference-in-differences estimates

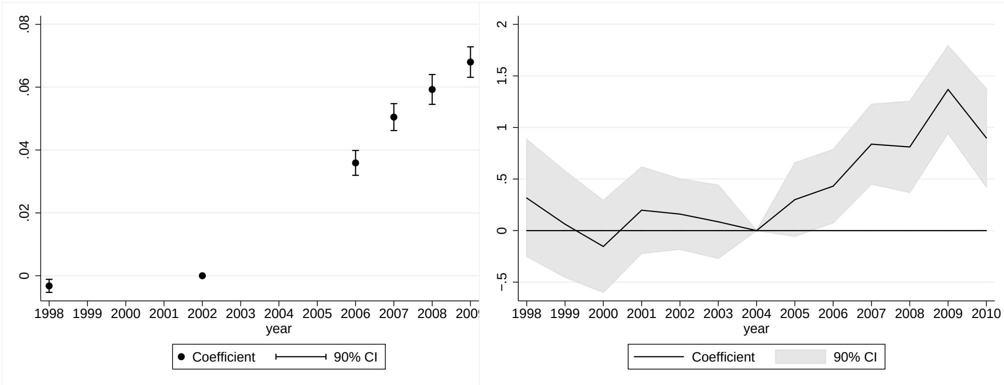

(i) Year-specific gaps in child care coverage

(ii) Year-specific gaps in birth rates

Notes: The figures show the impact of child care expansion starting in 2005 estimated in a difference-in-differences framework that allows for effects before and during the expansion. The treatment group consists of all counties with above median increase in child care coverage rates from 2002 until 2009, whereas the control group consists of counties with below median increase in child care coverage rates from 2002 until 2009. Standard errors are clustered at the county level. The figure in panel (i) depicts emerging differences in child care coverage between treatment and control group counties with the difference normalized to zero in 2002, the last year with data on child care coverage before the 2005 reform. The figure in panel (ii) depicts emerging differences in birth rates between treatment and control group with the difference normalized to zero in 2004, the last year with data on child care coverage before the 2005 reform.

trends diverged in the following years and the gap in coverage rates reached about 7 percentage points in 2010.

Panel (ii) of Figure 2 shows that treatment group counties had lower birth rates than control group counties over the whole period of observation. For both groups of counties, birth rates decreased from 1998 (49.0; 51.0) until 2006 (40.9; 42.8), followed by an upward movement from 2006 until 2010 (43.1; 44.5). The graph provides first suggestive evidence that treatment and control group counties followed a common fertility trend in the pre-treatment period. After 2005, we observe a slight departure from the common trend, with the treatment group level slowly starting to approach the control group level. This is compatible with the hypothesis that fertility increases with the provision of child care.

In a next step, we bring the data to the basic DID framework from equation (1). In particular, we regress the birth rate on a treatment group dummy, year fixed effects and interactions of the treatment group dummy with all year fixed effects without controlling for any other county-specific characteristics. We conduct the same exercise using the child care coverage rate as the dependent variable. The estimated interaction coefficients show year-specific deviations from common birth rates and child care trends. In Figure 3, we plot these interaction coefficients against years.

Panel (i) illustrates the child care coverage gap between treatment and control counties over time. While virtually no differences existed in 1998 and 2002, the gap steadily increased in the second half of the 2000s, reaching a maximum of about 7 percentage points in 2010. In Panel (ii), we focus on the dynamics of the difference in birth rates. It is evident that the birth rate gap did not systematically change from 1998 to 2004. However, from 2005 onwards, the average birth rate in the treatment group increased relative to that in the control group. By 2010, the birth rate in the treatment group had increased by 0.86 births more than in the control group. This pattern suggests a causal influence of child care provision on fertility.

The key identifying assumption of the DID approach is that birth rates in the treatment and control group follow the same trends in absence of the treatment. Figure 3 provides strong evidence for the plausibility of this assumption. Indeed, the estimated coefficients of the pre-treatment year interaction terms are small and far from any conventional significance levels. Thus, there are no systematic yearly deviations from the common birth rate trend in the pre-treatment period. An alternative way of analyzing the common trend assumption is to run a simple two-period placebo DID regression based on pre-reform data. In a non-reported exercise, we used the year 1998 as the baseline period and the year 2004 as the placebo treatment period. As one would expect from Figure 3, this placebo DID regression did not yield any significant treatment effects, which corroborates the validity of our identifying assumption. We also used all pre-treatment years to investigate whether the interaction of a linear time trend with the treatment group dummy is significant. The interaction coefficient turned out be very close to zero and far from any conventional significance levels.

6.2. Generalized difference-in-differences results

We now turn to the generalized DID model from equation (2), which uses a continuous treatment variable. Apart from county and year fixed effects, we control for a county's population density, the male employment rate, GDP per capita, the conservative vote share, and the share of highly educated women. Moreover, we capture emerging differences in the population age structure by including as additional covariates the age specific shares in the subsample of women aged 15 to 44 as well as the age specific shares of 45 year olds and older in the total population.

In Table 3, we report estimates for the effects of child care on birth rates both in year t as well as in year $t + 1$. In Columns 1 and 2, we estimate a stripped-down version which does not control for any time-varying county characteristics. We find positive and significant effects of child care coverage on birth rates. We next include the set of county covariates described above. Again, the model yields robust positive and significant effects of child care on birth rates.[18] The results in Column 3 suggest that increasing child care coverage by 10 percentage points leads to an increase in the birth rate in year t by 1.212 or 2.7% of the sample mean (44.150). In Column 4, the effect of a 10 percentage point increase in child care coverage on birth rates in year

18. The slight decrease of the point estimates can be entirely traced back to controlling for the population structure of female women in fertile age.

