Public Child Care and Mothers' Labor Supply—Evidence from Two Quasi-Experiments

Stefan Bauernschuster *, Martin Schlotter ‡

Forthcoming in the Journal of Public Economics

Abstract

Public child care is expected to assist families in reconciling work with family life. Yet, empirical evidence for the relevance of public child care to maternal employment is inconclusive. We exploit the introduction of a legal claim to a place in kindergarten in Germany, which was contingent on day-of-birth cut-off dates and resulted in a marked increase in kindergarten attendance of three-year olds in the following years. Instrumental variable and difference-in-differences estimations on two individual-level data sets yield positive effects of public child care on maternal employment. A set of placebo treatment tests corroborate the validity of our identification strategies.

Keywords: public child care, maternal employment

JEL Codes: J22; J13; D04

Acknowledgments: We are indebted to the handling editor Joseph Doyle, two anonymous referees, Martha Bailey, Rainald Borck, Oliver Falck, Martin Halla, Timo Hener, Christian Holzner, Helmut Rainer, Ludger Woessmann, Josef Zweimueller, participants of the CESifo Area Conference on Employment and Social Protection in Munich, the SOLE Annual Meeting in Chicago, the meeting of the German Economic Association in Goettingen, the EALE conference in Bonn, the RES conference in London, the E.ON Stipendienfonds & Research Council of Norway workshop on Family Policy in Dresden as well as seminar participants at the University of Munich, the University of Oslo, the University of Stockholm, and the ifo Institute for valuable comments and suggestions.

* University of Passau, Innstr. 27, 94032 Passau (Germany), Phone +49(0)851/509-2540, E-mail: stefan.bauernschuster@uni-passau.de; and CESifo, and Ifo Institute.
‡ Bavarian Ministry of Economic Affairs and Media, Energy, and Technology, Prinzregentenstr. 28, 81679 Munich (Germany), Phone: +49(0)89/2162-2275, E-mail: martin.schlotter@stmwi.bayern.de.

1 Introduction

Recent years have witnessed substantial political effort to increase public child care provision in many industrialized countries. Providing subsidized child care is supposed to promote reconciliation of work and family life for the mothers of young children and increase their labor force participation. However, empirical studies find mixed results for the actual effects of subsidized public child care on mothers' employment. For some countries, economists have identified clearly positive effects. For other countries, reforms which aimed at increasing availability or affordability of public child care had zero effects on maternal employment (or effects that are substantially smaller than the take-up rate). In particular, marginal decreases in the costs of child care do not affect maternal labor supply if employment rates and child care attendance rates are already high. Further, we should not expect substantial employment effects if newly available public child care slots simply crowd-out existing private child care arrangements. Crowding-out might be particularly relevant if public child care slots are still severely rationed after expansion. It might be that the newly provided public child slots are then primarily given to mothers already closely attached to the labor market, and that these mothers now just substitute private care with public care arrangements.

This paper contributes to the growing empirical literature on child care and maternal employment by providing first quasi-experimental evidence from a German public child care reform introduced in 1996. At that time, West Germany had long been known for its low female and, in particular, maternal labor force participation. Indeed, the employment gap between mothers of three to four year old children and women of a similar age but without any children amounted to 40 percentage points in 1995. To improve reconciliation of family and working life, expansion of public child care has been at the top of the political agenda for the past two decades. One of the most prominent political reforms in this context was the introduction of the legal claim to a place in kindergarten (Rechtsanspruch auf einen Kindergartenplatz) in 1996. Since that year, children from age three until school entry are eligible to attend highly subsidized half-day public child care. Thus, this reform intended to abandon any rationing of public child care for three to six year olds. While public child care attendance by five- and six-year old children was already at a very high level of 90 percent prior to the reform in West Germany, attendance by three- and four-year old children was substantially lower at 30 percent and 60 percent respectively. Despite the fact that public provision of child care has thus been far from universal (and non-existing for children under three), virtually no private market for child care had emerged. As expected the introduction of the legal claim to a place in kindergarten triggered a sharp increase in public child care attendance by three- and four-year olds from 1996 to 2001. To the best of our knowledge, the effects of this seminal German child care reform have not yet been analysed, although it might give us interesting insights which are relevant for systemizing recent findings in the literature on public child care and maternal employment. [1] In particular, our reading of the previous literature suggests that we should expect large and positive effects of public child care on maternal labor supply in West Germany since a) the reform provided highly subsidized universal public child care, b) public child care was far from universal prior to the reform, c) there was no private market for child care which could be crowded out by a public child care expansion, and d) maternal employment rates were low at the time the reform set in despite the fact the females were very well educated.[2]

The special features of the German policy reform in 1996 allow us to conduct both instrumental variables and difference-in-differences estimations within a single paper. In our first identification strategy, we exploit the fact that after introduction of the legal claim to a place in kindergarten in 1996, some municipalities were confronted with severe problems since they could not meet the increased demand for child care. This is why the German Federal Parliament (Deutscher Bundestag) adopted a legislative initiative proposed by the Federal Council of Germany (Bundesrat) that allowed communities to introduce day-of-birth cut-off rules for determining whether or not a child was eligible for public child care. The respective cut-off dates were the start dates of the "kindergarten year", which vary over federal states and years and typically coincide with the start dates of the school year. Due to these cut-off rules, in an extreme case, children turning three years old slightly after the cut-off date, i.e., the start date of the "kindergarten year", could not enter kindergarten right at their third birthday but only became eligible to attend kindergarten one year later, i.e., at the date the following "kindergarten year" started. These day-of-birth cut-off rules provide us with exogenous variation in attendance at German kindergarten, which we exploit in a 2SLS instrumental variable approach. Note that this approach identifies the local average treatment effect of public child care on maternal employment for the group of compliers, i.e., for those mothers who live in a municipality that applies the cut-off rule, who get a slot in public child care due to this rule and immediately take this slot but would not have been able to have a public child care slot otherwise.

---

[1] Most of the German evidence so far is based on predominantly descriptive studies or micro-simulations. Buechel and Spiess (2002) identify a significant positive association between regional public child care provision for three- to six-year old children and mothers' employment. Bick (2011) uses a micro simulation model to show that a lack of public child care in West Germany keeps women from labor market participation. Two recent papers exploit policy changes. Felfe and Lalive (2012) deal with child care for toddlers and child development and, in an aside, report some tentative evidence for positive effects of public child care on maternal employment. Finally, Gathmann and Sass (2012) analyze a reform in the East German state of Thuringia that increased prices for public daycare to find negative effects on female labor supply.

[2] If the lack of a private market for child care indicated a lack of demand, for example because society believes that maternal care is more beneficial for young children, this would of course mitigate any positive effects.

Drawing on the rich individual-level data of the German Socio-Economic Panel (SOEP), our instrumental variable estimations yield positive effects of public child care on maternal employment. Intention-to-treat estimates suggest that eligibility for public child care increases mothers' labor supply by 6 percentage points. Second-stage results show that if a mother's youngest child actually attends public child care as a result of the cut-off rule, this mother's probability of being employed increases by roughly 35 percentage points. This also means that about two thirds of these mothers use child care for other reasons than taking up a job.

Investigating the heterogeneity of our first-stage estimates gives us some sense about the validity of our approach and at the same time allows us to characterize the complier subpopulation for which the effects are identified. In particular, we show that the cut-off rules were mostly applied in West Germany during the first years after introduction of the legal claim to a place in kindergarten. The relevance of the cut-off rules becomes weaker with years passed since the reform and is virtually nonexistent in East Germany, where capacity constraints in public child care were not an issue. Further analyses suggest that the first stage tends to be stronger for more educated mothers, mothers with above median age, mothers having older children as well as mothers whose youngest child's age distance to his oldest sibling is above median. To investigate the validity of our instrumental variable approach, we run first-stage regressions in placebo treatment periods before 1996. Further, we perform placebo treatment tests drawing on the panel structure of the SOEP in which we investigate whether our instrument can predict maternal employment in the year preceding actual kindergarten entrance. All these specification tests corroborate the validity of our instrumental variable approach. In further robustness checks, we run piecewise linear regressions and gradually restrict the sample to observations very close to the cut-off. Of course, once we restrict the sample, the number of observations decreases resulting in large standard errors. But still, the point estimates of all these alternative specifications are very similar to those estimated in the standard specification.

In a second identification strategy, we exploit the marked increase in public child care provision for three- and four-year old children in the years following introduction of the legal claim to a place in kindergarten. To this end, we use data from the German Micro Census, Europe's largest household survey. In a difference-in-differences approach, we compare the employment ratios of mothers with three- and four-year old children in 1996 and 2001 with the employment ratios of mothers with older children and, as an alternative, also the employment ratios of women without children. The difference-in-differences estimations confirm the positive causal effects of public child care on maternal employment. To provide evidence that the key identifying assumption of our difference-in-differences approach is met, namely, that the treatment and control group follow the same time trends in the absence of the treatment, we run placebo treatment tests in pre-treatment periods. These placebo treatment tests show that our treatment and control groups do indeed follow the same time trend in the years preceding the actual treatment; thus, they corroborate the key assumption of our difference-in-differences model. We find that the results from the difference-in-differences model using the German Micro Census are remarkably similar to the instrumental variable results using the SOEP, although we identify average treatment effects on the treated (ATT) in the difference-in-differences approach and local average treatment effects (LATE) for the complier subpopulation in the instrumental variables approach. The reason for the fact that both approaches still yield very similar results could be that, in the end, in both identification strategies a very similar group of mothers is affected by the treatment, namely mothers living in West Germany in the late 1990s who get the opportunity to work because their three- or four-year-old child, who is also their youngest child, becomes eligible for public child care in a country where public child care was severely rationed before the reform, no market for private care existed and maternal employment was traditionally low despite the fact that women were very well educated.

