

DISCUSSION PAPER SERIES

DP15437

DO FAMILY POLICIES REDUCE GENDER INEQUALITY? EVIDENCE FROM 60 YEARS OF POLICY EXPERIMENTATION

Henrik Kleven, Camille Landais, Johanna Posch,
Andreas Steinhauer and Josef Zweimüller

LABOUR ECONOMICS

PUBLIC ECONOMICS



DO FAMILY POLICIES REDUCE GENDER INEQUALITY? EVIDENCE FROM 60 YEARS OF POLICY EXPERIMENTATION

Henrik Kleven, Camille Landais, Johanna Posch, Andreas Steinhauer and Josef Zweimüller

Discussion Paper DP15437
Published 09 November 2020
Submitted 04 November 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Labour Economics
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Henrik Kleven, Camille Landais, Johanna Posch, Andreas Steinhauer and Josef Zweimüller

DO FAMILY POLICIES REDUCE GENDER INEQUALITY? EVIDENCE FROM 60 YEARS OF POLICY EXPERIMENTATION

Abstract

Do family policies reduce gender inequality in the labor market? We contribute to this debate by investigating the joint impact of parental leave and child care, using administrative data covering the labor market and birth histories of Austrian workers over more than half a century. We start by quasi-experimentally identifying the causal effects of all family policy reforms since the 1950s on the full dynamics of male and female earnings. We then map these causal estimates into a decomposition framework a la Kleven, Landais and Sogaard (2019) to compute counterfactual gender gaps. Our results show that the enormous expansions of parental leave and child care subsidies have had virtually no impact on gender convergence.

JEL Classification: N/A

Keywords: N/A

Henrik Kleven - kleven@princeton.edu
Princeton University and CEPR

Camille Landais - c.landais@lse.ac.uk
London School of Economics and CEPR

Johanna Posch - johanna.posch@analysisgroup.com
Analysis Group

Andreas Steinhauer - andreas.steinhauer@ed.ac.uk
University of Edinburgh and CEPR

Josef Zweimüller - josef.zweimueller@econ.uzh.ch
University of Zurich and CEPR

Acknowledgements

We thank Francine Blau, Hilary Hoynes, Emmanuel Saez, Johannes Spinnewijn, Gabriel Zucman, and seminar participants at the AEA Meetings 2019, Bank of Italy, UC Berkeley, Bocconi, Crest-Ecole Polytechnique, EIEF, OeGB Innsbruck, PSE, National Bank of Belgium, Stockholm University, University of Salzburg, and Queen Mary. We gratefully acknowledge support from the Center for Economic Behavior and Inequality (CEBI) at the University of Copenhagen, financed by grant #DNRF134 from the Danish National Research Foundation. Camille Landais gratefully acknowledges support from the LSE International Inequality Institute and from ERC grant #679704-DYNAMICSS.

Do Family Policies Reduce Gender Inequality? Evidence from 60 Years of Policy Experimentation*

Henrik Kleven[†]
Camille Landais[‡]
Johanna Posch[§]
Andreas Steinhauer[¶]
Josef Zweimüller^{||}

November 2020

Abstract

Do family policies reduce gender inequality in the labor market? We contribute to this debate by investigating the joint impact of parental leave and child care, using administrative data covering the labor market and birth histories of Austrian workers over more than half a century. We start by quasi-experimentally identifying the causal effects of *all* family policy reforms since the 1950s on the full dynamics of male and female earnings. We then map these causal estimates into a decomposition framework a la [Kleven, Landais and Søgaard \(2019\)](#) to compute counterfactual gender gaps. Our results show that the enormous expansions of parental leave and child care subsidies have had virtually no impact on gender convergence.

*We thank Francine Blau, Hilary Hoynes, Emmanuel Saez, Johannes Spinnewijn, Gabriel Zucman, and seminar participants at the AEA Meetings 2019, Bank of Italy, UC Berkeley, Bocconi, Crest-Ecole Polytechnique, EIEF, OeGB Innsbruck, PSE, National Bank of Belgium, Stockholm University, University of Salzburg, and Queen Mary. We gratefully acknowledge support from the Center for Economic Behavior and Inequality (CEBI) at the University of Copenhagen, financed by grant #DNRF134 from the Danish National Research Foundation. Camille Landais gratefully acknowledges support from the LSE International Inequality Institute and from ERC grant #679704-DYNAMICSS.

[†]Princeton University, NBER, CEPR, and CEBI. Email: kleven@princeton.edu

[‡]London School of Economics and CEPR. Email: c.landais@lse.ac.uk

[§]Analysis Group. Email: johanna.posch@analysisgroup.com

[¶]University of Edinburgh and CEPR. Email: andreas.steinhauer@ed.ac.uk

^{||}University of Zurich and CEPR. Email: josef.zweimueller@econ.uzh.ch

1 Introduction

Is government intervention critical for reducing gender inequality in the labor market? Or is gender inequality driven primarily by factors outside the government's control? Recent research shows that the bulk of gender inequality can be attributed to the unequal impacts of parenthood on men and women (Bertrand, Goldin and Katz 2010; Kleven, Landais and Sogaard 2019; Kleven, Landais, Posch, Steinhauer and Zweimuller 2019), which suggests that family policies such as parental leave and child care have the greatest potential to affect gender inequality. A sizeable literature studies the impact of family policies on female labor market outcomes, typically focusing on the contemporaneous impacts of marginal reforms in isolation. Apart from cross-country evidence, there is little research trying to estimate the aggregate long-run impact of public policy on gender inequality (Olivetti and Petrongolo 2017). To make progress on this question, we investigate the joint impact of parental leave and child care over more than half a century. Our main conclusion is simple and striking: The enormous expansions of parental leave and child care subsidies have had virtually no impact on gender convergence.

To reach this conclusion, we rely on exceptional administrative data covering the labor market and birth histories of Austrian workers from 1953 to 2017. During this era, the Austrian gender gap in earnings fell by about 30 percentage points, while public policies supporting families with children were introduced and greatly expanded. Parental leave was introduced in 1961 and subsequently expanded (and in one instance scaled back) through a series of major reforms. At the same time, the provision of heavily subsidized public child care was rolled out across the country. Similar expansions of family policies have been implemented in other European countries and their implications for gender inequality is the subject of heated debate: Is access to job-protected parental leave and subsidized child care instrumental for boosting female labor force participation and reducing gender gaps? Are generous parental leave schemes counterproductive by inducing mothers to stay at home rather than investing in their careers? Or is the emphasis on policy solutions overstated, with gender inequality being driven mostly by equilibrium features of the labor market such as the temporal flexibility of jobs (Goldin 2014; Goldin and Katz 2016) and gender norms or culture (Bertrand 2011; Kleven and Landais 2017; Kleven, Landais, Posch, Steinhauer and Zweimuller 2019; Boelmann, Raute and Schönberg 2020)?¹

¹Another hypothesis is that gender gaps are endemic to biological gender differences, but Kleven, Landais and Sogaard (2020) provide evidence against the importance of such channels.

We take advantage of the richness of the Austrian administrative registers combined with quasi-experimental techniques to implement a unique, bottom-up approach to measuring counterfactual gender inequality absent family policies. Our approach proceeds in two steps. First, we identify the causal effects of *all* family policy reforms since the 1950s on the labor market outcomes of men and women. In a second step, we map these causal estimates into a decomposition framework to compute counterfactual gender gaps absent policy intervention.

Besides our long-run historical approach, we add two key innovations to the existing literature. The first innovation is to use “child penalties” in earnings as our main outcome. This ensures that we identify the full dynamic effect of each policy change on the labor market trajectories of women relative to men. The second innovation is to account for the potential interaction effects between parental leave and child care policies. Family policies are often designed as a bundle, with child care coverage starting around the time that parental leave eligibility is ending. As a result, we analyze the effects of parental leave and child care policies jointly, capturing any non-linearities brought about by their potential interaction.