Table 3. Estimates from a generalized difference-in-differences model

	Birth rate					
	t	t+1	t	t+1	t	t+1
	(1)	(2)	(3)	(4)	(5)	(6)
Child care coverage	14.968***	15.420***	12.115***	12.275***	13.169***	13.007***
	(3.550)	(3.156)	(2.637)	(2.462)	(3.076)	(2.895)
Revenue, debt, premises	No	No	No	No	Yes	Yes
Regional controls	No	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1,950	1,950	1,950	1,950	1,610	1,610
Number of counties	325	325	325	325	322	322
F-statistic	322.5	254.7	73.62	55.72	69.36	53.50

Note: The table shows the results of generalized difference-in-differences estimations. The outcome variable births per 1,000 women aged 15 to 44 is measured in period t columns (1), (3) and (5), and forwarded by one period in columns (2), (4) and (6). Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Debt and revenue of municipalities are not reported for the federal city states Hamburg and Bremen (including Bremerhaven). Revenue information is missing in 2001 from all 15 Schleswig-Holstein counties. Revenue information in 2009 is not included due to fragmentary raw data. Dwellings denotes controls for the number of newly built homes. Robust standard errors are clustered at the county level and given in parentheses. *** 1 percent significance level; ** 5 percent significance level; * 10 percent significance level.

$t+1$ is estimated to be 1.228 or 2.8% of the sample mean (44.106). The effects are economically meaningful and statistically highly significant.

Apart from the regional characteristics already included as covariates, one might want to control for local public finance. Local public finance might be a determinant of child care expansion. Moreover, prosperous municipalities might be able to provide an attractive environment for young couples to have children. In order to net out these potentially confounding factors, we include municipalities' gross revenue and debts as well as the number of new dwellings in a county as additional control variables.[19] Columns 5 and 6 in Table 3 show that the effect of child care on birth rates remains highly significant and positive. Indeed, the point estimates of 13.169 in t and 13.007 in $t+1$ are if anything larger than the estimates in Columns 3 and 4.

19. Unfortunately, the debt and revenue variables are not available for the city states of Hamburg and Bremen. Moreover, information on municipalities' gross revenue is missing for all counties in the state of Schleswig-Holstein in 2001; since gross revenue information is very fragmentary in 2009, we have to drop this year from our sample.

Figure 4. Child care effect on year-of-age specific birth rates

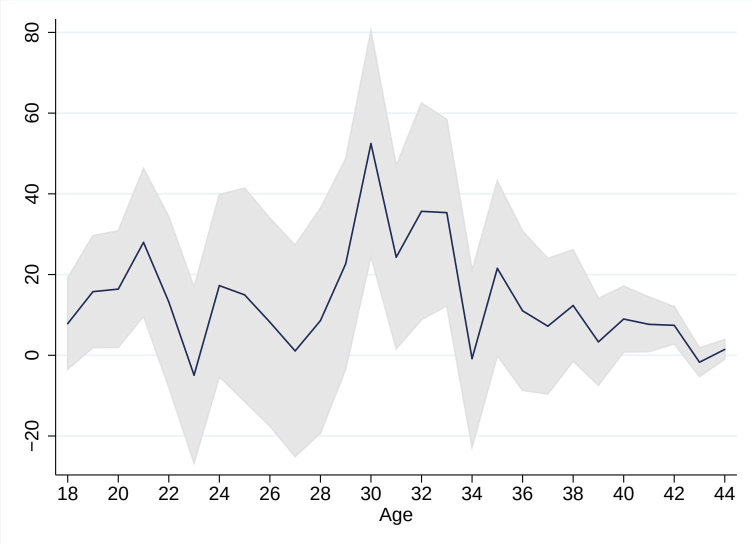

Notes: The figure shows the effect of child care coverage on year-of-age specific birth rates (18-44), number of births of women aged x per 1,000 women aged x, in period t+1 from generalized difference-in-differences regressions. The middle line denotes the average marginal effect, the grey area depicts 90% confidence intervals. All regressions are performed including the control variables as in the baseline estimation. County and year fixed effects are included. Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county.

6.2.1. Effect heterogeneity and children's birth outcomes. The birth registry data includes valuable information on individual births which we exploit in order to investigate the possibility of heterogenous reform effects. First, we run separate regression for all age-specific birth rates, and plot the point estimates together with 90 percent confidence intervals (Figure 4). The exercise reveals positive fertility effects almost throughout the entire age distribution. The effects turn out to be particularly strong and significant for women aged 29 to 33.[20]

Second, we investigate whether fertility responded to the increase in child care at the extensive or at the intensive margin. For children born to married mothers, the birth registry data provides information on the birth order of children.[21] We first split the dependent variable by mothers' marital status and find that the fertility effects are

20. We also find some significant effects on teen mothers, mostly aged 19. A closer inspection of the data reveals that effects on under-aged women, 15 to 17 years old, are much smaller and less robust. The rare incidence of these births calls for caution when interpreting the estimates for this subgroup.