The rest of the paper is structured as follows. Section 2 summarizes the empirical literature on public child care and maternal employment. Section 3 describes the pre-reform setting as well as the 1996 public child care reform in Germany and explains how we exploit this reform in two quasi-experimental settings. Section 4 provides information on the rich individual level data sets of the SOEP and the German Micro Census that we use in our analyses. In Section 5, we present instrumental variable results using the SOEP data and corroborate these results applying difference-in-differences techniques on the German Micro Census. Thereafter, we discuss the size of the effects and try to compare the effects to those identified in previous studies. Section 6 concludes.

2 Literature on public child care and maternal employment

Earlier empirical studies analysed the role of child care prices for mothers' employment (see Blau (2003) or Blau and Currie (2006) for reviews). However, there is considerable uncertainty about the magnitude of the effects, mostly due to questionable exclusion restrictions in these non-experimental settings (Blau 2003). More recent empirical papers apply quasi-experimental identification strategies to uncover causal effects of public child care on maternal employment.

The effects of the introduction of free kindergarten for five year olds in US public schools mainly in the 1960s and 1970s are evaluated by Cascio (2009). At that time, enrolment in public school was far from universal but private care was a viable alternative for many families. Cascio (2009) finds strong positive effects of the introduction of free kindergarten in US public schools on the labor supply of single mothers whose youngest child was five, whereas effects for married mothers cannot be detected. Using 1980 US Census data, Gelbach (2002) exploits quarter-of-birth variation in the access to free public preschool for five year olds. In general, Gelbach's (2002) effects tend to be smaller than the effects identified by Cascio (2009). One reason might be that public school attendance and maternal employment rates were already considerably higher in 1980 than in the 1960s and 1970s. Still using more recent data, Fitzpatrick (2010) applies a sharp regression discontinuity design to identify intention-to-treat effects (ITT) of free universal preschool for four year olds in the US states Georgia and Oklahoma. In general, she finds little effects on maternal labor supply. However, strong positive effects are identified for mothers from rural regions. Bassok et al. (2014) suggest that the public Pre-K programs in Georgia and Oklahoma did not crowd-out private child care arrangements. Fitzpatrick (2010) argues that her zero effects might rather be explained by female labor supply elasticities being smaller in recent years than in the past. Furthermore, individuals in her control group might benefit from other child care subsidies such as Head Start.

For Canada, Lefebvre and Merrigan (2008) use difference-in-differences techniques and exploit the stepwise introduction of inexpensive universal child care arrangements in the province of Quebec from 1997 until 2000. Based on Statistics Canada's Survey of Labor and Income Dynamics, the authors find that this program increased the labor market participation rate of mothers with at least one child aged one to five. Baker et al. (2008) analyze the same reform using data from the National Longitudinal Survey of Children and Youth (NLSCY). Their difference-in-differences estimates confirm the Lefebvre and Merrigan's (2008) positive effects.

Berlinski and Galiani (2007) exploit an Argentinian child care reform which provided free public preschool for three to five year olds in the mid to late 1990s and find that this reform increased the labor supply of mothers. In a related paper, Berlinski et al. (2011) confirm these positive effects by exploiting cut-off rules in admission to public child care for four year olds in Argentina. However, there is no effect for mothers whose four year old child is not her youngest child. Turning to an ethnic group of women whose labor market participation was particularly low before a child care reform, Schlosser (2007) analyses the effects of free universal preschool provision for three- to four-year olds on the labor supply of Arab mothers in Israel. Using Data from the Israeli Labor Force Survey from 1998 to 2003, she finds substantial positive effects of free universal child care on educated Arab mothers.

In contrast to Schlosser (2007), Lundin et al. (2008) focus on a country where female employment is traditionally high, namely Sweden. The authors exploit a Swedish child care reform from 2002 which reduced child care prices. Applying difference-in-differences techniques, Lundin et al. (2008) find zero effects of reduced child care prices on maternal employment. Taking into consideration that the child care attendance and maternal employment were at a very high level already before the reform, the zero effects do not come as a complete surprise.[3]

Havnes and Mogstad (2011) analyze the introduction of subsidized universal child care for three- to six-year old children in Norway from 1976 to 1979. Although this paper also focuses on a Scandinavian country, the setting was quite different from Lundin et al. (2008) since child care coverage and maternal employment was still comparably low in the 1970s. Nevertheless, Havnes and Mogstad (2011) find no positive effects of public child care on mothers' employment. Note that even after the expansion in the late 1970s, public child care for three to six year olds was still severely rationed. Therefore, if mothers who were already closely attached to the labor market prior to the expansion received the rationed slots, they might just have substituted informal care for formal public child care. This could explain the zero-effect of public child care on maternal employment in this paper.

Moving to continental Europe, Nollenberger and Rodríguez-Plánas (2011) investigate the effects of a Spanish reform which introduced free public child care for all three year olds in 1991. They exploit variation in public child care across time and states in a difference-in-differences setting and find positive employment effects that are particularly strong and persist for several years for mothers with a high-school degree, whereas no effects are found for college educated mothers.

Taken together, the current empirical literature on public child care and maternal employment suggests that marginal decreases in the costs of available public child care (Lundin et al. 2008) should not affect maternal labor supply if employment rates and child care attendance rates are already high. Further, we should not expect substantial employment effects if newly available public child care simply crowd-out existing private (Havnes and Mogstad 2011) or comparable public child care arrangements (Fitzpatrick 2010). Crowding-out might be particularly relevant if public child care slots are still severely rationed after expansion as in Havnes and Mogstad (2011). It might be that the newly provided public child care slots are then primarily given to mothers already closely attached to the labor market, and that these mothers now just substitute private care with public care arrangements. Moreover, any effects of public child care on female labor

[3] Bettendorf et al. (2012) exploit a similar reform which reduced child care prices on average by 50 percent in the Netherlands in the late 2000s. They find modest intention-to-treat effects on maternal labor supply for a combined treatment of price cuts in public child care and extensions of earned income tax credits.

force participation should be stronger for mothers whose youngest child is the marginal child who gets public care. If there are still younger children around, we do not expect strong effects on the respective mother's labor supply. Furthermore, the lower maternal employment and public child care coverage rates still are in a country where females are generally well educated, the larger the expected effects of providing inexpensive public child care.

3 The Expansion of Public Child Care in Germany as a Quasi-Experiment

In this paper, we exploit the expansion of Germany's public child care system in the late 1990s to identify effects of public child care on mothers' labor market outcomes. To establish causality, we need to overcome several econometric concerns. First, decisions about a mother's labor supply and her child's child care attendance are jointly determined. This simultaneity of mothers' choice about working and sending children to child care hampers causal identification of this relationship in simple OLS models (see, e.g., Coneus et al. 2009). Second, mothers who send their children to kindergarten could be systematically different from mothers who do not and these differences might also influence labor market outcomes. It is not possible to fully account for all these confounding factors by using control variables since some will remain insufficiently observed or completely unobserved. However, special features of the politically induced expansion of the public child care system in Germany in the late 1990s provide us with two quasi-experimental settings that we exploit in order to identify causal effects of public child care on maternal employment.

3.1 Child Care and Maternal Employment in Germany: the Pre-Reform Setting

In our analysis, we focus on the system of German kindergarten, i.e., the publicly provided child care institution children voluntarily attend from age three until school entrance. In contrast to the United States, German kindergarten is not integrated into the school system and is entirely included in the child care system.[4] The latter is historically characterized by large structural differences between West and East Germany (the former German Democratic Republic). In West Germany, child care was historically a private matter and the state was only marginally involved (Rauschenbach 2006). Mothers, grandparents, other relatives, and/or neighbors were the people chiefly responsible for the care of children from birth to school entrance. Most often, it was the mothers themselves who looked after their children (Tietze and Rossbach 1991). This stands in stark contrast to the situation in East Germany. Based on the socialist idea that

[4] To simplify matters, we will use *child care* and *kindergarten* as synonyms throughout the paper. This is unproblematic as we focus on only the care aspect of kindergarten. We are not interested in effects of possible skill accumulation during kindergarten or other quality aspects that would make the use of the notion *preschool* more consistent.

education and care of children should be controlled by the state from the very beginning, the former German Democratic Republic (GDR) created an elaborate child care system offering full provision of institutional care for all children between one and six years old. This comprehensive child care system largely survived the reunification of East and West Germany and thus there continued to be substantial differences in public child care between the two parts of the country in the early 1990s.

After reunification in 1990, West Germany provided places in kindergarten to about 67 percent of children aged three to six, whereas there was full provision in the East German states (Rauschenbach 2006). There is substantial heterogeneity in child care attendance across children's ages in West Germany. Data from the German Micro Census suggest that, in 1995, 90 percent of all five and six year olds in West Germany attended kindergarten, whereas this number is only 60 percent for the four year olds and 30 percent for the three year olds. Further, public child care for under three year old children does virtually not exist in West Germany in the 1990s. Indeed, expansion of public child care for under three year olds did not start before the mid-2000s.

If we speak about kindergarten in West Germany in the early 1990s, we usually mean half-day public child care for three to six year olds. Full-day public child care was scarce; and even if a public child care institution offers to care for the children in the morning and in the afternoon, this mostly means that there is no care provided during lunchtime. For example, in 1989, only 14 percent of all kindergarten slots were full-time slots, and 62 percent of those did not offer care during lunchtime (Spieß 1998). Kindergartens are closed on weekends and during longer school holidays. Note that in West Germany, although kindergartens are not officially part of the educational system, they are often used for educational reasons even if the mother has no job.