Leveraging a regression discontinuity design based on cutoff dates of the parental leave reforms, we first show that parental leave duration has significant negative effects on the earnings of women relative to men in the short run. We show that these effects are non-linear, being greater at lower baseline durations. However, for all parental leave reforms, we find precisely estimated zero effects in the long run. Leveraging a difference-in-differences design similar to [Duflo \(2001\)](#) based on large expansions of locally provided child care, we find precisely estimated zero effects of both nursery care and pre-school care on the earnings of women relative to men. We provide suggestive evidence that this null effect is driven by crowding out of other types of informal care, and the persistence of strong norms for maternal care. Finally, we find no interaction effect between child care and parental leave policies.

In the second step of our analysis, we use these estimates to back out counterfactual long-run series of gender inequality. We build on the framework developed by [Kleven, Landais and Sogaard \(2019\)](#), which decomposes gender inequality using cohort-specific estimates of child penalties. The intuition is simple: because we know how each policy reform causally impacted child penalties, we can reconstruct cohort-specific child penalties (and therefore aggregate gender inequality over time) in the absence of policy intervention. Our results imply that the dramatic expansion of family policies has had a remarkably small impact on gender inequality. Due to the negative effects of parental leave on labor supply in the short-run, family policies have been slightly counterproductive for gender convergence. The average earnings of Austrian women is 43.5% lower than the average earnings of men today, while they would have been 41.5% lower

if family policies had remained as they were in 1960. Sensitivity analyses suggest that this null effect is very precise. We also demonstrate that our conclusions are robust to accounting for frictions, equilibrium/spillover effects, and fertility responses.²

Our paper contributes to the large literature on gender inequality in the labor market (reviewed by [Bertrand 2011](#) and [Olivetti and Petrongolo 2016](#)) and especially to recent work showing the crucial role of parenthood ([Bertrand, Goldin and Katz 2010](#); [Kleven, Landais and Sogaard 2019](#); [Kleven, Landais, Posch, Steinhauer and Zweimuller 2019](#); [Kuziemko, Pan, Shen and Washington 2018](#); [Cortés and Pan 2020](#)). Our paper also contributes to a large literature on the impact of family policies (reviewed by [Olivetti and Petrongolo 2017](#)), including parental leave schemes (see e.g., [Lalive and Zweimüller 2009](#); [Schönberg and Ludsteck 2014](#); [Dahl, Løken, Mogstad and Salvanes 2016](#)) and child care subsidies (see e.g., [Baker, Gruber and Milligan 2008](#); [Havnes and Mogstad 2011](#)). We discuss how our findings relate to these literatures in more detail below. While we are far from the first to study these questions, our paper sheds new light by studying them within a unified quasi-experimental setting in order to quantify the aggregate importance of family policies for the long-run historical evolution of gender inequality.

Our bottom-up counterfactual methodology offers a credible approach to measuring the impact of policy interventions on the long-run evolution of gender inequality. The conclusions stand in sharp contrast to those that can be gleaned from cross-country approaches, which tend to find strong correlations between public policy and female labor market outcomes (e.g., [Ruhm 1998](#); [Blau and Kahn 2013](#); [Kleven 2014](#); [Olivetti and Petrongolo 2017](#)). The conceptual approach developed here could be used to study the role of policy intervention in other areas of inequality research (such as income and wealth inequality).

The remainder of the paper is organized as follows. Section 2 describes the Austrian institutional context and data. Section 3 presents quasi-experimental evidence on the impact of parental leave reforms on child penalties. Section 4 presents quasi-experimental evidence on the impact of local child care expansions on child penalties. This section also provides evidence on the potential mechanisms behind our null effect. Section 5 lays out our decomposition framework, and provides estimates of the counterfactual evolution of gender inequality under different policy trajectories. Section 6 concludes.

²Our null result is striking given that the subsidies to parental leave and child care in Austria (and in other European countries) create large changes in participation tax rates (accounting for all taxes, transfers, and subsidies) during the years following child birth. In fact, the implications of these policies for the extensive margin incentives of women with young children tend to be much larger than the implications of, for example, the famed EITC program in the U.S. See [Kleven \(2014\)](#) for a cross-country analysis of participation tax rates and labor supply.

2 Context and Data

2.1 Institutional Context

A (Typical) Gender-Conservative Environment Austria, like other German-speaking countries, is characterized by relatively large differences in labor market outcomes between men and women. The raw gender gap in earnings is roughly 40% and the female labor force participation rate is about 55% (in 2018). However, it is worth highlighting two factors that make Austria an interesting laboratory for studying the impact of family policies on gender inequality. First, Austria is not an international outlier in terms of gender gaps. English-speaking countries like the U.S. and the U.K. have broadly similar levels of gender inequality in earnings. Moreover, female labor force participation in Austria is higher than in a number of other European countries such as Italy, Spain, and France. Second, Austria has experienced a substantial reduction in gender gaps and a surge in female labor force participation over the past 50 years. This makes it a very useful setting to study whether this reduction in gender inequality is causally related to family policy expansions.

In addition to inequality in labor market outcomes, Austria is often singled out, together with other German-speaking countries, for having very conservative gender norms. Appendix Figure [A.I](#) probes the idea that gender norms are more conservative in Austria than elsewhere. The evidence presented is based on questions from the 2012 wave of the International Social Survey Program (ISSP) on the attitudes that people have towards market work by women with and without children. Two insights emerge from the figure. One is that gender attitudes are quite traditional in most countries: large fractions of people hold the view that women should stay at home when they have young children (panel (a)), and that, if they don't, their children will suffer (panel (b)). The other insight is that Austrians, while being on the more conservative side of the spectrum, are not so different from the citizens of other countries. Their views on whether mothers of young children should stay at home instead of seeking employment are very similar to the views observed in the U.S., the U.K., and France. Overall, the evidence in Figure [A.I](#) suggests that Austria is gender conservative, but not a strong outlier internationally. Austria is a typical gender-conservative environment, and as such, a great setting for our study. We are interested in better understanding whether family policies can change the evolution of gender inequality when gender attitudes are still traditional.

Family Policies Over the last 60 years, most developed countries have introduced and dramatically expanded two sets of family policies: parental leave and child care provision ([Olivetti and Petrongolo 2017](#)). Austria is a case in point.

In the aftermath of World War II, there was no system of parental leave in Austria. Women were simply entitled to *Mutterschutz*, a maternity insurance period of 8 weeks following delivery, akin to sick leave. In 1961, a paid and job-protected maternity leave of 52 weeks was introduced. Since then, the country has implemented four large reforms of its parental leave policies (in 1990, 1996, 2000, and 2008), which have considerably expanded the generosity and scope of parental leave.³ Today, the Austrian parental leave system stands out as one of the most generous in the developed world. Currently, parents are entitled to up to about 35 months of paid parental leave if that time is shared, or up to about 28 months if one parent takes the time alone. The replacement rate is relatively high: the median replacement rate is about 40% of net earnings. The take-up of parental leave comes with a job-protection guarantee of 24 months. The receipt of parental leave benefits is not subject to any work eligibility requirement prior to child birth.

Institutional child care in Austria is divided into nursery care (*Krippe*) covering children from ages 1 to 2 and pre-school care (*Kindergarten*) from ages 3 to 5. The provision of both forms of institutional child care has been greatly expanded over the past 50 years. In the beginning of the 1970s, less than 2% of children between the age of 1 and 2 attended nursery. Today, this number has risen to 25%. The fraction of children attending pre-school has risen from 40% to more than 90% over the same time period. Child care is provided by municipalities, following general guidelines set by state legislation.⁴ The median number of daily operating hours for Austrian child care institutions is 9 hours, with typical opening hours from 7:30am to 4:30pm.⁵ Institutional child care has always been heavily subsidized in Austria: out-of-pocket costs are among the lowest in developed nations. According to the OECD Family Database, out-of-pocket costs for a family with two children aged 2 and 3 in full-time daycare was just 3% of family net income in 2015, compared to an average of 13% across the OECD and 22% in the U.S.⁶ Note that subsidized institutional child care encompasses both public and church provided care. The private market for child care is very limited and essentially consists of a few private creches and child-minders (*Tagesmutter*).⁷

³Two further reforms in 2010 and 2017 provided additional flexibility to parents regarding the length of leave and the amount of the transfer. We do not evaluate these policy reforms here because the time-horizon to study future earnings dynamics is too short.

⁴State legislation typically determines the opening hours, the fee structure, and the total number of children that can be supervised by a teacher. For details see [Baierl and Kaindl \(2011\)](#).