21. Unfortunately, the data do not include unique identifiers for mothers, which would allow us to assign all births to a specific mother.

largely driven by married women (Online Appendix Table A9). This result suggests that the restriction to within-marriage births is not a major limitation. Also note that the vast majority of births in our period of observation (78.3%) occurs within marriages. We construct variables measuring the number of first births, second births, third births, and fourth births per 1,000 women of reproductive age in a county. The results for these birth order specific birth rates show that the child care expansion particularly increased the number of second and third births (Table 4). The highly significant point estimates suggest that a 10 percentage point increase in child care increases the incidence of second and third births by 3.9% and 7.5%, respectively. The effect on first births is 0.5% in t and 2.0% in $t+1$, but only the latter is statistically significant. Thus, the effects of this child care expansion on fertility are stronger at the intensive than at the extensive margin.[22]

Third, we are interested in the question whether increases in child care affected the health outcomes of newborns. This might be the case if the expansion of child care changes the composition of parents (Dehejia and Lleras-Muney, 2004). The birth registry data provides information on each newborn's length (in cm) and weight (in gram). In addition, we compute a low birth weight indicator variable which equals unity for birth weights below 2,500 grams. Moreover, we combine the birth weight and birth length information in a Ponderal index, defined as weight (in kg) by length (in meter^3). Based on this index, we generate two dichotomous variables indicating births in the lowest and highest 10 percentiles of the Ponderal index distribution. Using the county averages of these five birth outcome measures as dependent variables, we find no effects of child care on newborns' health outcomes in period $t+1$ (Table 5).[23] The insignificant point estimate for birth length is 0.15 cm, which corresponds to less than half a percent of the average birth length of 51 cm. The point estimate for birth weight of -45.07 grams, which corresponds to 1.4 percent of average birth weights of 3,335 grams, is also small and statistically insignificant. We also obtain small and insignificant estimates when using the low birth weight indicator and the two Ponderal index measures as dependent variables. Overall, the results suggest that the expansion of child care did not change the characteristics of mothers in a way detrimental to infant health.

Fourth, we investigate the extent to which the fertility effect varies with county-level characteristics. We find that the effect of child care on fertility is stronger in counties with a larger share of highly educated women and a lower conservative vote share (Online Appendix Table A8). Note, however, that this heterogeneity with respect to regional characteristics is not necessarily driven by individuals who carry the characteristics (*ecological fallacy*).

22. While there were no federal rules on how child care slots ought to be allocated among parents, we cannot rule out that mothers with previously enrolled children were preferentially treated in the allocation process. This and lower parental fees for higher order children may be reasons for the birth pattern in Table 4.

23. Estimates for period t are very similar and shown in Online Appendix Table A7.

Journal of the European Economic Association
Preprint prepared on 7 July 2015 using jeea.cls v1.0.

TABLE 4. Birth order specific birth rates

	Birth order specific birth rate							
	1st births		2nd births		3rd births		4th births	
	t	t+1	t	t+1	t	t+1	t	t+1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Child care coverage	0.767	3.057**	5.037***	5.362***	3.198***	3.121***	0.538*	0.409
	(1.522)	(1.430)	(1.181)	(1.239)	(0.683)	(0.659)	(0.306)	(0.299)
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Percent effect of 10pp childcare increase	0.005	0.020	0.039	0.041	0.075	0.073	0.047	0.036
N	1,950	1,950	1,950	1,950	1,950	1,950	1,950	1,950
Number of counties	325	325	325	325	325	325	325	325
F-statistic	122.6	62.91	94.92	89.08	33.48	27.13	13.73	9.076

Note: The table shows the results of generalized difference-in-differences estimations. The outcome variables are within-marriage birth order specific fertility rates, i.e., the number of 1st births births per 1,000 women (column 1), 2nd births (column 2), 3rd births (column 3), and 4th births (column 4). All outcome variable are forwarded by one period. Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Robust standard errors are clustered at the county level and given in parentheses. *** 1 percent significance level; ** 5 percent significance level; * 10 percent significance level.

Table 5. Investigating marginal birth outcomes

	Birth length (cm) t+1 (1)	Birth weight (grams) t+1 (2)	Low birth weight t+1 (3)	Ponderal below p10 t+1 (4)	Ponderal above p90 t+1 (5)
Child care coverage	0.154	-45.067	-0.012	0.034	-0.040
	(0.283)	(33.008)	(0.013)	(0.031)	(0.025)
Regional controls	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes
N	1,945	1,945	1,945	1,945	1,945
Number of counties	325	325	325	325	325
F-statistic	5.07	13.88	2.60	2.83	22.68

Note: The table shows the results of generalized difference-in-differences estimations on outcomes in period t+1. Outcome variables birth length and birth weight are averages over all births in a county. Low birth weight is the county average of a dummy variable equal to one for birth weights below 2,500 grams. Ponderal index measures are county averages of indicators that are equal to one if the Ponderal index is below the 10th percentile resp. above the 90th percentile of the German Ponderal index distribution from 1998 to 2010 ($Ponderal = weight(kg)/height(m)^3$). Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Outcomes are missing in 1999 through 2008 for Aachen and in 1999 for Hannover, as we cannot recover the means after the counties' borders have changed. Robust standard errors are clustered at the county level and given in parentheses. *** 1 percent significance level; ** 5 percent significance level; * 10 percent significance level.