We call child care institutions "public" if they receive public subsidies; subsidies can be received by public-sector as well as voluntary youth welfare organisations. In 1990, 30 percent of all child care institutions for three to six year olds in West Germany were under the organization of public-sector authorities, i.e., municipalities, and 69 percent under the organization of youth welfare organisations. This means that 99 percent of all child care institutions in West Germany were publicly funded in 1990. Accordingly, there was virtually no private market for child care despite the fact that provision of child care was far from universal (Spieß 1998). At given subsidy levels, demand for child care greatly outreached supply resulting in waiting lists at public child care institutions. Due to the heavy subsidization of public child care and numerous regulations about the quality of staff, child-staff-ratios, construction norms, hygiene and health precautions, and partly curricula, it is not surprising that no private market had emerged. Therefore, parents mainly had to rely on informal child care as a supplement to public child care; indeed, on weekdays, about 30 percent of all children aged three to seven were (also) cared for by friends and relatives (in particular by grandparents) in West Germany in 1991 (Spieß 1998).

West Germany has long been known for its low female, and in particular maternal labor force participation. Data from the German Micro Census suggest that, in 1995, only 45 percent of all mothers whose youngest child was between three and four were employed. The employment rate of mothers whose youngest child was ten or eleven was slightly below 65 percent, while the employment rate for women without any children aged 29 to 36 was slightly above 85 percent. Thus, the employment gap between mothers of three to four year old children and women of a similar age but without any children amounted to 40 percentage points. Also note that two thirds of employment of mothers of three to four year olds is only part-time employment.

A further institutional detail, which might be worth mentioning in this context, is Germany's generous maternity leave regulation. After a series of expansions in leave coverage starting in the 1970s and continuing through the 1980s, mothers could benefit from job-protected leave for 36 months after birth and maternity leave payments for 24 months after birth from January 1993. Job-protected leave means that the employer must not dismiss a mother from the day she knows of the woman's pregnancy until 36 months after childbirth. While mothers received their average net pre-birth-income for the first two months after childbirth (or 600 Deutschmarks in case the mother did not work prior birth), maternity leave payments amounted to 600 Deutschmarks from the third month until the twenty-fourth month. In 1986, 600 Deutschmarks corresponded to roughly 25 percent of average pre-birth earnings. Ludsteck and Schönberg (2012) find that expansions in leave coverage had negative effects on labor market outcomes of mothers in the short run. However, they also show that these expansions did not affect long-run maternal labor supply measured six years after childbirth.

3.2 The Introduction of the Legal Claim to a Place in Kindergarten

The mid to late 1990s witnessed substantial political effort toward increasing public child care coverage in West Germany. One of the crucial reasons behind this effort was the increase in the number of well-qualified women who wanted to participate in the labor market. Women's increasing desire (or need) for labor force participation meant they were no longer as available to care for children before they reach school entrance age. Thus, the trade-off between family or career being faced by young mothers revealed the insufficiency of the child care system and consequent economic and social problems. Contemporaneously, people started becoming aware of the importance of the early life period for child development, leading to an even stronger focus on better public care provision.

The crucial step toward full provision of public child care for three- to six-year olds in West Germany was the government's introduction of a legal claim to a place in kindergarten (Rechtsanspruch auf einen Kindergartenplatz), which became effective on 1st January 1996. Since that year, every child has been eligible to attend a center-based kindergarten from his or her third birthday until school entrance.[5] The legal claim is found in §24 SGB VIII (Achtes Buch Sozialgesetzbuch) and covers highly subsidized half-day care.[6] On average, roughly 80 percent of the costs of a place in public child care were covered by subsidies, with the remaining 20 percent financed by parents' child care fees (Statistisches Bundesamt 2004). Altogether, the German government spent 10.5 billion Euros on public child care in 2002, 10 billion of which were spent on operating expenses and the remaining 0.5 billion on investments. Note that the official statistics only include government spending aggregated over all forms of public child care; however, kindergarten care was by far the most prominent expense at that time. While providing half-day care was the minimum legal requirement, local authorities were generally requested to adjust their public child care supply according to the parents' needs. To meet the increased demand for public child care, a substantial and rapid expansion of center-based care places became necessary.

Due to the complicated organizational system in Germany, implementation of the legal claim to a place in kindergarten turned out to be a difficult process. The new law had been passed at the federal level, but it was the states and, especially, the municipalities that were responsible for organizational and financial implementation of it. Decision makers at the municipal level were confronted with completely new responsibilities and were unable to quickly meet the new and officially mandated demand for child care supply. As a consequence, by the end of 1995, the German Federal Parliament (Deutscher Bundestag) adopted a legislative initiative proposed by the Federal Council of Germany (Bundesrat) that introduced cut-off rules: municipalities could decide that children will not necessarily become eligible to attend kindergarten right at the day of their third birthday, but will instead become eligible at the start of the "kindergarten year" following their third birthday. The start of a "kindergarten year" typically coincides with the start

[5] The claim was originally formulated as a social measure accompanying a new abortion regulation in Germany and intended to encourage mothers to carry children to term. Moreover, as pointed out by an anonymous referee, the wording "legal claim" might have the non-neutral connotation that kindergarten attendance is something very valuable, which could induce demand for public child care beyond what we would see if it just became available through other means.
[6] As a comparison, government spending on universities amounted to 19.5 billion Euros in 2001.

of a school year, which is in August or September depending on the federal state and the year.[7] So, children turning three slightly after the cut-off date, i.e., the date of the start of the "kindergarten year", had, at worst, to wait almost a year to actually become eligible to attend kindergarten. This rule relieved some of the pressure on municipalities as it gave them some extra time to provide the child care supply required by the new law.

The cut-off rule was intended as a temporary measure and the period of its applicability should have ended at the end of 1998; however, the use of cut-off dates for child care entrance continued to be common practice until quite recently. When looking at child care attendance rates for the different age cohorts from three to six, we still find the lowest rates for three-year-old children (Konsortium Bildungsberichterstattung 2006, p. 38). According to the 12th Kinder- und Jugendbericht (BMFSFJ 2005, p. 298), this is partly attributable to supply-side shortages and the ongoing use of cut-off rules in several municipalities.

The cut-off rule was not enforced universally as some municipalities had child care supply sufficient to provide for all children aged three up to the age school entrance and did not need the rule. Moreover, some families might have been able, for various reasons, to place their children in child care at the age of three or even earlier even though the children, to be in compliance with the rule, should have been older when entering. Other mothers preferred to care for their children themselves and therefore decided to delay kindergarten entrance until their children were older, completely irrespective of the rule.

Eventually, the introduction of the legal claim to a place in kindergarten induced an expansion in public child care at the extensive but also at the intensive margin. By 2002, about 25 percent of all kindergarten slots in West Germany were full-time slots offering lunch and an additional 50 percent of all kindergarten slots were full-time slots without lunch (Riedel 2005, p.132).

3.3 Empirical Strategies

Instrumental Variable Approach

Unfortunately, there are no data on which municipalities adopted cut-off rules in which years. However, the mere existence of these cut-off rules provides us with a quasi-experimental setting that can be exploited in an instrumental variable approach. This is because the cut-off rules create variation in child care attendance that is reasonably exogenous to confounding factors determining child care and mothers' labor market participation simultaneously. We exploit these

[7] Although the start of the kindergarten year is not completely restricted to a specific date and generally possible in every month, most children enter kindergarten at the start of the kindergarten year.

cut.off rules and instrument actual child care attendance D_i with the eligibility Z_i at the last kindergarten start. Our quarter-of-birth style instrument Z_i is a dummy variable indicating whether a mother's youngest child was above cut-off age (>=36 months) at the last kindergarten start and was therefore eligible for child care; it is zero if a mother's youngest child was below cut-off age at the last kindergarten start. This leads to the following model predicting actual public child care attendance:

$$D_i = \mu + \delta Z_i + \phi X_i + \chi_i \qquad (1)$$

The coefficient δ on Z_i shows the relevance of our instrument and indicates the share of children who enter kindergarten in compliance with the cut-off rule; X is a vector of covariates including the youngest child's age (in months) and χ is the error term.[8] We use the predicted values $\hat{D}_i$ from Equation (1) in Equation (2) to obtain a local average treatment effect (LATE) τ of public child care attendance of the youngest child on mothers' labor market outcomes, while X is a vector of covariates including the youngest child's age (in months) and π is the error term:

$$Y_i = \eta + \tau \hat{D}_i + \nu X_i + \pi_i \qquad (2)$$

In other words, the cut-off rule creates a discontinuous jump in kindergarten attendance at the age threshold of 36 months. We use this exogenous variation in public child care attendance D_i that is due to the age cut-off rule Z_i in order to estimate the effect of kindergarten attendance on maternal employment. Note that in this main specification, we restrict the youngest child's age to have the same effect to the left and to the right of the cut-off. Later, we will also consider relaxing the assumption that the slopes are equal on both sides of the cut-off.

In order for τ to yield the local average treatment effect for the complier subpopulation, our quarter-of-birth-style instrument Z_i (cf. Angrist and Krueger 1991) must meet the exclusion restriction. This means that the cut-off rule must not affect maternal employment directly nor via any other channel but kindergarten attendance. Problems might arise if the age of the youngest child has a direct effect on maternal employment or an indirect effect on maternal employment running via channels other than child care. For example, there could be an indirect effect if informal caregivers such as grandparents, friends, or neighbors can more easily care for slightly older children (just above three years old) than for younger children (just below three years olds), which, at the same time, could make maternal employment more likely. Further, mothers'

[8] In all our IV regressions, standard errors are clustered at the individual mother level; clustering the standard errors by the child's age does hardly change the results; if at all, standard errors slightly decrease.

attitudes toward career and family life could change with the youngest child's age. In this case, the exclusion restriction in a model that does not control for age would be violated if, as the youngest child grows older, mothers are more likely to opt for child care and simultaneously become employed. This is why we explicitly include the age of the youngest child (in months) as a covariate in our instrumental variable approach. To check the robustness of our results, we not only control for the youngest child's age (in months) linearly but also include the square term of the child's age and allow for heterogeneous effects of the youngest child's age to the left and to the right of the cut-off. In addition to the child's age, we can control for several demographic and socioeconomic characteristics of the mother (age, highest educational degree, migration status) and of her partner (age, highest educational degree, migration status, employment status, labor income). Likewise, we include information about the household size, the youngest child and his or her siblings (youngest child's gender, number of siblings, and distance in months to oldest sibling), as well as state and year dummies. Further note that in all regressions, we restrict our sample in a way which ensures that maternity leave has already expired for all mothers once we observe them.