⁵See [Baierl and Kaindl \(2011\)](#).

⁶<http://www.oecd.org/els/family/database.htm>. Note that small variations in prices exist across states, and sometimes across municipalities.

⁷Figures from a microcensus survey in 2002 reveal that less than 3% of children under age two were taken care of by *Tagesmutter*.

2.2 Data

Labor Market Outcomes & Birth History Our study relies on unique data covering the labor market and birth histories of Austrian workers from 1953 to 2017. Our main data come from the Austrian Social Security Database (ASSD, see [Zweimüller, Winter-Ebmer, Lalive, Kuhn, Wuellrich, Ruf and Büchi 2009](#) for details). This is a matched employer-employee dataset, which contains detailed information on employment and earnings by calendar year from 1972 to 2017. The ASSD register covers the universe of paid employment spells, with the notable exclusion of civil servants, self-employed and farmers. This means that about 85% of the Austrian population is covered by ASSD, and this coverage rate has been stable over time.

Earnings recorded in the ASSD correspond to the amount on which workers pay social security contributions, which is subject to a yearly-adjusted cap. As a result, our earnings variable is top-coded for about 5% of workers. To deal with top-coding, we replace all right-censored earnings with mean earnings above the cap predicted from a Pareto distribution, the parameters of which we obtain by gender and age groups using uncensored income tax data.⁸ Besides earnings, the ASSD includes information on the uptake of all social insurance programs along with demographic variables such as gender, year of birth, year of death, nationality, and all child birth events of women (e.g., date of birth of each child).

ASSD data exist before 1972, but has not been digitized. However, detailed tabulations on earnings by gender, drawn from the original data, were regularly published. This enables us to consistently compute the evolution of gender inequality on the population covered by the ASSD since 1953. Furthermore, we use additional micro-level earnings information before 1972 from the *Rueckwirkend Erfasste Versicherungsdaten* (REV) register. The REV retrospectively digitised, for ASSD workers retiring in the 1980s and 1990s, earnings and employment spells prior to 1972 for the purposes of pension calculations.⁹ The selective nature of this data means that only certain cohorts are well covered. Fortunately, for birth cohorts 1920-1950, which are primarily those having children between 1955-1972, the REV covers about 80% of workers, as shown in Appendix Figure D.V. Although the REV does not contain direct information on child births, it has information on the uptake of maternity insurance (*Mutterschutz*), which has been available to Austrian women ever since the 1950s. As a consequence, the REV data enables us to consistently estimate child penalties for women giving birth all the way back to the mid-1950s. We are not aware of any existing work estimating child penalties over such a long time period.

⁸We provide further details on our methodology and its robustness in Appendix section D.1.

⁹The earnings information is provided in the same way as in the ASSD (except there is no information about the firm), and we adjust for topcoding using the same Pareto estimates.

Importantly, both the ASSD and REV contain child birth information only for women. We can nevertheless link children to their fathers for a large subset of child births. To this end, we use the child benefit register (*Beihilfedaten*). This register contains the universe of child benefit claims filed between 1995-2012. As the take-up of child benefits is almost complete, we are able to match around 90% of children born between 1995 and 2012 cohorts. Each claim lists the social security number of the claimant (usually the mother) as well as that of her partner. We pick the first entry in a child’s life and set the man listed on that claim as the father.

Child Care Provision Our analysis also relies on granular information on child care provision at the municipal level. We have obtained administrative data collected by Statistics Austria, which records, for every year since 1988 and for each of the nearly 2,000 Austrian municipalities, the number of child care facilities (broken down by nurseries and pre-schools), their opening hours (full-day or half-day), and the number of teachers in each facility. From this data, we are able to construct precise measures of child care supply at the municipal level. We provide details about the construction of these measures in section 4.

2.3 Child Penalties

Our analysis of the impact of family policies on gender inequality over the last 60 years relies on two inputs. First, we need to estimate how the full dynamics of parental outcomes respond to the arrival of children, what we refer to as *child penalties*. Second, we need to identify the causal effect of all family policy reforms on these child penalties. In this section, we describe how child penalties are defined and estimated, and document their magnitude in the Austrian context.

To estimate the impact of children on the labor market trajectories of mothers and fathers, we adopt the event-study specification proposed by [Kleven, Landais and Sogaard \(2019\)](#). For each parent in the data, event time t is indexed relative to the year of the first child birth. Denoting by Y_{ist}^g the outcome for individual i of gender g in year s and at event time t , we run the following regression separately for men and women:

$$Y_{ist}^g = \boldsymbol{\alpha}^g \mathbf{D}_{ist}^{Event} + \boldsymbol{\beta}^g \mathbf{D}_{ist}^{Age} + \boldsymbol{\gamma}^g \mathbf{D}_{ist}^{Year} + v_{ist}^g \quad (1)$$

On the right-hand side, we use boldface to denote vectors. The first term includes event time dummies, indexed such that $t = 0$ denotes the year of arrival of the first child.¹⁰

¹⁰For notational convenience, we get rid of the transpose operator and denote by $\boldsymbol{\alpha} \mathbf{D}$ the scalar product of two vectors of same dimension, $\boldsymbol{\alpha}$ and \mathbf{D} .

We omit the event-time dummy at $t = -1$, implying that the event-time coefficients measure the impact of children relative to the year just before the first child. The second term includes age dummies (to control for lifecycle trends) and the third term includes year dummies (to control for time trends). We are able to identify the effects of all three sets of dummies because, conditional on age and year, there is variation in event time driven by variation in the age at which individuals have their first child.¹¹ Our main outcome variable is gross labor earnings, specified in levels.¹² We convert the estimated level effects into percentages to obtain the child penalty P_t^g for gender g at event time t :

$$P_t^g \equiv \frac{\hat{a}_t^g}{E[\tilde{Y}_{ist}^g | t]} \quad (2)$$

where \tilde{Y}_{ist}^g is the predicted outcome when omitting the contribution of the event time dummies. This implies that P_t^g can be interpreted as the percentage loss of average earnings due to having children.¹³

Figure 1 shows estimated child penalties for different outcomes in our ASSD sample. The magnitude of earnings penalties for mothers is very large (panel (a)). Ten years after the birth of the first child, Austrian mothers suffer a 50% earnings drop relative to fathers. Panels (b), (c) and (d) show that this effect reflects a combination of penalties along three margins: the extensive margin of labor supply, the intensive margin of labor supply, and the wage rate. Although these patterns are qualitatively similar to those documented in Scandinavia, the magnitudes are significantly larger in Austria. As shown in Appendix Figure A.II, long-run earnings penalties in Scandinavia lie in the 20-30% range, about half the size of the long-run penalties observed in Austria and Germany.

On average, over our period of interest, Austrian mothers have thus faced large career costs from having children. In the next two sections, we turn to analyzing how family policies have affected the evolution of these large penalties.

¹¹Kleven, Landais and Sogaard (2019) lay out the identification assumptions underlying this approach, compare its results to alternative approaches in the literature, and provide evidence of its ability to identify the causal effect of parenthood.

¹²We specify equation (1) in levels rather than in logs to be able to keep the zeros in the data due to nonparticipation.

¹³Importantly, under the identifying assumptions, equation (2) corresponds to $\frac{\hat{a}_t^g}{E[\tilde{Y}_{ist}^g | t]} = \frac{E[y_{1,t}|t] - E[y_{0,t}|t]}{E[y_{0,t}|t]}$, where $y_{1,t}$ is the actual earnings at event time t and $y_{0,t}$ is the counterfactual earnings absent children. In contrast, a log specification would estimate the average percentage earnings loss, which differ if treatment effects are correlated with counterfactual earnings. That is, if for example the effect of children on earnings is higher or lower for women who would have had higher earnings absent children. P_t^g is therefore the relevant metric when assessing the effect of children on the aggregate gender pay gap.

3 The Impact of Parental Leave on Child Penalties

As mentioned above, Austria, like most developed countries, has considerably expanded the scope and generosity of paid parental leave (PL). We now turn to the question how the dramatic expansion in parental leave duration since the 1960s affected the labor market outcomes of Austrian mothers and fathers?