6.2.2. Timing of the fertility response. We exploit variation in child care coverage over time to identify effects on an outcome (births) that can only react with a certain lag. Therefore, it is essential to get the timing of cause and effect right. We have decided to estimate the effect of child care coverage in period t both on births in period t as well as on births in period $t+1$. Recall that child care is measured at the beginning of March and births are measured as the sum of births in a specific year. As child care centers typically align their service to school years (August or September to June or July), the coverage rate measured in March may partly reflect coverage from August or September in the preceding year. Furthermore, utilized child care slots in March must already have been established at some time in the preceding 12 months. As a consequence, fertility may already react to changes in child care in period t. Using births in period $t+1$ as an additional outcome variable means that we allow some more time for conception and pregnancy—in sum 10 to 22 months after child care is observed in March. Although this is in line with the literature on fertility responses to public policies (see, e.g., Rindfuss et al., 2010), it is important to test empirically whether our specification is valid.

Figure 5. Timing of the dependent variable

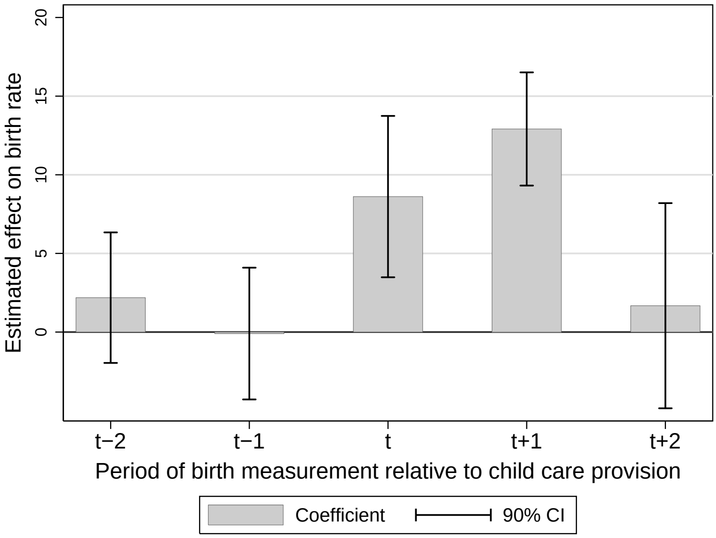

Notes: The bars indicate the effect of child care coverage in period t on births per 1,000 women in the period according to the x-axis. Horizontal lines indicate 90-percent confidence intervals. All five regressions are independently estimated using the generalized difference-in-differences approach. Control variables including child care coverage are lagged by one year and are included in period x-1. Only for x-1=t the control variable child care coverage is identical to the variable of interest child care coverage in period t.

In Figure 5, we show effects of child care measured in period t on births in period x, while controlling for child care and all other covariates in period $x-1$. Accordingly, in the fourth column ($x = t+1$) we observe the effect of child care in t on births in $t+1$, which is exactly the effect from the generalized difference-in-differences model on births in $t+1$ (12.275). To check whether our timing specification is reasonable, we now shift the outcome and control variables on the timeline. The middle bar reveals that even if we control for the relevant child care in $t-1$, child care in t has a positive and significant effect on births in t.

The expansion of child care within counties might not be independent over periods, which raises concerns about spurious estimates. To be sure that we do not only pick up pure child care expansion trajectories, we estimate the effect of child care in t on births in $t+2$ while controlling for child care in $t+1$. As can be seen from the far right bar, there is no significant effect of child care in t on births in $t+2$. We apply the same procedure for the other direction on the timeline and find that there is no effect of child care in t on births in $t-1$ (and $t-2$ respectively), conditional on child care in $t-2$ (and $t-3$ respectively). Indeed, the coefficients are very close to zero and far away from any conventional significance levels. This result also provides evidence against any reverse causality concerns. Taken together, our decision to investigate the effects of child care in t on births in t and $t+1$ is not only in line with the previous literature but also empirically well founded.

Table 6. Age at birth by birth order

	Mothers' age at birth							
	1st births		2nd births		3rd births		4th births	
	t	t+1	t	t+1	t	t+1	t	t+1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Child care coverage	-0.091	-0.088	0.713*	0.137	0.200	-0.227	1.628	2.269*
	(0.426)	(0.387)	(0.414)	(0.423)	(0.639)	(0.738)	(1.161)	(1.190)
Regional controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1,950	1,950	1,950	1,950	1,950	1,950	1950	1,950
Number of counties	325	325	325	325	325	325	325	325
F-statistic	111.8	88.04	119.1	104.1	42.03	31.84	15.81	11.18

Note: The table shows the results of generalized difference-in-differences estimations. The outcome variables are mothers' average age at birth by within-marriage birth order, i.e., the average age of mothers at 1st births (column 1), 2nd births (column 2), 3rd births (column 3), and 4th births (column 4). All outcome variable are forwarded by one period. Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Robust standard errors are clustered at the county level and given in parentheses. *** 1 percent significance level; ** 5 percent significance level; * 10 percent significance level.

Journal of the European Economic Association
Preprint prepared on 7 July 2015 using jeea.cls *v1.0.*

6.2.3. Effects on completed fertility? Our dependent variable is a period fertility measure. It shares most features with the total fertility rate and is as such suitable to analyze the short-run effects of policy changes. Yet, in contrast to completed (or cohort) fertility, which reflects the actual number of births per woman measured after the reproductive age, it can be distorted by changes in the timing of births. If couples, as a response to the child care expansion, decide to have children earlier in life but not to have more children over the course of life, we might see a short-term increase in period fertility without there being a long-term effect on completed fertility.