Difference-in-Differences Approach

Introduction of the legal claim to a place in kindergarten led to a substantial expansion of child care facilities, particularly in West Germany, in the late 1990s and early 2000s. Figure 1 shows the development of child care attendance in West Germany from 1991 to 2003.[9] Child care attendance by five- and six-year old children was already at the high level of slightly under 90 percent in 1991 and remained relatively stable throughout the following years. Three- and four-year old children, in contrast, were substantially less likely to attend child care in the early 1990s. However, from 1996 to 2001, child care attendance by three- and four-year-old children increased sharply. We exploit the considerable rise in public child care attendance by three- and four-year-old children in West Germany in our second identification strategy by applying difference-in-differences techniques.

Figure 1 about here

The fact that the public child care reform was at the national level does not allow us to exploit state specific variation in the timing of the reform, which would probably be the most convincing strategy. However, we can still exploit temporal variation across child age groups. To this end, we define mothers whose youngest child is three or four years old as our treatment group. As a

[9] We observe few dynamics and universally high attendance in child care by all children between the ages of three and six in East Germany, as can be observed in Figure A.1 of the Online Appendix.

control group, we use mothers whose youngest child is aged 10 or 11 years.[10] Since these older children were attending school during the whole period of observation, they were not affected by the legal claim to a place in kindergarten. There was also no school reform (e.g., an expansion of all-day schools) or any change in the provision of after-school care that might have influenced the time these older children were exposed to public care arrangements (including schools). As an alternative control group we use women without any children under the age of 18. We identify the treatment effect by computing the difference in the employment trend from 1996 to 2001 between mothers with three- to four-year olds and mothers with 10- to 11-year olds (or women without children) in West Germany. The treatment effect θ can be expressed by the following double difference equation:

$$\begin{aligned}&\theta = E[Y_i(1)] - E[Y_i(0)] = \\ &\{E[Y_i|T_i = 1, D_i = 1] - E[Y_i|T_i = 0, D_i = 1]\} - \\ &\{E[Y_i|T_i = 1, D_i = 0] - E[Y_i|T_i = 0, D_i = 0]\},\end{aligned} \tag{3}$$

where $E[\cdot]$ is the expectation operator and Y_i is the labor market outcome of mother i. $D_i = 0$ for the control group of mothers with the youngest child aged between 10 and 11 years, and $D_i = 1$ for the treatment group of mothers with three- to four-year-old children. $T_i = 0$ for the year 1996, and $T_i = 1$ for the year 2001.[11] In a multivariate setting, the treatment effect θ is identified as the coefficient on the interaction term between D and T in the following regression:

$$Y_{it} = \sigma + \omega T_t + \vartheta D_i + \theta (T_t D_i) + \varepsilon_{it}, \tag{4}$$

ω captures time effects common to the treatment and control group, ϑ captures baseline differences between the treatment and control group, and ε_{it} is the error term. The inclusion of further covariates in this model is straightforward. Instead of merely looking at two years, 1996 and 2001, we additionally estimate an extended version of Equation (4) to illustrate the relative development of maternal employment over the whole period from 1991 to 2001. However, note that robust standard errors may be downward biased particularly in this setting where we use data from many time periods (Bertrand et al. 2004).

[10] This is because ten to eleven year old children form the youngest possible pair of ages which is not contaminated by the reform and can therefore be considered a control group. Using mothers with five- or six-year-old children as a control group would lead to spurious results since this group could benefit from the legal claim to a place in kindergarten. For example, a child who was five years old in 2001 was three years old in 1999 and thus part of the treatment group. This boils down to a violation of the standard stable unit treatment value assumption (SUTVA) (Rubin 1978).

[11] Due to a conspicuous drop in child care attendance of four year olds in 1995 (see Figure 1), we decided not to use the year 1995 as the baseline for the main specifications. However, we checked that our results are not sensitive to using 1995 or 1993 as an alternative baseline year

The key identifying assumption of this difference-in-differences estimation strategy is that the treatment and control group follow the same employment trends in absence of the treatment. In other words, there are no unobserved variables that change over time and result in differential effects on the labor market outcomes of mothers with children aged three or four as compared to mothers with children aged 10 or 11 (or women without children). A simple way to partially test the plausibility of this assumption is a placebo treatment test in the years preceding the actual treatment that can show deviations from the common trend in pre-treatment years. In the post-treatment years, the only reason why the trend of the treatment group departs from the trend of the control group must be the treatment. This assumption would obviously be violated if there was another policy reform which affected treatment and control group differently. To the best of our knowledge, no such policy occurred in the post-treatment years. Still, one might argue that maternity leave expansions in the 1980s and early 1990s (the last one was in January 1993) could be a threat to identification. In particular, the control group of mothers of ten to eleven year olds in 2001 had benefitted from more generous maternity leave coverage than mothers of ten to eleven year olds in 1996. If extended maternity leave had negatively effects on mothers' labor supply, this would negatively affect the trend in the control group and thus bias our difference-in-differences estimates upwards. However, Ludsteck and Schönberg (2012) find that maternity leave expansions did not affect mothers' long-run labor outcomes, measured six years after childbirth. Thus, it seems justified to assume that maternity leave expansions did also not affect mothers' labor outcomes measured ten to eleven years after childbirth.[12]

4 Data on Child Care Attendance and Maternal Employment

To identify the causal effect of child care attendance on maternal labor supply, we draw on two different individual-level data sets. For our instrumental variable approach, we use the German Socio-Economic Panel (SOEP); the difference-in-differences estimations are based on data from the German Micro Census, Europe's largest household survey.

4.1 The German Socio Economic Panel (SOEP)

The SOEP is a representative and longitudinal household survey that provides annual information about households and their members on a wide range of socioeconomic and demographic characteristics. The survey started in West Germany in 1984; since 1991, it has been

[12] Furthermore, we do not find any evidence for fertility effects of the child care reform, which would lead to selective sampling of the control group. Indeed, mothers in our sample whose children are eligible for kindergarten do not differ by the number of children observed in their last available SOEP interview from mothers whose children are not eligible for kindergarten.

extended to East German households. The SOEP allows us to match mothers with their children as well as their partners. This feature enables us to characterize families along many socioeconomic and demographic dimensions.

For our instrumental variable approach, it is especially important to have information about child care attendance and, most crucially, about the month of birth of the children. Since we know the month of the interview, we can compute the age of a child in months at the time of the interview and also at the time of the start of the last kindergarten year. We use two outcome variables measuring maternal employment. The first is a dichotomous variable that equals 1 if the mother is working; 0 otherwise. The second is a variable that measures a mother's weekly working hours. We include an extensive set of covariates in our model to capture various aspects of the household's socioeconomic situation. Specifically, we use information on the mothers' age, years of education, and migration background, as well as information about their partner's age, years of education, migration background, occupational status, and labor income. We also know whether the family owns or rents the flat/house in which they live and in which state of Germany this flat/house is located. Furthermore, we use variables that characterize the children, such as age, gender, number of siblings, and age of siblings. Finally, we include year dummies to capture year-specific effects.

Figure 2 about here

In Germany, the beginning of a school year lies between August and September, slightly varies by state and year, and, most importantly, coincides with the normal start of a kindergarten year. We have exact information on the date school (and therefore kindergarten) begins in each German state for every year of our period of observation and can link this information to our household data. Since there are no strict rules preventing children from starting kindergarten also in other months, there is some uncertainty with respect to the month a child actually starts kindergarten. However, Schlotter (2011) uses data from the Children's Panel of the German Youth Institute (DJI-Kinderpanel) to show that almost two-thirds of all children report actual kindergarten entrance in August or September. Entrance in other months is almost equally distributed at a low level of about 5 percent (see Figure 2). Using the state-year specific information on the normal start of a kindergarten year, we compute the age in months at the last possible kindergarten start, both for children who attend kindergarten and for those who do not.

By using the SOEP waves from 1991 until 2005 in our estimations, we can not only analyze an extensive period after introduction of the legal claim to a place in kindergarten, but also several years before this reform, which should give us a comprehensive picture of the effects of this reform on maternal employment. For our main estimations, we limit our sample to mothers whose youngest child was born between 1992 and 2000 as these cohorts are the ones most likely affected by the cut-off rules introduced in 1996. Furthermore, we only include mothers whose youngest child, born between 1992 and 2000, was at least 36 months old at the time of the interview and at most 48 months old at the last possible kindergarten start. Thus, we focus precisely on the group of children whose child care attendance was dependent on the existence of a cut-off date. Moreover, note that this restriction ensures that maternity leave has already expired for all mothers in our sample. Furthermore, excluding mothers with children younger than 36 months seems reasonable since the presence in the household of a very young sibling might prevent mothers from entering the labor market, irrespective of the child care attendance of their older child. As most interviews take place in spring and therefore roughly half a year after kindergarten start in August/September, we allow for an interval long enough to give the mother a chance to search for and find a suitable job. The full sample consists of 1,228 mothers and 1,936 mother-child observations.