To investigate this question, we start by focusing on the parental leave regime that was in place in 2008. This regime offered parents the possibility for households to choose between three PL options: (i) a long parental leave (the “30+6” option), with 30 months for the mother and 6 months for the father at a daily flat benefit of 14.53€; (ii) an intermediate option of 24 months, split between 20 months for the mother and 4 months for the father, (the “20+4” option), at a daily benefit of 20.8€; (iii) a short parental leave of 18 months, with 15 months for the mother and 3 months for the father at the daily rate of 26.6€ (the “15+3” option).¹⁴ In Figure 2, we split parents of children born in 2008 into three groups, according to the PL option they selected into, and report the child penalty P_t^g estimated separately on each of the three groups. The figure first shows that the earnings dynamics of fathers remains unaffected by the arrival of children, irrespective of the length of the PL option selected by the household. In contrast, the graph displays stark differences in the child penalties across PL options for mothers. The magnitude and persistence of the drop in their earnings after child birth strongly correlates with the length of their PL. Nine years after child birth, women who opted into the “30+6” option exhibit penalties that are still up to 15 percentage points larger than women who selected into the “15+3” option.

The large correlation observed in Figure 2 between PL duration and female child penalties raises the question: is it causal? In other words, how much of this large correlation is explained by the causal effect of PL duration on child penalties? And how much is due to differential selection of couples into PL options? To answer this, we now turn to the four large reforms that exogenously changed the duration of parental leave over the past 60 years.

3.1 The 1990, 1996 and 2000 Reforms

Empirical Strategy We start by focusing on the three reforms that took place in 1990, 1996 and 2000. These reforms share important institutional similarities that allow us to study them using the exact same empirical strategy.

First, all three reforms focused on the *duration* of the paid parental leave. The 1990

¹⁴Households need to choose one of these three options in the 8 weeks following the birth of the child. Their choice is legally binding and cannot be altered ex-post.

reform increased the potential duration of parental leave from 12 to 24 months. The 1996 reform introduced the rule that each parent had to take at least 6 month of leave, which meant a reduction in the maximum duration of *maternal* leave from 24 to 18 months. Finally, the 2000 reform increased again the maximum duration of maternal leave from 18 to 30 months.¹⁵ Importantly, the other key parameters of the PL system, that is, job protection provisions, and the level of benefits, were all kept constant over the period 1990 to 2000. In terms of job protection provisions, since 1990, a woman has the right to return to her previous employer up to 24 months after the birth of her child. As far as benefits are concerned, they were, until 2008, a flat amount, indexed on wage inflation, that did not depend on individuals' previous earnings, nor on the duration of the leave.¹⁶ As a result, the 1996 and 2000 reforms identify the effects on child penalties of both extending and reducing the duration of PL, conditional on all other parameters of PL policies, while the 1990 reform identifies the effect of extending both the duration of paid leave and job protection.

A second common feature of these reforms is that their implementation relied on a strict eligibility cutoff, based on the date of birth of the child, with no grandfathering. For the 1990 reform for instance, mothers of children born after July 1st 1990 became eligible to the new regime, and received a maternal leave entitlement of 24 months. Mothers of children born before this date remained entitled to 12 months of leave. The corresponding cutoff dates for the 1996 and 2000 reforms are July 1st 1996 and July 1st 2000. This implementation structure naturally lends itself to a regression discontinuity design, using the date of birth of the first child relative to the cutoff date as the assignment variable.¹⁷

Regression Discontinuity Evidence We are interested in the effect of parental leave reforms on the full dynamics of female outcomes around the event of the first child

¹⁵Before 1990, the parental leave benefits were only available to mothers. The 1990 reform offered the possibility for couples to split the total credit of 24 months between the mother and the father. Fathers were then entitled to take up to 6 months out of the 24 months, but could only take their leave in the second year of the child. In practice, the take up rate of paternal leave was zero. We therefore treat the 1990 reform as an increase in the potential duration of maternal leave from 12 to 24 months. The 1996 reform kept the total duration of parental leave to 24 months, of which one parent could not take more than 18 months. We therefore treat the 1996 reform as a decrease in the potential duration of maternal leave from 24 to 18 months. The 2000 increased the maximum duration of parental leave to 36 months, of which one parent could not take more than 30 months. We therefore treat the 2000 reform as an increase in the potential duration of maternal leave from 24 to 18 months.

¹⁶Note also that prior work history requirements did not change over the period 1990 to 2000, stipulating that an individual needed to have worked (and paid social security contributions) for at least 52 weeks during the two years prior to birth to be eligible to benefits. The 2000 reform introduced the possibility for mothers without prior work history to collect child benefits for up to 30 months. We restrict our sample to women meeting the 52 weeks work requirement to make our estimates comparable across all three reforms.

¹⁷Density tests systematically confirm the absence of manipulation of the timing of births around the three cutoff dates.

birth.¹⁸ To explore this effect, Appendix Figure B.I starts by displaying, for all event years $t \in [-4; 10]$, the average earnings of mothers by weekly bins of the running variable, in an 8-months window around the cutoff date of the 1990 reform. Starting with panel (a), we see that earnings 4 years before birth do not display any sign of discontinuity at the cutoff date in the running variable. Similarly, panels (b) to (d) show no sign of discontinuity in earnings in event year -3 to event year -1. The evidence from these five panels validate the credibility of our RD design, by corroborating that there is no selection in pre-birth earnings around the cutoff date of the reform. Moving to panels (f) and (g), corresponding to earnings in event years 1 and 2, we find a strong discontinuity in earnings at the cutoff. Because earnings are measured at the calendar year level, these two event years correspond to the moment were the take up of the parental leave extension has bite.¹⁹ The PL extensions appears to cause a significant drop in annual earnings of more than 1,000€ in both of these calendar years. Strikingly, these significant negative effects fully disappear in event year 3, as shown in panel (h). And we find no sign of discontinuity in the earnings of mothers for any event year after that, as evidenced by panels (i) to (o). Note also that in panel (e), we see no evidence of an effect of the PL extension on earnings in the year of birth of the child.²⁰ This suggests the absence of any extensive margin response to the reform. In other words, women who were not taking maternal leave are not more likely to take a leave because of its longer duration.

To confirm the visual diagnostic from Appendix Figure B.I, we estimate the effects of the 1990 reform on earnings Y_{it}^w for each event year t using the following event-study specification:

$$Y_{it}^w = \alpha D_{it}^{Event} + \alpha^T D_{it}^{Event} \cdot Treat_i + \alpha^B D_{it}^{Event} \cdot Birth_i + \delta D_{it}^{Event} \cdot Treat_i \cdot Birth_i + \beta D_{it}^{Age} + v_{it} \quad (3)$$

To leverage the RD design for identification, we estimate this specification on a sample of mothers who give birth in a narrow window of four months around July 1st, and compare women who gave birth just after the cutoff date ($Treat = 1$) to women

¹⁸We focus the following analysis on mothers, as fathers' outcomes appear unaffected by PL regimes, as shown in Figure 2. For completeness, we confirm at the end of this section that the estimated effects of the PL reforms for fathers are precisely zero.

¹⁹For women who give birth in July 1990, event year 1 corresponds the period going from January to Dec 1991, that is the period going from month 6 to month 18 after birth. Similarly, event year 2 corresponds to the period going from 18 to 30 months after birth.