Therefore, we now investigate whether the evidence so far can be interpreted as a positive long-term fertility effect, or whether it might rather reflect changes in the timing of births. If intended births are brought forward, *ceteris paribus*, the age at which women give birth to children should decrease. However, a simple comparison of age at birth is not sufficient. In particular, since our estimates suggest a larger effect for higher order births, the average age at birth increases mechanically. Therefore, to test for the possibility of births being brought forward, we need to assess the age at birth by each parity separately. Accordingly, we use the generalized DID model to estimate the effect of child care on mothers' age at birth separately for first births, second births, third births, and fourth births. As can be seen from Table 6, we find no evidence that women get children earlier in life in response to increases in child care. Quite to the contrary, the average age of women significantly increases both at second and fourth births.[24] Thus, the fertility effect of child care we identified does not seem driven by births brought forward in time. Instead, the evidence is suggestive of positive effects on long-term completed fertility.

6.2.4. Maternity leave reform as a confounder? In 2007, the federal government enacted a major parental leave reform. The reform constituted a shift from a means-tested parental leave benefit targeted at lower-income families to a system which benefits higher-earning women by tying parental leave benefits to pre-birth earnings. Since this federal reform applied to all German counties, year fixed effects absorb the effects of the reform unless these effects vary systematically between counties. In particular, if the reform had a larger impact on birth rates in counties that increased child care coverage more strongly, our estimates would be upwardly biased. We now investigate this issue. Raute (2014) shows that the parental leave reform had positive effects on fertility, driven mainly by women in the upper-end of the education and income distribution. Thus, it is important to allow the relationship between female education and fertility to change after the parental leave reform in 2007. We therefore now interact the share of highly-educated women at the county level with a post-2007 dummy. As can be seen from the first two columns of Table 7, the interaction coefficient is positive and reaches significance for the fertility outcome measured in t. The pattern remains robust if we additionally control for local public finance in

24. Estimates for fourth births should be interpreted with caution as they are based on few observations. Less than 4% of all births within marriages are fourth births.

TABLE 7. Taking account of the 2007 maternity leave reform

	Birth rate			
	t	t+1	t	t+1
	(1)	(2)	(3)	(4)
Child care coverage	10.846***	11.469***	12.007***	12.435***
	(2.707)	(2.552)	(3.156)	(2.982)
Maternity leave reform	3.749**	2.382	3.279*	1.616
× female high education	(1.687)	(1.638)	(1.751)	(1.765)
Revenue, debt, premises	No	No	Yes	Yes
Regional controls	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
N	1,950	1,950	1,610	1,610
Number of counties	325	325	322	322
F-statistic	76.00	56.77	69.79	53.54

Note: The table shows the results of generalized difference-in-differences estimations. The maternity leave reform dummy is one for all years 2007 and later and zero otherwise. The outcome variable births per 1,000 women aged 15 to 44 is measured in period t columns (1) and (3), and forwarded by one period in columns (2) and (4). Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Debt and revenue of municipalities are not reported for the federal city states Hamburg and Bremen (including Bremerhaven). Revenue information is missing in 2001 from all 15 Schleswig-Holstein counties. Revenue information in 2009 is not included due to fragmentary raw data. Dwellings denotes controls for the number of newly built homes. Robust standard errors are clustered at the county level and given in parentheses. *** 1 percent significance level; ** 5 percent significance level; * 10 percent significance level.

Columns 3 and 4. Thus, this result is consistent with Raute's (2014) findings. More importantly, however, we still find positive and highly significant effects of public child care on fertility across all four specifications. Indeed, the effect of public child care coverage on fertility stays similar in size as compared to the main results in Table 3. Apart from the parental leave reform, we are not aware of any other policy intervention which might have affected fertility.

6.2.5. Selective migration. Although we control for unobserved time-invariant heterogeneity between counties and include an extensive set of time-varying county characteristics, selective migration of potential mothers might confound the estimates. In particular, if women who are pregnant or plan to have a child systematically move to counties that strongly increase child care, the regressions would yield upwardly biased estimates. If this type of selective migration is a relevant concern, it should show up in the data in the form of higher gross in-migration flows or lower gross out-migration flows in counties that substantially increased coverage rates. Unfortunately, we do not have information on age-specific female in- and out-migration flows.

Table 8. Gross migration flows

	Child care coverage (1)	S.E. (2)	Regional Controls, Year & county FE (3)	N (4)	Counties (5)	F-stat (6)
Dependent var						
In-migrants 18-29 (f), t	-0.001	(0.002)	Yes	1,625	325	7.46
In-migrants 18-29 (f), t+1	0.002	(0.002)	Yes	1,625	325	4.62
In-migrants 30-49 (f), t	-0.001	(0.002)	Yes	1,625	325	10.77
In-migrants 30-49 (f), t+1	-0.001	(0.001)	Yes	1,625	325	3.92
In-migrants 18-49 (f), t	-0.003	(0.003)	Yes	1,625	325	7.18
In-migrants 18-49 (f), t+1	0.001	(0.003)	Yes	1,625	325	5.46
Out-migrants 18-29 (f), t	-0.002	(0.001)	Yes	1,625	325	55.61
Out-migrants 18-29 (f), t+1	-0.001	(0.001)	Yes	1,625	325	21.39
Out-migrants 30-49 (f), t	-0.003	(0.002)	Yes	1,625	325	3.37
Out-migrants 30-49 (f), t+1	-0.002*	(0.001)	Yes	1,625	325	4.05
Out-migrants 18-49 (f), t	-0.004	(0.003)	Yes	1,625	325	6.50
Out-migrants 18-49 (f), t+1	-0.004	(0.002)	Yes	1,625	325	11.22