4.2 The German Micro Census

We also use data from the German Micro Census, the largest household survey in Europe. This data set provides annual information on a representative 1 percent sample of the German population with a focus on demographic and labor market relevant variables.[13] Entire households are sampled, enabling us to observe both mothers and their children. We use the Micro Census waves 1991, 1993, 1995, and 1996 until 2001.[14] In these years, all variables are measured in one specific week in April. The census data do not provide information as detailed as that we obtain from the SOEP; in particular, the children's months of birth are not observed. Yet, the larger number of observations of the Micro Census and the availability of variables such as a mother's employment status and child care attendance of the children (as well as several additional covariates), which can be observed in all waves and are comparable over time, make this data set a complement to the SOEP data. In each of the Micro Census cross-sections, we focus on households with mothers whose youngest child is three or four years old, since this is the group of families affected by introduction of the legal claim to a place in kindergarten in 1996 and the massive expansion of public child care in subsequent years. As with the SOEP, using only families whose youngest child is three or four years old ensures that mothers' employment outcomes are not influenced by the presence in the household of an even younger child. As a

[13] The Scientific Use Files that we can use are a 70 percent subsample of the original 1 percent Micro Census sample.
[14] Unfortunately, the Scientific Use Files for the Micro Census waves 1992 and 1994 are not available and we have to use the incomplete times series of 1991, 1993, and 1995 to study pre-reform trends.

control group we use mothers whose youngest children are 10 or 11 years old and thus do not benefit from the increased supply of public child care in the years after 1996. Alternatively, we use women without children as our control group.

As key outcome variable, we use a woman's employment, where the definition of employment is quite broad and includes vocational training, marginal employment, jobs to supplement welfare benefits, and the like. In addition, the data set contains a variable that assesses child care attendance by the children (reported by the head of the household) at the time of the interview.[15] In our estimations we can also include several covariates, such as the woman's age, highest school degree, and nationality. Information on state of residence of the household and on family type facilitates separate analyses for East and West Germany, as well as for single mothers and mothers with partners, respectively.

5 Results

In what follows, we provide evidence on the effects of public child care on mothers' labor market outcomes from the instrumental variable approach based on the SOEP and the difference-in-differences approach based on the Micro Census.

5.1 Results from an Instrumental Variable Approach

We start our instrumental variable approach by plotting kindergarten attendance rates by the child's age (in months) at the start of the last kindergarten year. Note that kindergarten attendance is measured in spring, i.e., roughly half a year after the kindergarten year started. The upper panel of Figure 3 shows that, in general, the older the child at the start of the last kindergarten year, the more likely she is to attend kindergarten roughly half a year later. Our instrumental variable approach exploits the fact that kindergarten attendance sharply increases at the cut-off age of 36 months. Whereas 62 percent of the children aged 35 months at the start of the last kindergarten year attend kindergarten half a year later, this number jumps to 79 percent for children aged 36 months at the start of the last kindergarten year. This 17 percentage point difference between children who differ by just one month in age corresponds to a substantial increase of kindergarten attendance by more than 27 percent.

[15] The Micro Census question only asks whether the child attends public child care or not and thus does not allow us to differentiate by public child care facility for different age groups (Kinderkrippe for those aged zero to three years, Kindergarten for three- to six-year old children, and Kinderhort for after school care for students). However, since we know the age of the children, we can clearly identify child care attendance by children aged three and four as kindergarten attendance, which is the child care facility we are interested in due to its expansion in 1996 and subsequent years.

Now, we predict kindergarten attendance in a linear probability model using all background characteristics, which are also listed in Table 1, as explanatory variables. The lower panel of Figure 3 uses the results of this model and plots predicted kindergarten attendance against the child's age (in months) at the start of the last kindergarten year. While we do observe a discontinuous jump in *actual* kindergarten attendance at the 36 months cut-off, we do not observe a jump in *predicted* kindergarten attendance. Indeed, predicted kindergarten attendance runs absolutely smoothly through the threshold. Thus, this exercise shows that the jump in actual kindergarten attendance at the cut-off age of 36 months cannot be explained by a discontinuous change in observable characteristics that might predict child care attendance. Rather, the jump in actual public child attendance seems to be driven by the exogenous cut-off rule shock.

Figure 3 about here

To further investigate this finding, Column 1 of Table 1 sets out descriptive statistics of mothers whose youngest child was 36 months or older at the time of the last kindergarten start (in August or September) and thus eligible for child care due to the cut-off rules. We compare these statistics to those of mothers whose youngest child was slightly younger than 36 months at the time of the last kindergarten start and thus not eligible for child care (Column 2 of Table 1). As seen from Table 1, the groups are not statistically different in terms of mothers' characteristics (such as years of schooling, migration background), partners' characteristics (years of schooling, employment status, labor income, migration background), and other child and family characteristics (youngest child's gender, number of siblings, child's distance to oldest sibling in months, number of household members). This provides support for the validity of our identification strategy, which relies on the assumption that the groups systematically differ in eligibility for child care but not in other socioeconomic characteristics. Virtually the only variables where we find, by mere construction, significant differences between the groups are the age variables, i.e., mothers' age, partners' age, and age of the youngest child. To avoid the possibility that pure maturity effects are driving our results, we control for these age variables in our 2SLS estimations. Further, we also include all the other variables reported in Table 1 as covariates in our multivariate instrumental variable framework.

Table 1 about here

Turning to a multivariate setting, we start by estimating first-stage regressions predicting child care attendance D of the youngest child at the time of the interview by the cut-off rule Z and a vector of covariates X (see Equation (1)). Column 1 of Table 2 shows the results of a regression of kindergarten attendance on the cut-off rule instrument where we only control for the youngest child's age (in months). We can confirm that our instrument is indeed a strong and relevant predictor of actual child care attendance. Eligibility due to the cut-off rule increases the probability of child care attendance of the youngest child by 17.8 percentage points. This first stage effect turns out to be virtually unaffected by the inclusion of all the individual level controls presented in Table 1 as well as state and year fixed effects (Column 2 of Table 2).

During its socialist years, East Germany developed one of the most advanced child care systems in the world, which, by and large, continued to exist after German reunification. As a consequence, in East Germany, child care was already available for very young children (i.e., those far below the age of three) in the late 1990s. In West Germany, however, child care for children below the age of three was virtually non-existent and child care for children aged three to six was still scarce during our period of observation. Therefore, we expect the age cut-off rule to be more relevant for West German than for East German mothers. Running the first-stage regressions separately on the subsamples of East and West German mothers, we see that this is, indeed, the case (Columns 3 and 4 of Table 2).

The age cut-off rule was most often applied during the first years after introduction of the legal claim when there was often a severe shortage of child care places for three- to six-year-old children. Despite the fact that the period of its applicability should have ended at the end of 1998, cut-off dates continued to be used in many municipalities until quite recently (BMFSFJ 2005), even though supply-side shortages are no longer much of an issue. We split the sample into two subsamples, one of them covering the years from 1996 to 1998, the other covering the years from 1999 to 2005. We expect the first stage to be stronger in the first years after the reform, and weaker, though still present, in the later years. Indeed, as we observe in Columns 5 and 6 of Table 2, the coefficient of the cut-off rule instrument is substantially larger in the years from 1996 to 1998 than in the years from 1999 to 2005.[16] Since a large part of the identifying variation in public child care attendance comes from West Germany in the first years after the legal claim to a place in kindergarten was introduced, we should interpret our IV results as short-run effects of public child care in the years following the introduction of the legal claim in West Germany.

Table 2 about here

Our instrument is designed to reveal effects of the cut-off rules applied in the aftermath of introducing the legal claim to a place in kindergarten in 1996. Therefore, we should not see any

[16] Slightly varying the cut-off year does not affect the general pattern that the first stage coefficient is larger in the earlier than in the later years.

correlation between our instrument and child care attendance before 1996. We test this by estimating our first-stage regression (Equation (1)) on a sample of mothers whose youngest child was slightly below or slightly above the cut-off age of 36 months at the start of the last kindergarten year in the years 1991 to 1995. As expected, Column 7 of Table 2 shows that our instrument does not predict actual child care attendance in these years.

In further first stage regressions, which are available from the authors upon request, we analyze heterogeneity in the first stage effects along several characteristics of the mother, her partner and the family. This should give us a more elaborated picture of the complier subpopulation for which our instrumental variable estimates identify treatment effects. In particular, we run the first stage regressions on specific subsamples and compare the size of the first stage coefficients to each other. It turns out that mothers with above median years of schooling, mothers with above median age, mothers also having older children as well as mothers whose youngest child's age distance to his oldest sibling is above median react more strongly to the eligibility cut-off.

Let us now turn to our main outcome variable, which is maternal employment measured roughly half a year after kindergarten start. The upper panel of Figure 4 plots actual maternal employment rates against the age of the youngest child (in months) at the last kindergarten start. Thus, Figure 4 boils down to a graphical representation of the intention to treat effects (ITT). We can see that, in general, the older the child at the start of the last kindergarten year, the more likely it is that her mother is employed half a year later. Most importantly, however, we observe a jump in employment at the cut-off age of 36 months. While 46 percent of all mothers whose youngest child was 35 months at the start of the last kindergarten were employed half a year later, this number jumps to 59 percent for mothers whose youngest child was 36 months old at the start of the last kindergarten year. In addition, although the pattern is rather noisy, it is suggestive of a more permanent change in employment, controlling for a trend in age, as we cross the cut-off; this is in line with the permanent change in child care attendance depicted in Figure 3. To make sure that the employment jump is not driven by a discontinuous change in employment relevant characteristics at the cut-off, the lower panel of Figure 4 plots predicted maternal employment against the youngest child's age in months at the start of the last kindergarten year. The fact that we do not observe a jump in predicted employment at the cut-off age of 36 months indicates that the jump in actual employment is not driven by a discontinuous change in the composition of mothers. Figure 5 repeats this exercise for actual and predicted weekly working hours. We get the very same pattern of results; the discontinuous jump at the threshold can be seen very clearly for *actual* working hours but not for *predicted* working hours. Moreover, there appears to be a more permanent change in *actual* working hours, which is again consistent with the permanent change in child care attendance depicted in Figure 3.