²⁰The strong positive slope in the relationship between earnings in event year 0 and the running variable is explained by the fact that earnings are measured at the calendar year level in the ASSD data. Event year 0 therefore corresponds to calendar year 1990. As women who give birth in the earlier months of 1990 spend more time in parental leave in that calendar year, their earnings level is mechanically lower.

who gave birth just before the cutoff date ($Treat = 0$). We control for seasonality using women who gave birth in a similar four-months window around July 1st in 1989 ($Birth = 1989$). The impact of the reform on earnings at each event time t is identified by the parameters δ_t . To make magnitudes interpretable, we scale $\hat{\delta}_t$ by the average counterfactual earnings $E[\tilde{Y}_{it}^w | t, Birth_i = 1990]$, obtained from specification 3 when omitting the contribution of the event dummies. The statistic $\hat{\delta}_t / E[\tilde{Y}_{it}^w | t]$ captures the percentage-point change in the child penalty at event time t caused by the reform. Results, reported in Panel (a) of Figure 3, first confirm the absence of selection on earnings prior to child birth. They also confirm the absence of any significant effect of the reform on earnings in event year 0. The negative effect of the 1990 PL extension is entirely concentrated in event year 1 and 2, when the take up of the additional leave occurs: for both years, the extension from 12 to 24 months of maternal leave causes a 10% drop in female earnings. We find no significant effect of the 1990 reform on the earnings of women beyond event year 2. Overall, we find that the 1990 PL extension caused an average 4.04% (0.62%) reduction in the annual earnings of mothers over the first 5 years after the birth of their first child. In Panel (b) of Figure 3, we draw counterfactual child penalties under the pre and post 1990 PL regimes. We start from the baseline child penalties estimated on the sample of mothers who gave birth before the 1990 reform, to which we add, for each event year, the estimates of the effects of the reform from panel (a). This gives us the counterfactual child penalty under the 24 months maternal leave regime. The graph shows clearly that extending maternal leave from 12 to 24 months causes a larger child penalty in the short run, but has absolutely no impact on the child penalty in the long run.

Using the same empirical strategy, we can now turn to measuring the effects of the 1996 reform. This reform reversed part of the 1990 extension, by reducing the duration of maternal leave from 24 to 18 months. Visual inspection of the relationship between earnings and the assignment variable for all event years in Appendix Figure B.II reveals no sign of discontinuity at the July 1st 1996 cutoff except for event year 2. In that year, we see an increase in earnings of more than 2000€ for women who gave birth just after the cutoff date, and are therefore exposed to the shorter maternal leave regime. Panel (c) of Figure 3 reports the associated dynamic RD estimates of the effects of the 1996 reforms, which confirm three key findings from the 1990 reform. First, we see, once more, no sign of extensive margin response to the reform. The reduction in parental leave duration did not significantly affect the earnings of mothers in event year 0 and 1. Put it differently: mothers did not reduce their propensity to take *any* leave in response to the 1996 reform. Second, the effects of the reform are fully concentrated in event year 2, that is, on the period where the variation in parental leave take up has bite.²¹ Finally,

²¹For women who gave birth around the July 1st 1996 cutoff, event year 2 is defined as calendar year going from January to December 1998, which corresponds to 18 month to 30 months after birth.

the reform does not seem to have any long term effects on the careers of mothers: the increase in labor supply and earnings in event year 2 generated by the reform does not translate into any positive earnings gains in the long run.

Panel E of Figure 3 replicates the same analysis for the 2000 reform, which extended back the maternal leave duration from 18 to 30 months. The same three striking findings emerge again. The only noticeable difference is the smaller magnitude of the effect of the 2000 reform overall. Compared to the 1990 reform, which also extended the duration of the parental leave by one year, we find that the 2000 reform caused a significantly smaller decline of 2.60%, compared to 4.04%, in the average earnings of mothers in the first 5 years after the birth.

Finally, for completeness, in Appendix Figure B.IV, we investigate the effect of these parental leave reforms on the careers of fathers. Since 1990, fathers have been entitled to take up to 6 months of parental leave. And it is important to investigate whether variations in the duration of leave available to mothers generated substitution effects towards paternal care. Following the same methodology as we did for mothers, we compare fathers of children born on both sides of the eligibility cut-off dates, and find, for all three parental leave reforms, that they had precisely zero effect on the dynamics of male earnings.

3.2 The Introduction of Parental Leave in 1961

We now turn to the introduction of paid maternal leave in January of 1961. This reform entitled mothers to receive maternity benefits for up to one year after the birth of a child.²² Maternity benefits were set to the level of unemployment insurance benefits, and receipt was conditional on not working or owning arable land. The maternity leave scheme came with job protection. Empirical analyses of reforms that introduce both job protection and pecuniary benefits have been scant.²³ The main reason is that these reforms were implemented a long time ago in most countries, prior to the era usually covered by administrative registry data. In our context, however, we are able to leverage the REV register, which has information on labor market outcomes around the time of the 1961 reform.

To estimate the causal effect of this reform on the earnings dynamics of mothers, we need to account for the fact that the reform was grandfathered. That is, women who gave birth *before* January 1961 were also eligible for the new leave scheme, which they could use up until the first birthday of their child. As a result, scheme eligibility did not

²²Maternity benefit payments (*Karenzurlaub*) technically start at the end of the maternity insurance period (*Mutterschutzgesetz*), which covers women for 8 weeks following delivery.

²³These reforms are nevertheless particularly interesting in the context of current debates in the US regarding the desirability of a paid and federally mandated maternity leave.

display a sharp discontinuity for mothers of children born after January 1961. Indeed, Appendix Figure B.V shows that take-up of maternity benefits featured a gradual increase starting with children born in January 1960. To deal with this issue, instead of the RD approach used above, we analyze the 1961 reform using a DiD approach comparing different birth cohorts over time. Specifically, we compare the outcomes Y_{ist}^w of women whose first child was born between May and August 1961 and who were therefore fully eligible for maternity leave benefits ($T = 1$) to women whose first child was born between May and August 1959 and who were fully ineligible ($T = 0$). We implement this approach by running the following event-study specification on the sample of women who had their first child in either 1959 or 1961:

$$Y_{it}^w = \alpha D_{it}^{Event} + \alpha^T D_{it}^{Event} \cdot T_i + \beta D_{it}^{Age} + \gamma D_{it}^{Year} + v_{it} \quad (4)$$

Figure 4 shows the time series of DiD estimates $\frac{\hat{\alpha}_i^T}{E[\tilde{Y}_{it}^s | t, T=1]}$, corresponding to the percentage-point impact on child penalties for women who became eligible to 12 months of maternal leave. Panel A shows the effects for employment, while panel B shows the effects on earnings. The results indicate that the qualitative effects of the 1961 reform are strikingly similar to those observed for the other parental leave reforms. There is a sharp negative effect on both employment and earnings, concentrated in the first 12 months after birth, i.e. in event years 0 and 1.²⁴ These negative effects quickly disappear as women exhaust their maternal leave entitlement; we observe no significant effect in the long run.

The causal interpretation of these results relies on the assumption that there is no confounding trend in child penalties by birth cohort. To validate this assumption, Appendix Figure B.VI shows placebo results from specification (4) estimated on a sample of mothers who gave birth in 1962 ($T = 0$) or 1963 ($T = 1$). This figure corroborates our identifying assumption, showing that there is no significant trend in child penalties across birth cohorts.²⁵

3.3 Why Does Parental Leave Not Hurt Mothers' Careers?

The lack of systematic evidence on the long-run career effects of parental leave has generated mixed views among scholars. On the one hand, some have voiced concerns that generous parental leave schemes, by keeping mothers out of the labor market for too long, may permanently hurt their careers. This view is predicated on findings that, in

²⁴For women who gave birth between January 1961 and December 1961, the period of the first 12 months after birth spans between the beginning of event year 0 (1961) and the end of event year 1 (1962).

²⁵In a similar vein, we implemented doughnut RD specifications excluding births in 1960, to flexibly account for trends in outcomes by birth cohorts, and found similar results.

the short run, parental leave schemes induce women to stay out of the labor force and reduce labor market earnings (e.g. [Lalive and Zweimüller 2009](#), [Rossin-Slater, Ruhm and Waldfogel 2013](#), [Baum and Ruhm 2016](#)). These short-term work interruptions may translate into longer-term effects through experience channels (such as human capital accumulation or signaling) or persistent preferences changes. On the other hand, others have argued that job-protected parental leave fosters female labor market attachment, thus leading to positive long-run effects on the career trajectories of mothers. This view has been challenged by findings in [Ejrnaes and Kunze \(2013\)](#), [Schönberg and Ludsteck \(2014\)](#), [Dahl, Løken, Mogstad and Salvanes \(2016\)](#), and [Bailey, Byker, Patel and Ramnath \(2019\)](#), who find either small or negative effects of parental leave schemes in the longer run.

Our analysis of 50 years of parental leave reforms in Austria offers an unambiguous answer: parental leave policies have no significant effects on the labor market outcomes of mothers in the long run. But why? And how does this fit with the existing evidence?