Note: The table shows the results of generalized difference-in-differences estimations. The outcome variables are denoted in rows and all are defined as fractions of the population within a county. Regional control variables include the county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure control variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44 and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Estimates in rows are from independent regressions. Robust standard errors are clustered at the county level and given in parentheses. *** 1 percent significance level; ** 5 percent significance level; * 10 percent significance level.

However, we can analyze in-migration and out-migration patterns of two age groups, namely 18 to 29 year old and 30 to 49 year old women. Thus, we now estimate the generalized DID model using the ratio of female in-migrants aged 18 to 29 in t over the total population in $t-1$ as the dependent variable (Row 1 in Table 8). Furthermore, we also look at the ratio of female in-migrants aged 18 to 29 in $t+1$ over the total population in t as the outcome variable (Row 2). Moreover, we use the ratio of female in-migrants aged 30 to 49 in t and $t+1$ (Rows 3 and 4) as well as the ratio of female in-migrants aged 18 to 49 in t and $t+1$ (Rows 5 and 6). Similarly, we compute the respective out-migration ratios and use these variables as dependent variables (Rows 7 to 12). We do not find any evidence that child care affects in-migration of women of reproductive age. Only for 30 to 49 year old out-migrants in $t+1$ a marginally significant negative effect appears. With only one out of twelve coefficients reaching marginal significance at the 10 percent level, we cautiously conclude that selective migration does not seem to play a major role in our setting.

Figure 6. Child care and female employment: Difference-in-differences estimates

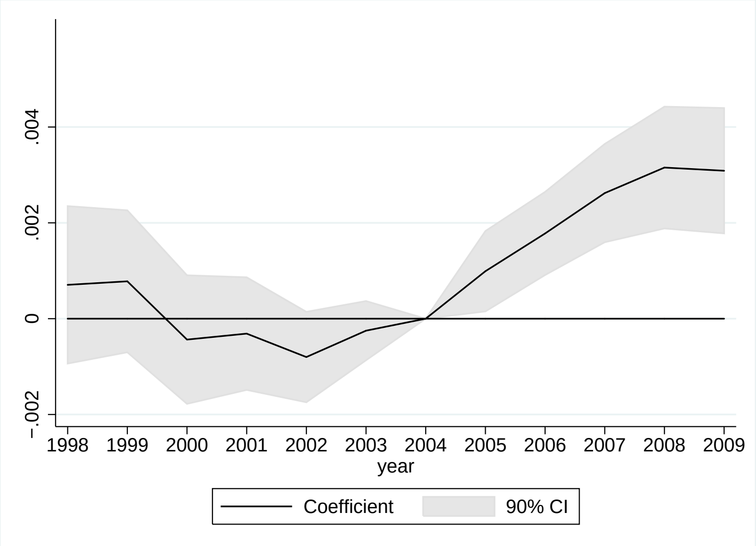

Notes: The figure shows the impact of child care on female employment estimated in a difference-in-differences framework that allows for effects before and during the child care expansion which started in 2005. The treatment group consists of all counties with above median increase in child care coverage rates from 2002 until 2009, whereas the control group consists of counties with below median increase in child care coverage rates from 2002 until 2009. Apart from county fixed effects, the model includes as controls population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the interpolated share of high educated females until age 44 as well as an extensive set of age structure controls. Age structure variables include the year-of-age share of 15 to 44 year old women over all women aged 15 to 44. and the year-of-age shares of 45 to 74 year old and 75 plus years old people over the population in each county. Standard errors are clustered at the county level.

6.2.6. Labor supply effects. A central result that emerges from public finance theory is that countries which support families through child care facilities rather than child cash payments are likely to have both higher fertility and higher female labor supply (Apps and Rees, 2004). The underlying intuition comes from a differential effect on female labor supply and hence on the tax base. Child care provision encourages female labor supply, increases the tax base, and so increases size of the policy changes that can be implemented. Cash subsidies discourage female labor supply via an income effect, reduce the tax base, and so only allow for small policy changes. Against this background, it is interesting to ask whether the child care reform under consideration also had a positive impact on female employment. Thus, we now estimate the basic DID model from equation (1) using the female employment rate as the dependent variable. Apart from county fixed effects, we control for a county's population density, GDP per capita, the male employment rate, the interpolated conservative vote share, the interpolated share of highly educated women as well as the extensive set of age structure controls. Figure 6 plots the results. What we find are employment effects which are remarkably similar to the fertility effects presented in Figure 3. While treatment and control group counties follow the same female employment trend during the pre-treatment years, there is a deviation from this common trend for the post-reform years of 2005-2009. The overall picture, therefore, is one of both higher fertility and higher female employment due to the expansion of public child care.[25]