Figure 4 about here

Figure 5 about here

We now estimate ITT effects from a reduced form regression, where we regress a mother's labor market outcomes Y on our age cut-off instrument Z to identify the effect of child care eligibility on maternal employment. In addition to these ITT estimates, we also present results from 2SLS regressions where we use only the exogenous part in the variation of child care attendance that is determined by our age cut-off rule instrument (see Equation (2)). This allows us to identify the causal effect of the youngest child attending child care on maternal labor supply Y for the complier subpopulation. In all regressions, we first only control for the youngest child's age in months. In a second step, we then include the full set of control variables X (mother's age, years of schooling, and migration background; partner's age, years of schooling, migration background, employment status, and net labor income, the size of the household, the youngest child's age and gender, number of siblings, and distance (in months) to his or her oldest sibling; as well as state and year dummies).

Table 3 reports results for our dichotomous outcome variable *maternal employment*. Columns 1 through 3 present the estimates where we only control for the youngest child's age in months, while we include the full set of control variables X in Columns 4 through 6. The reduced-form estimate in Column 1 shows that eligibility for child care increases a mother's probability of being employed by 6.4 percentage points. If we include state and year fixed effects as well as all our individual level controls, the ITT estimate remains virtually unaffected: eligibility for child care increases a mother's probability of being employed by 6.5 percentage points (Column 4 of Table 3). The first-stage estimates in Columns 2 and 4 of Table 3 have already been presented in Table 2. As we can see again, these estimates are unaffected by the inclusion of the full set of control variables. Moreover, the results confirm that our cut-off rule instrument is highly relevant and does not suffer from weak instrument problems (with robust F test statistics of 30.2 and 37.5, respectively). Since the reduced form as well as the first stage estimates are unaffected by the inclusion of control variables, it is not surprising to see that so are the second stage estimates. Column 6 of Table 3 suggests that child care attendance by the *youngest* child causally increases the probability of the mother being employed by 36.6 percentage points. Assuming linearity, this means that a 10 percentage point increase in public child care attendance of three to four year olds increases employment of mothers whose youngest child is three to four by 3.7 percentage points. On the other hand, this also means that roughly two thirds of these mothers use child care for other reasons than taking up a job.

In Table 4, we run the same exercises for our alternative outcome variable *weekly working hours*. Again, Columns 1 through 3 present the estimates where we only control for the youngest child's age in months, whereas we include the full set of control variables X in Columns 4 through 6. Once more, the inclusion of the full set of control variables leaves the ITT and 2SLS estimates virtually unaffected. From Column 4 of Table 3 we can see that eligibility for child care results in an average increase of 2.5 working hours per week. The first-stage estimate from Columns 5 of Table 4 again proves that our age cut-off instrument is highly relevant and does not suffer from weak instrument problems (with a robust F test statistic of 32.4), while the second-stage estimates show that child care attendance by the youngest child causally increases working hours per week by 14.3 (Column 6 of Table 4). Thus, our multivariate 2SLS estimates yield univocally positive effects of public child care on maternal employment.[17]

In order to get still more feeling for the size of the estimated effects, we once again turn to the ITT estimates from Column 4 of Table 3 and Column 4 of Table 4. Relating these coefficients to the mean value of the respective outcome variable for the group of mothers having the instrument switched off, we find that the intention-to-treat effects correspond to an increase of maternal employment by 14.1 percent and an increase of working hours by 23.2 percent. These results suggest that public child care has not only positive effects at the extensive but also at the intensive margin. An additional ordered logit reduced form regression corroborates this result. We observe significant reductions in non-employment and a strong increase in part-time work, while the significant increase in full-time work is about half the size of the part-time effect. For detailed results, see Table A.2 of the Online Appendix.

Remember that the legal claim to a place in kindergarten refers to half-day public child care only; however, in the end, the reform induced an expansion in public child care not only at the extensive but also at the intensive margin to meet the parents' needs. Thus, the effect on full-time employment might partly be traced back to the increased availability of full-time public child care. Another complementing explanation for this result might be that informal child care by friends and relatives might play a supplementary role and is not crowded out by the availability of public child care. Similar arguments have been raised by Gathmann and Sass (2012) for East Germany.

[17] In additional estimations presented in Table A.1 of the Online Appendix, we included mothers with partners as well as single mothers in our sample, dropped all partner covariates, and added a dummy indicating whether or not a mother is a single mother. The results are not different from those presented in Table 3 and Table 4.

Table 3 about here

Table 4 about here

The panel structure of the SOEP allows us to perform yet another specification test. If the coefficients estimated in Table 3 and Table 4 do indeed reveal the causal effect of public child care in kindergarten, we should not see any effects of public child care in period *t* on maternal employment in period *t-1*. To test this prediction, we run our reduced-form and 2SLS models with a minor modification, namely, that we now use a mother's employment status in *t-1* and weekly working hours in *t-1*, respectively, as our outcome variables. The results of this placebo treatment test are shown in Table 5. As expected, we find no significant effects of public child care in *t* on maternal employment in *t-1*. Indeed, the point estimates are not only far from being statistically significant, they are slightly positive in the employment equation and slightly negative in the weekly working hours equation, albeit always close to zero. This shows that mothers whose youngest child turned three shortly after the cut-off date are not different from mothers whose youngest child turned three shortly before the cut-off date in terms of employment before the cut-off rule becomes actually relevant. Thus, this placebo treatment test corroborates the validity of our instrumental variable approach.[18]

Table 5 about here

We have also run piecewise linear regressions, which allow the association between the youngest child's age at kindergarten start and the outcome variable to differ on both sides of the cut-off. While the slopes are different in the first stage, they do not statistically differ from each other in the second stage. We find that the child care point estimate is virtually unchanged in this specification but less precisely measured. Moreover, in further 2SLS regressions, we gradually restrict the sample to observations close to the eligibility cut-off. In particular, we first only use observations of mothers whose youngest child is between 30 and 42 months old at the start of the last kindergarten year. Then, we further restrict the sample to mothers whose youngest child is between 32 and 40 months old at the start of the last kindergarten year. And finally, we even narrow down the sample to mothers whose youngest child was between 34 and 38 months old at the start of the last kindergarten year. Of course, observation numbers substantially decrease in these exercises where we move closer to the cut-off. Accordingly, also the precision of the estimates clearly deteriorates. However, it is reassuring to see that the point estimates still depict large and positive effects of public child care attendance on maternal employment; indeed, the

[18] At the same time, however, large standard errors do not make these estimates statistically distinguishable from the main estimates.

point estimates fluctuate around the coefficient estimated in our main specification for the full sample. For detailed results of these regressions, see Table A.3 of the Online Appendix.[19]

Taken together, our 2SLS estimates suggest that child care attendance has significantly positive causal effects on mothers' employment probability and on their weekly working hours. Thus, public child care is not a mere substitute for other informal child care arrangements in Germany. This result is also supported by additional instrumental variables estimations which show that public child care indeed reduces the mother's time used for child care by roughly four hours a day.[20] However, even if our main estimates reach conventional significance levels, they are not very precise due to large standard errors. This is a typical problem with 2SLS estimates especially in the presence of relatively low numbers of observations. Despite the fairly low number of observations, the SOEP is the best data available for our 2SLS estimations since we need exact information on the month of birth of a child, which is not available in the German Micro Census. However, we can use the large samples of the Micro Census in a difference-in-differences approach, which might provide a useful context for our instrumental variable effects. The results of this alternative approach are presented in the next section.

5.2 Results from a Difference-in-Differences Framework

Between 1996 and 2001, there was a considerable increase in kindergarten attendance, from 53 percent to 70 percent, by three- and four-year-old children whereas kindergarten attendance for five and six year old children remained flat at a high level of 90 percent (Figure 1).[21] After 2001, the expansion of public child care places for three to four year olds stopped. In line with this observation, we use the Micro Census waves 1996 and 2001 for our basic difference-in-differences estimations. We define 1996 as the baseline year and 2001 as the post-treatment year. Mothers whose youngest child is three or four years old constitute our treatment group. As a control group, we use mothers with a youngest child aged 10 or 11 years and, as an alternative, two different groups of women without children under the age of 18. First, we use women aged 29 to 36 years without children; in terms of age, these women are comparable to our treatment group since they are in the two middle quartiles of the age distribution of mothers with children aged three or four. . As an additional control group, we drop the strict age restriction and use all women aged between 18 and 60 without children under the age of 18. In each wave we observe

---

[19] In a related robustness check, we additionally included the square of the age of the youngest child in months as a control variable. The point estimate on public child care gets insignificant due to large standard errors. Yet, the point estimate even gets slightly larger. Finally, we have also checked the sensitivity to the inclusion of observations right at the cut-off, i.e. mothers whose youngest child is exactly 36 months old at the kindergarten start. The point estimates are still large (0.258) but insignificant and smaller than the estimates from all other specifications.

[20] The results from these regressions are available from the authors upon request.

[21] As already noted earlier, formal child care for under three year olds was virtually non-existing in West Germany in the 1990s. Indeed, expansion of formal child care for under three year olds did not start before the mid-2000s.

more than 5,000 mothers whose youngest child is three or four years old (treatment group), at least 4,000 mothers whose youngest child is aged 10 or 11, more than 7,000 women aged 29 to 36 years without children , and more than 47,000 women aged 18 to 60 years without young children (control groups). For descriptive statistics for the treatment group and all control groups, both for the pre-treatment year 1996 as well as the post-treatment year 2001, refer to Table A.4 of the Online Appendix.

Figure 6 shows the development of women's employment over time for the treatment and control groups in West Germany. While employment levels obviously differ for our treatment and control groups, the pre-treatment trends from 1991 to 1996 are nearly identical across all groups of women. After 1996, we observe a similar increase in employment for our control groups of women without children and women whose youngest child is aged 10 or 11. In contrast, the rise in employment for mothers of children aged three or four years—our treatment group—is considerably larger. To highlight this relative increase in employment for our treatment group, we added a dashed line to Figure 6 that shows the counterfactual development of the treatment group had it not been treated. This counterfactual trend is the parallel shift of the employment development of mothers with 10- or 11-year-old children by the 1996 baseline difference in employment between this group and our treatment group. In other words, we assume that the pre-treatment-difference in employment between treatment group and control group would have been stable in subsequent years in absence of the treatment. We clearly observe that the increase in actual employment for our treatment group is more pronounced than is predicted by the counterfactual trend, which provides graphical evidence of a causal effect of public child care on maternal employment.