Non-Linear Effects of Parental Leave Extensions First, our study of Austrian reforms allows us to shed light on non-linearities in the effects of parental leave policies, which are usually held responsible for the differences in estimates across contexts and reforms (see [Rossin-Slater 2017](#)). One key source of non-linearity stems from the interaction between the duration of job protection and the duration of cash benefits ([Lalive, Schlosser, Steinhauer and Zweimüller 2013](#)). As discussed in [Olivetti and Petrongolo \(2017\)](#), it could be that the availability of some job protection, relative to no protection at all, ensures continuity of employment and discourages transitions out of the labor market, while further extensions simply delay return to work without further gains in employment and earnings.

Using our estimated effects of the 1990, 1996 and 2000 reforms on the average earnings of women in the first five years after birth, and assuming additivity in the effects of marginal extensions, we can compute the effects of 6 months increments in parental leave duration, keeping job protection fixed at 24 months. Results, reported in [Figure 5](#) confirm that the effects of PL extensions are indeed highly non-linear, and strongly decrease with baseline duration. Extending the duration of paid leave from 12 to 18 months has a larger negative effect than extending duration from 18 to 24 months. Furthermore, extending the duration of leave from 24 to 30 months, that is, beyond the duration of job protection, has no effect on female outcomes. The latter result happens to be entirely driven by compliance rates. In other words, there is a strong interaction between job protection and the duration of cash benefits: unless it is covered by job protection, the take up of additional leave by mothers appears close to zero. We also report in [Figure 5](#) the estimate of the 1961 introduction of parental leave, which captures

the effect of a 10 month extension (from 2 to 12 months) of paid parental leave, covered with job protection. The effect is larger than the effect of a 12 month extension from 12-24 months with job protection, confirming again that the marginal effect of parental leave duration decreases significantly as a function of baseline duration.

But interestingly, the strong non-linearity evidenced by Figure 5 only applies to short term effects, and seems entirely driven by decreasing marginal take up rates, as the duration of PL increases. As far as long term effects are concerned, neither the duration of cash benefits nor the presence of job protection seem to make any difference: the long run labor market outcomes of mothers are systematically unaffected by PL policies.

Parental Leave & Experience Effects What could explain the absence of longer term effects of PL policies? In the next section, we investigate the possibility that interactions between institutional child care provision and PL policies drive these results in the Austrian context. But, the absence of any long run dynamic effects on mothers' careers still raises one puzzle: how is it compatible with the presence of significant negative effects in the short run? Experience effects would suggest that spending an additional year out of the labor force would negatively affect future earnings. So are returns to experience in the labor market simply zero for Austrian mothers?

Indeed, a large fraction of mothers work in relatively low-skill occupations in Austria, where experience effects have been estimated to be small. But interestingly, Appendix Figure B.VII shows that even among mothers in the top quartile of the pre-birth earnings distribution, PL reforms have no long term effects. A more natural explanation is that our estimates identify *local average treatment effects* among the compliers to the reforms. And compliance is likely to be higher among mothers who do not incur a large cost on their career for taking additional leave. In other words, this suggests that, when they make their parental leave decisions, individuals do internalize the long term consequences of such choices on their careers.

The Role of Subsequent Fertility Our estimates of the effects of PL policies do not control for subsequent fertility. In practice, fertility could be a significant mediator of the long term effects of parental leave on child penalties, and contribute to mask other dynamic effects of the policy.²⁶

To investigate the mediating role of subsequent fertility, we estimate the dynamic effects of PL reforms on the restricted sample of mothers with a completed fertility of one. Panel A of Appendix Figure B.VIII plots the results for the 1990 reform, and does

²⁶If, for instance, longer PL duration is associated with lower or delayed subsequent fertility, this could reduce the longer run penalty, thus counteracting potentially negative experience effects of taking a longer maternal leave.

not reveal any difference in dynamic patterns compared to the whole sample of mothers. This suggests that subsequent fertility does not have any significant mediating impact on the earnings' responses to PL extensions. We further confirm this diagnostic by plotting in Panel B the average completed fertility (measured as of 2017) of mothers, by bins of the weekly birth date of their first child, in an eight-months window around the cutoff date of the 1990 PL reform. The absence of any discontinuity at the 1st of July cutoff indicates that eligibility to the 24-months parental leave did not, in fact, have any significant effect on completed fertility.²⁷

We should nevertheless highlight that women who give birth to their first child on both sides of the reform cutoff dates face different PL regime for the first birth, but face on average similar PL regimes for any subsequent birth. As a consequence, our estimates so far have only identified how child penalties respond to the PL regime applicable *to the first child*. To compute the total effect on child penalties of moving, in steady state, from one PL regime to another, one needs to account for the effect of PL extensions for births of higher parity. If this effect is significant, subsequent fertility will mechanically increase the total effect of PL extension on child penalties. To illustrate this, we turn to mothers giving birth to their second, third and fourth child, and replicate, for each parity, our dynamic estimates of the impact of PL reforms on earnings. We then construct the total steady state impact of a reform by adding dynamic estimates at all parities, weighted by the average completed fertility, and average timing between parities, of women who had their first birth around the time of each reform. Results are reported in Appendix Figure B.IX. We find that accounting for the effects of PL reforms for all subsequent births only slightly increases the impact of parental leave in the medium run. But we still cannot detect any significant effect of parental leave in the long run. Overall, even when considering their steady state impact, we find that parental leave reforms have had only small effects on the careers of mothers.

²⁷These results are reminiscent of [Dahl, Løken, Mogstad and Salvanes \(2016\)](#) who, in the context of Norway, also find that being exposed to a longer parental leave for the first child has no effect on subsequent fertility. [Lalive and Zweimüller \(2009\)](#), studying the Austrian 1990 and 1996 reforms, find significant effects on fertility timing and argue this is consistent with incentives created by the parental-leave-renewal rule: to qualify for a new parental leave, mothers giving births to a new child – while still on parental leave for their previous child – are not subject to the previous-work requirement. In line with [Lalive and Zweimüller \(2009\)](#), we find a significant impact on the *timing* of subsequent births but no effect on completed fertility. Moreover, the fertility effects are quantitatively too small to alter our conclusions on the impact of parental leave rules on long-run child penalties. We do not find any impact on fertility of the 2000 parental-leave expansion that abolished the previous-work requirement. This is consistent with the idea that fertility timing responses of previous reforms were driven by the renewal rule.

4 The Impact of Child Care Provision on Child Penalties

Over the past fifty years, while parental leave generosity was drastically expanded in Austria, so was the provision of public child care. Sixty years ago, nurseries were a rarity, and the majority of Austrian municipalities did not have a kindergarten. The fraction of children between the age of 3 and 5 who were enrolled in a kindergarten was lower than 40% in 1970. Today it is above 90%. How did this expansion of child care provision affect the careers of women relative to men? And how did the effects interact with the expansion of parental leave generosity?

4.1 The Expansion of Child Care Provision

Measuring Child Care Provision

We start by constructing long time series of child care provision for each Austrian municipality. The central government collects, in its supervisory role, granular information on all nurseries & kindergarten for each municipality and year. The reports, available since 1988, contain information on the precise location of each child care institution, whether it is public, private or operated by the church, its opening hours, its number of certified teachers and their contracts (part-time vs full time), as well as the legal maximum number of children per certified teacher (which varies by type of institution and is subject to state legislation). The data also reports the number of children attending each institution, by age, and by number of hours of attendance. We convert this rich information into two indices that capture the level of child care coverage for children of age 1 to 2 (nursery care index), and for children of age 3 to 5 (pre-school care index), in each municipality and year.

$$\text{Index 1-2} = 100 \times \frac{\# \text{ FTE Child Care Spots for Children Age 1-2}}{\# \text{ Children of Age 1-2}}$$

$$\text{Index 3-5} = 100 \times \frac{\# \text{ FTE Child Care Spots for Children Age 3-5}}{\# \text{ Children of Age 3-5}}$$