6.3. Cost-Benefit Analysis: Child Care versus Child Benefits

In a final step, we ask whether the fertility effects identified in this paper could have been reached at a lower cost by alternatively increasing child benefits. Gauthier and Hatzius (1997) estimate a long-run cross-country elasticity of fertility with respect to child benefits of 0.16.[26] Currently, Germany spends €40 billion per year on child benefits.[27] Accordingly, increasing child benefits by 1 percent would cause additional expenses of €400 million and result in a 0.16% increase in fertility. Alternatively, the €400 million could be used to provide 58,823 public child care slots, which would result in an increase of the child care coverage rate by about 2.9 percentage points. Relying on our findings, an increase of child care coverage rates by 2.9 percentage points yields an increase in fertility by 0.82%. Thus, the fertility effect of spending €400 million on public child care is about five times larger than the effect of spending the money on increases in child benefits. In addition, public child care, in contrast to child benefits, also increases maternal employment. As a result, state revenues from income tax payments and social security contributions increase. This effect further increases the relative effectiveness of child care versus child benefits.

7. Concluding Remarks

The question of whether family policies, such as affordable child care, can positively affect the private choice to have children has gained importance over the past decades. However, most of what we know about the link between child care provision and fertility comes from studies that have suffered from the limited magnitude of the available policy variation. In this paper, we have made a first step to overcome this problem by evaluating the impact on fertility of a major German child care reform. The reform we study led to a significant expansion of child care slots for young

25. The employment effects we have uncovered are consistent with the evidence in Milligan (2014), which shows an increase in the employment rates for German mothers with young children between 2000 and 2010. An interesting avenue for future research is to examine whether the expansion of child care in Germany has helped women break through the "glass ceiling". In particular, Blau and Kahn (2013) argue that some family-friendly policies (e.g., long, paid parental leaves) have the unintended side effects of generating a reliance on part-time employment for women and lower female representation in high-level positions, and it would be interesting to analyze whether child care provision has the opposite effect.

26. More recent studies yield comparable results. Milligan (2005) finds a benefit elasticity of 0.11 for Canada, while Cohen et al. (2013) estimates a benefit elasticity of 0.19 for Israel.

27. Government expenditure for child benefits and child tax credit were €39.95 billion in 2010. Source: http://www.bmfsfj.de/RedaktionBMFSFJ/Abteilung2/Pdf-Anlagen/familienbezogene-leistungen-tableau-2010.

children, and our empirical strategy exploits the temporal and spatial variation in child care coverage induced by this expansion. First, we apply a basic difference-in-differences model that compares a treatment group of counties with above-median child care expansion to a control group of counties with below-median child care expansion over time. Second, we use a generalized difference-in-differences estimator that utilizes the full variation in treatment intensities. Results from both specifications show consistently that the provision of public child care has a significant positive effect on fertility. In particular, our results suggest that a 10 percentage point increase in child care coverage leads to an increase in birth rates of 2.8%. Results are neither confounded by regional fundamentals and demographics nor by selective migration into counties with strong child care expansion. We find no evidence that our results are driven by births brought forward in time, which suggests a positive effect on completed fertility. We show that effects of child care on birth rates materialize mainly at the intensive margin. Finally, we find no indication that the children born in response to the reform have inferior health outcomes. Taken together, the results presented in this paper suggest that there is nothing inevitable about very low fertility rates: policies that facilitate the combination of parenthood and employment hold the promise of being an effective way to boost birth rates where these rates are considered to be too low.

References

Apps, Patricia and Ray Rees (2004). "Fertility, Taxation and Family Policy." *Scandinavian Journal of Economics*, 106(4), 745–763.

Baker, Michael, Jonathan Gruber, and Kevin Milligan (2008). "Universal Childcare, Maternal Labor Supply, and Family Well-Being." *Journal of Political Economy*, 116(4), 709–745.

Bauernschuster, Stefan and Martin Schlotter (2015). "Public Child Care and Mothers' Labor Supply: Evidence from Two Quasi-Experiments." *Journal of Public Economics*, 123, 1–16.

Becker, Gary S (1960). "An Economic Analysis of Fertility." In *Demographic and Economic Change in Developed Countries*. Universities-National Bureau Conference Series No. 11, Princeton: Princeton University Press.

Bien, Walter, Thomas Rauschenbach, and Birgit Riedel (2006). *Wer betreut Deutschlands Kinder?* DJI-Kinderbetreuungsstudie, München.

Björklund, Anders (2006). "Does Family Policy Affect Fertility?" *Journal of Population Economics*, 19(1), 3–24.

Blau, Francine and Lawrence Kahn (2013). "Female Labor Supply: Why Is the United States Falling Behind?" *American Economic Review: Papers & Proceedings*, 103(3), 251–256.

Cascio, Elizabeth U. (2009). "Maternal Labor Supply and the Introduction of Kindergartens into American Public Schools." *Journal of Human Resources*, 44(1), 140–170.

Cigno, Alessandro (1993). "Intergenerational Transfers without Altruism: Family, Market and State." *European Journal of Political Economy*, 9(4), 505–518.

Cohen, Alma, Rajeev Dehejia, and Dmitri Romanov (2013). "Financial Incentives and Fertility." *Review of Economics and Statistics*, 95(1), 1–20.

Dehejia, Rajeev and Adriana Lleras-Muney (2004). "Booms, Busts, and Babies' Health." *Quarterly Journal of Economics*, 119(3), 1091–1130.

Del Boca, Daniela (2002). "The Effect of Child Care and Part Time Opportunities on Participation and Fertility Decisions in Italy." *Journal of Population Economics*, 15(3), 549–573.