Figure 6 about here

In Columns 1, 3, and 5 of Table 6 we provide the basic difference-in-differences estimates from models where no additional covariates are included. The coefficient on the interaction term (After treatment * Treatment group) corresponds to the treatment effect θ of Equation (4). We find a positive effect of the increase in public child care between 1996 and 2001 on the employment of mothers of three- and four-year-old children ranging from 5.0 to 8.2 percentage points. Controlling for additional covariates in Columns 2, 4, and 6 hardly affects the estimates. Taking into consideration that public child care coverage for our treatment group increased from 55 to 70 percent from 1996 to 2001, the interpretation of these coefficients is straightforward: relying on the rather conservative estimates from Column 2, for example, a 10 percentage point increase in public child care coverage for three- to four-year olds leads to a 3.4 percentage point increase in employment of mothers of three- to four-year-old children. This estimate is reassuringly similar to the 2SLS results, where we find that a 10 percentage point increase in public child care coverage leads to a 3.7 percentage point increase in maternal employment.

Table 6 about here

Deviations from a common trend between treatment and control groups in the years preceding the actual treatment would make questionable the validity of our difference-in-differences approach. The largest threat to our identification would be if mothers with children aged three or four become more eager to enter the labor market and if it this eagerness that led to child care expansion in the years after 1996. To assess the validity of the key identifying assumption of our difference-in-differences approach, namely, that the time trends are the same for treatment and control groups in absence of the treatment, we conduct placebo treatment tests in the pre-treatment periods. More precisely, we use the Micro Census waves 1991 (pre-placebo-treatment period) and 1995 (post-placebo-treatment period) and analyze whether treatment and control groups follow the same trends during that period.

According to our placebo treatment estimates reported in Table 7, such reverse causality is not plausible. Starting with mothers whose youngest child is 10 or 11 as a control group, our placebo treatment estimates turn out to be insignificant. If at all, the coefficients on the interaction term (After treatment * Treatment group) are rather negative (Columns 1 and 2 of Table 7) This suggests that adjusting for group specific trends in our main difference-in-differences specification would lead to even larger estimates than those presented in Table 6. Accordingly, the results from Columns 1 and 2 of Table 6 should be interpreted as lower-bound estimates. In all other specifications using women without young children as our control group, the coefficients of the placebo interactions are virtually zero and far from any conventional significance levels. This means we do not observe any placebo treatment effects, i.e., mothers with children aged three and four and women without young children follow the same time trends in the years preceding the child care expansion (Columns 3–6 of Table 7).[22]

Table 7 about here

We also combined all pre- and post-treatment years from 1991 to 2001 in an extended difference-in-differences model. Note that robust standard errors might be downward biased in this special setting with many periods (Bertrand et al. 2004). We use the year 1996 as the baseline year and

[22] Although this result provides strong evidence in favor of the validity of the common trend assumption, it should be noted that we can of course not completely rule out that unobserved time-varying variables might affect treatment and control group differently in the post-reform years and thereby confound our estimates.

include interactions of the treatment group variable with all year dummies, while controlling for the main treatment group and year effects as well as age, education, and nationality. This exercise allows us to investigate the pre-treatment trends and the temporal emergence of the effects in more detail. Again, we use mothers whose youngest child is 10 or 11, women aged 29 to 36 without any children, and women aged 18 to 60 without any young children as the control groups. The detailed results of these estimations can be found in Table A.5 of the Online Appendix. A graphical depiction of the results is presented in Figure 7, where the data points represent the interaction coefficients of the difference-in-differences model. Since 1996 is the baseline year, the difference between treatment and control groups is normalized to zero in 1996. We observe that this difference is roughly constant through all pre-treatment years. There is no indication of a systematic divergence in trends already before the actual treatment sets in; indeed, the point estimates are very close to zero in most of the cases. This lends further support for the validity of the key identifying assumption that the post trends in the control groups serve as an estimate of what would have happened in the treated group in the absence of the reform. Only after the policy reform went into effect in 1996, we observe a gradual divergence in trends. Just as public child care attendance steadily increases for the treatment group from 1996 to 2001 (see Figure 1), so does relative employment of the treatment group as compared to the control groups. Note that the 2001 effect depicted in Figure 7 of course corresponds to the estimates from our basic difference-in-differences model (see Columns 2, 4, and 6 of Table 6). Taken together, this graph shows that the degree of the post-treatment divergence is clearly more pronounced than the ordinary fluctuation in the pre-treatment years, and that the increase in relative employment of the treatment group after the reform follows the gradual expansion in public child care. Thus, this exercise corroborates our previous results.

Figure 7 about here

In further regressions, we estimated the difference-in-differences model for mothers with partners and single mothers separately and found that the reduced form point estimates are larger for single mothers than for mothers with partners. The results of these additional estimations can be found in Table A.6 of the Online Appendix. Further inspection of the data also reveals that the treatment, i.e., the increase in public child care attendance of the three to four year olds, was also slightly larger for the group of single mothers than for the group of married mothers. Yet, also taking this difference in treatment intensity into account, the effect for single mothers still tends to be slightly larger than for mothers with partners.

Furthermore, we experimented with using mothers of three to four year olds from East Germany as an alternative control group. As can be seen from Figure A.1, child care attendance of three to four year old children in East Germany was flat and at a very high level during our period of observation, while child care attendance of three to four year old children in West Germany substantially increased from 1996 until 2001. Difference-in-differences regressions using this alternative control group yield results which are very similar to the ones from Table 6. However, the pre-treatment trends are rather different for these two groups. In particular, we observe a stark decline in East German employment (starting from a high level of more than 80 percent in 1991) during the first years after reunification, whereas the employment trend is rather flat in West Germany. This pre-treatment trend difference does not allow us to make too rigorous conclusions from this exercise. Very similar arguments apply to another exercise where we used mothers of under three year olds from West Germany as an alternative control group. Note that child care attendance for under three year olds in West Germany is flat and at a very low level of under 10 percent during our period of observation. Difference-in-differences regressions using this control group again yield results which are very similar to those presented in Table 6. However, also in this case, the pre-treatment trends differ between treatment and control group. In particular, while the trends rather converge from 1993 until 1996, they diverge again in the late 1990s to reach a maximum employment gap in 2001. Therefore, we also do not over-interpret the results from this exercise.

Taken together, our difference-in-differences estimates confirm the positive effects of public child care on maternal employment that we find in our instrumental variable approach. Indeed, estimates from the different empirical approaches are remarkably similar, although they do not necessarily measure the same effect. In particular, the IV setting identifies local average treatment effects (LATE) for the complier subpopulation whereas the difference-in-differences approach estimates average treatment effects on the treated (ATT). The reason for the fact that both approaches still yield very similar results could be that, in the end, in both identification strategies a very similar group of mothers is affected by the treatment, namely, those who have the opportunity to work because their three- or four-year-old child, who is also their youngest child, becomes eligible for public child care.[23]

[23] A further concern could be the existence of general equilibrium effects: the expansion of public child care might indeed increase labor market participation of mothers with children aged three or four. At the same time, however, rising employment of those mothers could lead to a crowding out of mothers with older children or women without children. In such a case, public child care expansion might enhance labor market participation of mothers whose youngest child is aged three or four, but *not* women's employment as a whole. We are confident that this is only a minor issue given the small group of mothers whose youngest child is aged three or four who would—by this

5.3 Discussion

Germany, in particular West Germany, is a country with traditionally low female and particularly maternal employment although women are generally very well educated. At the same time, public child care for three and four year olds was severely rationed before the reform analysed in this paper while virtually no private market for child care existed. Providing universal highly subsidized public child care should therefore yield large and positive effects on maternal employment, in particular for mothers whose *youngest* child becomes eligible for child care, which is exactly the group of mothers we focus on. Our reduced form estimates suggest that eligibility for public child care increases these mothers' employment by 6 percentage points. Our second stage IV estimates, which compare mothers whose youngest child attends kindergarten for the arguably exogenous reason of the cut-off rule to mothers whose youngest child cannot attend kindergarten due to the cut-off rule, suggest that a 10 percentage point increase in public child care attendance rates increases mothers' employment by 3.7 percentage points. Thus, roughly two thirds of mothers whose youngest child becomes eligible for public child care use the place in public child care for other reasons than taking up a job.

The effects identified in the instrumental variable approach are local average treatment effects for a compliant subpopulation of mothers in communities that actually had to apply cut-off rules possibly because they were less prepared and therefore confronted with substantial excess demand for public child care. Thus, we should in general be cautious in generalizing these results. However, the finding that the average treatment effects on the treated identified in the difference-in-differences strategy are very similar to the instrumental variable results, rather suggests that the effects identified in the 2SLS approach might also generalize to communities that did not use the cut-off rules to assign children to kindergarten yet still experienced the policy change of better access to public child care.