To get the numerator of index 1-2 (resp. index 3-5), we first multiply the number of full time equivalent teachers working in each nursery (resp. kindergarten) by the legal maximum number of children that a single teacher can supervise in this institution. We then weight this number by the opening hours of the institution: a fulltime kindergarten receives a weight of 1 for instance, while a half-day kindergarten receives only weight of 0.5. We then sum these numbers across all nurseries (resp. kindergarten) in each municipality to obtain the number of full time equivalent child care spots available

for children of age group 1-2 (resp. of age group 3-5). We then divide the numerator by the number of children of that age group in the municipality. This gives us a direct measure of the probability that a child of a given age range has access to a full time spot in a child care institution, for each municipality and year.²⁸ Finally, we also combine the information from our two indices into an aggregate index of child care, that we call Index 1-5:

$$\text{Index 1-5} = 100 \times \frac{\# \text{ FTE Spots for Age 1-2} + \# \text{ FTE Spots for Age 3-5}}{\# \text{ Children of Age 1-5}}$$

Appendix Figure C.I provides evidence of the massive expansion of child care in Austria since the 1980s. It displays the evolution of the nursery care and pre-school care indices computed for Austria as a whole. Panel (a) focuses on the nursery care index. It shows that the coverage of nursery care was extremely low in 1988: the probability that a child between the age of 1 and 2 had access to a full time spot in nursery care was just 5% then. The index increased significantly after that, especially during the 2000s. But despite its 7-fold increase, the index remained at a relatively low absolute level, below 35, in 2015. Panel (b) shows that the coverage of pre-school care is generally much higher. It started in 1988 from a much higher level (around 55) and expanded quickly throughout the 1990s and early 2000s, to reach almost full coverage today.

4.2 The Effect of Local Child Care Expansions on Child Penalties

To assess the impact of child care provision on the dynamics of female labor market outcomes, we start by comparing the child penalties of women who give birth in municipalities with high vs low child care coverage. We rank municipalities according to the value of their aggregate child care index (Index 1-5) in 1990 and split the sample between municipalities below vs above the median. We then compute for each subsample the child penalty for mothers whose first child was born between July 1990 and June 1996, i.e. under the two years parental leave regime. To control for differences in pre-birth characteristics of mothers across the two subsamples, we use inverse-probability-weighting based on pre-birth earnings, pre-birth employment, and age.²⁹ Results, reported in Figure 6, show that child penalties are significantly lower for mothers in high

²⁸Note that we topcode each index at 100, as some communes offer spots in childcare for more than the number of children in their own municipality.

²⁹Concretely, we predict, using a probit model, the probability π_i of being in a municipality with above median child care provision. We use as predictors the number of days of employment in the year before birth, a set of dummies for the age of the mother at first birth, and the log of individual earnings plus one, $\log(y_{it} + 1)$, for event years $t = -5$ to $t = -1$. Based on these probits, we predict the probability that a municipality has above median child care provision in 1990, and we weight mothers in municipalities with above median coverage by $E(\pi_i) / \hat{\pi}_i$ whereas the weight for mothers in municipalities with below median index is $E(1 - \pi_i) / (1 - \hat{\pi}_i)$.

coverage municipalities: their earnings penalty is 5.11 percentage point lower on average over the first five year after birth. Furthermore this effect seems to be quite persistent and stable over the longer run. If we scale this difference in earnings penalty by the difference in child care coverage across the two groups, we obtain a large and significant TOT effect of .158 (.002), suggesting that increasing child care coverage from 0 to 100% would translate into a 16 percentage point reduction in the child penalty.

To interpret this evidence as causal, however, one would need to assume away selection on unobservables characteristics correlated with child penalties. In practice, women may decide to live in municipalities with high child care coverage precisely because they would like to invest in their career after the birth of their child. A closer look at the geographical distribution of child care in 1990 confirms that selection is likely to be significant. Appendix Figure C.II, which provides a heat map of the child care index in 1990, shows a stark contrast between very high coverage rates in Vienna and its surrounding municipalities, and very low coverage rates in the rest of Austria. This strong urban/rural divide suggests that women in the high and low coverage municipalities of Figure 6 were likely to evolve in radically different environments in terms of preferences and gender norms.

Identification Strategy: Using Local Expansions

While child care was extremely segmented geographically in 1990, the rest of Austria did eventually catch up with Vienna, unleashing the massive expansion of child care provision observed in the 1990s and 2000s. But, interestingly, when catching up with Vienna and its surroundings, Austrian municipalities did so at very different rates. Figure 7 maps the change in child care index for each municipality between 1990 and 2000 in panel (a), and between 2000 and 2010 in panel (b). It reveals a tremendous amount of spatial heterogeneity in the intensity and timing of the expansion. Some municipalities saw a rapid increase in child care provision in the 1990s, while their neighbors or other geographically similar municipalities only expanded child care much later in the 2000s. Some municipalities saw a steep increase, while others did not expand child care provision much.

This large spatial variation in local child care expansion represents a promising source of identification. Conceptually, we want to capture the macro effects of child care expansion on the dynamics of female outcomes, accounting for potential equilibrium effects in the labor market or for endogenous changes in preferences and norms, in the spirit of [Duflo \(2001\)](#). To the extent that local labor markets are not fully integrated and that treatment effects do not diffuse much across space, spatial variation in child care expansion across municipalities can meaningfully capture macro effects and is there-

fore appealing from an identification standpoint.³⁰

However, spatial variation in child care expansion is also likely to be endogenous to the evolution of female labor market outcomes. For instance, increases in local female employment may push municipalities to expand their supply of institutional child care. To deal with this concern, we isolate episodes of large and sudden increases in child care provision at the municipal level. We focus on events where the nursery care index or the pre-school care index of a municipality increases by more than 20 points in one single year. These events are driven by large supply shocks such as the construction of a new facility, or the expansion of existing ones and the hiring of new teachers. We define a municipality as “treated” when it experiences one such episode. Municipalities that never experienced an increase of 20 points or more in their child care index constitute our control group. Our strategy relies on comparing the outcomes of women in treated municipalities, before and after an episode of child care expansion, to those of women in control municipalities. To probe whether treated and control municipalities are well balanced across observables dimensions, Appendix Table C.I provides descriptive statistics by treatment status, and reveals that prior to local childcare expansion, both groups of municipalities are remarkably similar in terms of population structure and labor market outcomes.

In Figure 8, we start by focusing on the evolution of child care provision and of the take-up of institutional care, which constitute the first stage of our quasi-experimental design. Panel (a) plots the evolution of the nursery care index around the time of a nursery care expansion separately for treated and control municipalities. The graph shows that the increase of nursery care provision in treated municipalities has four important characteristics: it is unanticipated, sudden, large, and persistent. By unanticipated, we mean that the index was extremely stable in the years prior to the event, and we cannot detect any significant differences in pre-trends between treated and control municipalities. By sudden, large and persistent, we mean that the year of the event is characterized by a surge of 40 points in the value of the index in treated municipalities, and the index remains stable at this high level thereafter. For control municipalities we do not see any specific trend upward in nursery care provision: the nursery care index remains stable throughout, at a low level, around 8.

One concern with our quasi-experimental design is that child care expansions may be endogenously driven by the dynamics of local demand. An increase in the demand for child care may, for instance, create a shortage in supply, prompting a decision by local authorities to expand the provision of institutional child care. Panel (b) of 8 dis-

³⁰Evidence from [Lalive, Landais and Zweimüller \(2015\)](#) indeed suggests that local labor markets are quite geographically segmented in Austria. Studying the response to a large region-specific UI extension program for older workers, they find significant equilibrium effects in treated regions, but very limited geographical spillovers to neighboring municipalities in untreated regions.

plays the evolution of the uptake of nursery care around an expansion and alleviates these endogeneity concerns: the graph does not reveal any differential pre-trend in take up between treated and control municipalities. Furthermore, the panel shows that the dynamics of take up closely mimics the dynamics of child care provision. The large increase in the provision of nursery care in treated municipalities translates into a large increase in the uptake of nursery care, which is also sudden and persistent. On average, the supply shocks we focus on are accompanied by an immediate trebling of the fraction of children going to nursery care, from 8% to 28%. The implied take up rate of the additional full time equivalent spots created is a bit below 50%.