Del Boca, Daniela, Silvia Pasqua, and Chiara Pronzato (2009). "Motherhood and Market Work Decisions in Institutional Context: a European Perspective." *Oxford Economic Papers*, 61(1), i147–i171.

Demeny, Paul (1986). "Population and the Invisible Hand." *Demography*, 23(4), 473–487.

Destatis (2010). *Kindertagesbetreuung Regional 2009: Ein Vergleich aller 413 Kreise in Deutschland*. Statistisches Bundesamt, Wiesbaden.

DJI (2005). "Zahlenspiegel 2005 - Kinderbetreuung im Spiegel der amtlichen Statistik." Tech. rep.

Dorbritz, J. (2008). "Germany: Family Diversity with Low Actual and Desired Fertility." *Demographic Research*, 19(17), 557–598.

Ermisch, John (2003). *An Economic Analysis of the Family*. Princeton: Princeton University Press.

Felfe, Christina and Rafael Lalive (2012). "Early Child Care and Child Development: For Whom it Works and Why." CESifo Working Paper Series 4043, CESifo Group Munich, URL http://ideas.repec.org/p/ces/ceswps/_4043.html.

Fenge, Robert and Volker Meier (2005). "Pensions and Fertility Incentives." *Canadian Journal of Economics*, 38(1), 28–48.

Folbre, Nancy (1994). "Children as Public Goods." *American Economic Review*, 84(2), 86–90.

Gauthier, Anne Helene and Jan Hatzius (1997). "Family Benefits and Fertility: An Econometric Analysis." *Population Studies*, 51(3), 295–306.

Gelbach, Jonah (2002). "Public Schooling for Young Children and Maternal Labor Supply." *The American Economic Review*, 92(1), 307–322.

Groezen, Bas van, Theo Leers, and Lex Meijdam (2003). "Social Security and Endogenous Fertility: Pensions and Child Allowances as Siamese Twins." *Journal of Public Economics*, 87(2), 233–251.

Hank, Karsten and Michaela Kreyenfeld (2003). "A Multilevel Analysis of Child Care and Women's Fertility Decisions in Western Germany." *Journal of Marriage and Family*, 65(3), 584–596.

Hank, Karsten, Michaela Kreyenfeld, and C Katharina Spieß (2004). "Kinderbetreuung und Fertilität in Deutschland." *Zeitschrift für Soziologie*, 33(3), 228–244.

Hank, Karsten, Katja Tillmann, and Gert G. Wagner (2001). "Außerhäusliche Kinderbetreuung in Ostdeutschland vor und nach der Wiedervereinigung: ein Vergleich mit Westdeutschland in den Jahren 1990-1999." MPIDR Working Paper

2001-03.

Havnes, Tarjei and Magne Mogstad (2011a). "Money for Nothing? Universal Child Care and Maternal Employment." *Journal of Public Economics*, 95(11-12), 1455–1465.

Havnes, Tarjei and Magne Mogstad (2011b). "No Child Left Behind: Universal Child Care and Children's Long-Run Outcomes." *American Economic Journal: Economic Policy*, 3(2), 97–129.

Huesken, Kathrin (2011). "Kita vor Ort: Betreuungsatlas auf Ebene der Jugendamtsbezirke 2010." Tech. rep.

Lefebvre, Pierre and Philip Merrigan (2008). "Child-Care Policy and the Labor Supply of Mothers with Young Children: A Natural Experiment from Canada." *Journal of Labor Economics*, 26(3), 519–548.

Leibenstein, Harvey (1957). *Economic Backwardness and Economic Growth*. New York: Wiley & Sons.

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

Milligan, Kevin (2014). "The Road to Egalitaria: Sex Differences in Employment for Parents of Young Children." *CESifo Economic Studies*, 60, 257–279.

Mörk, Eva, Anna Sjögren, and Helena Svaleryd (2013). "Childcare Costs and the Demand for Children—Evidence from a Nationwide Reform." *Journal of Population Economics*, 26, 33–65.

NICHD – Early Child Care Research Network (2003). "Does Amount of Time Spent in Child Care Predict Socioemotional Adjustment During the Transition to Kindergarten?" *Child Development*, 74(4), 976–1005.

NICHD – Early Child Care Research Network (2004). "Multiple Pathways to Early Academic Achievement." *Harvard Educational Review*, 74(1), 1–29.

Population Reference Bureau (2007). "2007 World Population Data Sheet."

Raute, Anna (2014). "Do Financial Incentives Affect Fertility: Evidence from a Reform in Maternity Leave Benefits." University College London mimeo.

Rindfuss, Ronald R., David K. Guilkey, S. Philip Morgan, and Øystein Kravdal (2010). "Child-Care Availability and Fertility in Norway." *Population and Development Review*, 36(4), 725–748.

Sinn, Hans-Werner (2004). "The Pay-as-you-go Pension System as Fertility Insurance and an Enforcement Device." *Journal of Public Economics*, 88(7), 1335–1357.

Takayama, Noriyuki and Martin Werding (2011). *Fertility and Public Policy: How to Reverse the Trend of Declining Birth Rates*. Cambridge, MA: MIT Press.

World Bank (2009). "World Development Indicators 2009." Tech. rep.

Wrohlich, K. (2008). "The Excess Demand for Subsidized Child Care in Germany." *Applied Economics*, 40(10), 1217–1228.