Yet, it is difficult to compare the size of our effects to the effects identified in other studies. First of all, many previous studies present intention-to-treat (reduced form) effects of a child care reform. Reduced form coefficients identified for different reforms in different countries can hardly be compared to each other since the size of the first stage, i.e., the effect of the reform on the take-up, might differ substantially across studies. Therefore, trying to compare the effect size of our study to previous studies, we focus on our second-stage (local) average treatment effects which identify the effect of child care attendance on mothers' labor supply. However, only few

argument—have to crowd out the large group of mothers and women who are not affected by the public child care expansion.

authors use data which actually allow them to compute (local) average treatment effects. Secondly, the second-stage coefficients identified in our paper come with large confidence intervals, which make a rigorous comparison of effects across studies very difficult. Thirdly, instrumental variables estimates identify the effect for the subpopulation of compliers, which might naturally differ across instruments, periods of observation, and countries, and at the same time be different from the effect for the average population. Fourthly, note that we focus on a very special group, namely mothers whose *youngest* child is three to four. The effect of public child care for this group should in general be stronger than the effect of public child care for *all* mothers who have a three to four year old child (and might also have younger ones). Taking these issues into consideration, we find that the size of our point estimates is slightly smaller than the effects of the introduction of free kindergarten for five year olds in U.S. public schools in the 1960s and 1970s on the labor supply of single mothers whose youngest child is five years old (Cascio 2009). Note that private child care was a viable alternative to public child care in the U.S. whereas it was not in West Germany, which is why we would expect larger effects in Germany than in the US. However, at the same time, the children we study are younger than the children in Cascio's (2009) study, which would be one reason why effects might be smaller in Germany than in the US.

6 Conclusions

This paper provides evidence of positive causal effects of public child care on labor market participation of mothers in Germany, a country with traditionally low female, and in particular maternal employment rates. To establish causality, we exploit a policy reform from 1996 which introduced a legal claim to a place in kindergarten for all children from age three until school entry. Prior to the reform, public child care coverage for three and four year olds was severely rationed in West Germany. At the same time, there was virtually no private market for child care. We rely on two identification strategies and two individual-level data sets. First, we use an instrumental variable approach that exploits age cut-off rules determining eligibility for public child care; these instrumental variable 2SLS estimations are based on rich individual-level data from the German Socio-Economic Panel (SOEP). Second, we apply difference-in-differences techniques exploiting the increase in public child care provision for three- and four-year-old children in Germany from 1996 to 2001 using data from the German Micro Census.

From our reading of the previous international literature on public child care and given the peculiarities of the German institutional setting, we expected to find positive effects of public child care on maternal labor supply in Germany. And indeed, both identification strategies yield positive effects of public child care on maternal employment. Our intention-to-treat estimations suggest that kindergarten eligibility increases the labor supply of mothers whose youngest child is three to four years old by 6 percentage points. Regressions identifying (local) average treatment effects suggest that 10 percentage point increase in public child care attendance increases employment of mothers whose youngest child is three to four years old by roughly 3.7 percentage points. On the other hand this means that roughly two thirds of these mothers use child care for other reasons than taking up a job. Several robustness checks, including placebo treatment estimations, corroborate the validity of our empirical approaches. From a social welfare perspective, it is clear that those that take advantage of public child care are better off. Public child care increases maternal employment and thereby mothers' labor income. [24] However, this comes at the cost of raising taxes to finance child care subsidies.[25]

In sum, our results show that expansion of public child care was an effective policy for increasing young mothers' labor market participation. Many studies on public child care effects and mothers' labor supply identify reduced form intention-to-treat effects; this makes it difficult to compare the size of the effects across these studies since the take-up rate might differ substantially across countries, and reforms. For our second-stage instrumental variable effects, however, we find that the point estimates are similar to the effects of the introduction of free kindergarten in U.S. public schools on the labor supply of single mothers whose youngest five year old child is affected by the reform (Cascio 2009). Still, large standard errors do not allow us to make too rigorous statements about the comparability of results..

Note that we identify labor market effects only for mothers whose youngest child is at least three years old at the time of the interview. Recent family and child care policies are aimed at improving coverage for children younger than three, for whom public child care attendance is still much lower than it is for three- to six-year olds. Evidence as to the labor market effects of public child care for this age group is nearly non-existent and thus presents an opportunity for further research. Also, it would appear worthwhile to more rigorously extend the scope of research on public child care to discover its effects on fertility. A few recent papers begin to investigate this issue, but to date the empirical evidence on how public child care provision affects fertility is scarce and inconclusive.

[24] For instrumental variable estimates of the effect of public child care on mothers' labor income, please see Table A.7 of the Online Appendix.

[25] In this respect, it might be interesting to know that cautious back-of-the-envelope calculations reveal that about 60 percent of the operating expenses of public child care net of parents' fees (i.e., public subsidies), are covered by income taxes and social security contributions induced by the increase in maternal employment.

Literature

Angrist, Joshua D. and Alan B. Krueger. 1991. Does Compulsory Schools Attendance Affect Schooling and Earnings, *Quarterly Journal of Economics*, 106(4), 979–1014.

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

Bassok, Daphna, Fitzpatrick, Maria, and Susanna Loeb. 2014. Does State Preschool Crowd-Out Private Provision? The Impact of Universal Preschool on the Childcare Sector in Oklahoma and Georgia, Journal of Urban Economics, 83, 18-33.

Berlinski, Samuel and Sebastian Galiani. 2007. The Effect of a Large Expansion of Pre-Primary School Facilities on Preschool Attendance and Maternal Employment, *Labour Economics*, 14, 665–680.

Berlinski, Samuel, Galiani, Sebastian, and Patrick J. McEwan. 2011. Preschool and Maternal Labour Market Outcomes: Evidence from a Regression Discontinuity Design, *Economic Development and Cultural Change*, 59(2), 313–344.

Bertrand, Marianne, Duflo, Esther, and Sendhil Mullainathan. 2004. How Much Should We Trust Difference-in-Differences Strategies?, *Quarterly Journal of Economics*, 119(1), 249–275.

Bettendorf, Leon, Jongen, Egbert, and Paul Muller. 2012. Childcare Subsidies and Labour Supply: Evidence from a Large Dutch Reform, CPB Discussion Paper, No.217.

Bick, Alexander. 2011. The Quantitative Role of Child Care for Female Labor Force Participation and Fertility, Frankfurt, mimeo.

Blau, David M. 2003. Child Care Subsidy Programs, in: Robert Moffit (Hrsg.), *Means Tested Social Programs*, Chicago: University of Chicago Press.

Blau, David M. and Janet Currie. 2006. Preschool, Day Care, and After School Care: Who's Minding the Kids? in: E. A. Hanushek and F. Welch (Hrsg.), *Handbook of the Economics of Education*, 2, 1163–1278.

BMFSFJ. 2005. *Zwoelfter Kinder- und Jugendbericht - Bericht über die Lebenssituation junger Menschen und die Leistungen der Kinder- und Jugendhilfe in Deutschland*. Berlin: Bundestagsdrucksache 15/6014.

Buechel, Felix and C. Katharina Spiess. 2002. Kindertageseinrichtungen und Müttererwerbstätigkeit: Neue Erkenntnisse zu einem bekannten Zusammenhang, *Vierteljahrshefte zur Wirtschaftsforschung*, 71(1), 95–113.

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

Coneus, Katja, Goeggel, Kathrin, and Grit Muehler. 2009. Maternal Employment and Child Care Decision, *Oxford Economic Papers*, 61, i172–i188.

Felfe, Christina and Rafael Lalive. 2012. Child Care and Child Development – What Works for Whom?, University of St. Gallen, mimeo.

Fitzpatrick, Maria Donovan. 2010. Preschoolers Enrolled and Mothers at Work? The Effects of Universal Pre-Kindergarten, *Journal of Labor Economics*, 28(1), 51–85.

Gathmann, Christina and Bjoern Sass. 2012. Taxing Childcare: Effects on Family Labor Supply and Children, IZA Discussion Paper No. 6440.

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

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

Konsortium Bildungsberichterstattung. 2006. Bildung in Deutschland. Ein indikatoren-gestützter Bericht mit einer Analyse zur Bildung und Migration. Bielefeld.

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.

Ludsteck, Johannes and Uta Schönberg. 2012. Expansions in Maternity Leave Coverage and Mothers' Labor Market Outcomes after Childbirth, *Journal of Labor Economics*, forthcoming.

Lundin, Daniela, Mörk, Eva, and Björn Öckert. 2008. How Far Can Reduced Childcare Prices Push Female Labour Supply? *Labour Economics*, 15(4), 647–659.

Nollenberger, Natalia and Núria Rodríguez-Planas. 2011. Child Care, Maternal Employment and Persistence: A Natural Experiment from Spain, IZA Discussion Paper No. 5888.

Rauschenbach, Thomas. 2006. Wer betreut Deutschlands Kinder – Eine einleitende Skizze, in: W. Bien, T. Rauschenbach and B. Riedel (eds.), *Wer betreut Deutschlands Kinder?* BeltzVerlag, Weinheim und Basel, 10–24.

Riedel, Birgit. 2005. Das institutionelle Angebot für Kinder von 3 Jahren bis zum Schuleintritt (Kindergartenalter), in: Deutsches Jugendinstitut (ed.), Zahlenspiegel 2005: Kindertagesbetreuung im Spiegel der Statistik, DJI, München, 127–142.

Rubin, Donald. 1978. Bayesian Inference for Causal Effects: The Role of Randomization, *Annals of Statistics*, 6, 34–58.

Schlosser, Analia. 2007. Public Preschool and the Labor Supply of Arab Mothers: Evidence from a Natural Experiment, University of Jerusalem, mimeo.

Schlotter, Martin. 2011. Age at Preschool Entrance and Noncognitive Skills before School - An Instrumental Variable Approach, Ifo Working Paper No. 112.

Spieß, Katharina, 1998. Staatliche Eingriffe in Märkte für Kinderbetreuung. Analysen im deutschamerikanischen Vergleich, Frankfurt a.M. & New York: Campus Verlag.

Statistisches Bundesamt. 2004. *Kindertagesbetreuung in Deutschland: Einrichtungen, Plätze, Personal und Kosten 1990 bis 2002*, Wiesbaden: Statistisches Bundesamt.

Statistisches Bundesamt. 2011. *Statistiken der Kinder- und Jugendhilfe: Ausgaben und Einnahmen 2009*, Wiesbaden: Statistisches Bundesamt.

Tietze, Wolfgang and Hans-Günther Rossbach. 1991. Die Betreuung von Kindern im vorschulischen Alter, *Zeitschrift für Pädagogik*, 37(4), 555–602.