Panel (c) and (d) replicate the same analysis for sudden expansions in the provision of pre-school care. Both panels reveal similar dynamic patterns as those observed for nursery care expansions. The main important difference relates to the baseline levels of the control municipalities. In the case of nursery care expansions, control municipalities are at a lower baseline level than treated municipalities, the latter getting ahead of the curve in terms of child care provision. To the contrary, in the case of pre-school care expansions, treated municipalities were originally lagging behind, and are using expansions to catch up with control municipalities.

To identify the effects of the large expansions of child care provision presented in Figure 8, we use a difference-in-difference framework. We compare the evolution of the earnings of mothers in treated municipalities ($T_i = 1$) and control municipalities ($T_i = 0$) in the years s around a child care expansion (we normalize the year of the expansion event to $s = 0$). Because we want to map out the effects of child care on the full dynamics of labor market outcomes around child birth, our outcomes of interest are the earnings Y_{ist} of mothers who are observed exactly t years after/until the birth of their first child. For each t , we are interested in the following model:

$$Y_{ist} = \alpha_t D_{ist}^{Exp} + \alpha_t^T D_{ist}^{Exp} \cdot T_i + \beta D_{ist}^{Age} + \gamma D_{ist}^{Year} + v_{ist} \quad (5)$$

where D^{Exp} is a vector of event time dummies with respect to child care expansion. Note that we control, as before, for calendar year fixed effects and age fixed effects. Identification relies on the usual parallel trend assumption that absent child care expansion, the evolution of earnings of mothers observed t years after birth would have been the same in treated and control municipalities.

For the purpose of estimation, we stack the difference-in-difference models (5) for each event time t relative to child birth, into one unique specification:

$$\begin{aligned}
Y_{ist} = & \sum_j (\alpha_j D_{ist}^{Exp}) \cdot \mathbf{1}[D_{isj}^{Event} = 1] + \sum_j (\alpha_j^T D_{ist}^{Exp} \cdot T_i) \cdot \mathbf{1}[D_{isj}^{Event} = 1] \\
& + \beta D_{ist}^{Age} + \gamma D_{ist}^{Year} + v_{ist}
\end{aligned} \tag{6}$$

This stacked model allows us to calculate counterfactual earnings \tilde{Y}_{ist} , that is, earnings without child penalty. We then standardize the coefficient estimates $\hat{\alpha}_{st}$ and $\hat{\alpha}_{st}^T$ by $E[\tilde{Y}_{ist} | t]$. To get a visual representation of our difference-in-difference design, we finally plot, for given values of t , the standardized coefficient estimates $\hat{\alpha}_{st}$ and $\hat{\alpha}_{st} + \hat{\alpha}_{st}^T$ in a ten year window around the expansion event ($s = -5$ to $s = 5$). In this graphical representation, the coefficients $\hat{\alpha}_{st}$ capture the effect on the child penalty, t years after birth, of being in a control municipality, as a function of time s relative to child care expansion. And the corresponding effect in treated municipalities is given by $\hat{\alpha}_{st} + \hat{\alpha}_{st}^T$. When presenting our results, we set $s = -1$ as our baseline category, and therefore express all results relative to the year prior to child care expansion. To account for spatial correlation in unobserved labor market shocks, we cluster standard errors at the municipal level.

The Limited Effects of Large Local Child Care Expansions

We start in Figure 9 by presenting the results for nursery care expansions. In panel (a), we focus on women who are one year before birth ($t = -1$), and report their estimated earnings in control municipalities vs treated municipalities, from five years before, to five years after the expansion shock. The graph first shows no evidence of a differential trend in earnings before expansion. This helps alleviate concerns that the supply shock was an endogenous response to the evolution of female labor market outcomes in treated municipalities. Secondly, the graph reveals that the expansion had no effect on the outcomes of women prior to giving birth. In other words, the increased access to child care did not lead to anticipatory effects among women about to have children.

In panel (b), we turn to the child penalty one year and two years after birth, which corresponds to the period when children are enrolled in nursery care. As in panel (a), and in support of our identification strategy, we do not detect any sign of differential trends before the expansion. But strikingly, neither do we find any significant difference in the child penalty of women in treated municipalities *after* the expansion. Despite a large increase in the take up of nursery care, the evolution of earnings one to two year after birth remains virtually identical in treated and control municipalities. The corresponding diff-in-diff estimate indicates that nursery care expansions decreased the average penalty in the first two years after birth by 1.1 (1.4) percentage points. If

we scale this estimate by the change in nursery care index, we find that increasing municipal nursery care coverage from 0 to 100% would reduce the child penalty by only 3.1 (4.0) percentage points.

To further corroborate this surprising null result, we replicate diff-in-diff graphs similar to panel (a) and (b) for child penalties at every horizon t , from 3 years before birth to 5 year after birth. The results, reported in Appendix Figure C.III confirm that the large municipal expansions of nursery care had no detectable effect on child penalties. For every single panel, we find that the child penalty estimates for treated and control municipalities closely line on top of each other. In the last panel, we plot the evolution of the average penalty in the 5 years following the birth of a child. These estimates similarly imply that an increase in coverage from 0 to 100% would only reduce the child penalty by 3.0 (4.3) percentage points in the first five years after birth.

One concern is that the presence of significant frictions in the labor market, or slow-moving gender norms may significantly hinder the response to nursery care expansion in the short run. But interestingly, results are identical if we look at longer run effects. In panel (c) of Figure 9, we plot for all event times t the coefficients $\hat{\beta}_{st}$ evaluated at $s = 10$, ten years after the expansion. These estimates capture the long run causal effect of nursery care expansions on the full dynamics of the child penalty. Note that panel (c) is therefore the mirror image, for nursery care reforms, of panels (a), (c) and (e) of Figure 3 in the case of parental leave reforms. The graph shows without ambiguity that, even in the long run, these large increase in the supply of municipal nursery care have had no significant effect on child penalties. Based on these estimates, we draw in panel (d) counterfactual child penalties under the pre and post nursery care expansion regimes: the two profiles are indistinguishable.

One potential explanation for the absence of effects of these large increases in the local provision of nursery care is the availability of long maternal leaves during our period of analysis. Mothers can enrol their child in nursery while still being on maternal leave. The possibility to combine instutional care with maternal leave is no longer available when looking at pre-school years, when children are between 3 and 5 years old. Yet, Figure C.IV, shows that episodes of large pre-school expansions had no significant impact on child penalties either. Panels (a) and (b) show that the earnings, be it just before birth, or 3 to 5 years after birth, evolved in the exact same way in treated and control municipalities. Looking in panels (c) and (d) at the long run causal impact, ten years after the expansion, we find that child penalties were not significantly impacted by large increases in the local supply of pre-school care.

4.3 Interaction Between Child Care and Parental Leave Policies

Since child care provision is not found to significantly affect child penalties, this implies that there should be no significant effects on female earnings arising from the interaction between child care supply and parental leave policies. To confirm this, we look in Appendix Figure C.V at the effects of the 1990 extension of parental leave, comparing women living in municipalities with high vs low levels of child care coverage. We first split municipalities according to the value of their aggregate index of nursery and pre-school care (index 1-5) in 1990. Panel (a) shows that the estimated effect of the reform on child penalties is identical for women in municipalities with above vs below median levels of child care provision. In panel (b), we focus specifically on the provision of nursery care and in panel (c) on the supply of pre-school care: in both cases, we do not detect any significant differences in the impact of the parental leave reform between municipalities with high vs low provision of child care.

4.4 Why Does Child Care Not Improve Mothers' Careers?

The existing literature on the impact of child care policies on mothers' labor market outcomes has produced surprisingly mixed results. Some studies estimate positive and sometimes large effects of child care subsidies on female labor supply (e.g., [Baker, Gruber and Milligan 2008](#); [Berlinski and Galiani 2007](#); [Cascio 2009](#); [Eckhoff Andresen and Havnes 2018](#)). These micro studies are corroborated by cross-country evidence. At the same time, a number of other studies provide a much less optimistic view on the potential for child care provision to improve female labor supply. Studies from the U.S. ([Fitzpatrick 2010, 2012](#); [Barua 2014](#)), Norway ([Havnes and Mogstad 2011](#)), France ([Goux and Maurin 2010](#)), and Spain ([Nollenberger and Rodríguez-Planas 2015](#)) find that child care expansions did not generate any significant improvements in maternal employment. Or if they did, the effects were small and typically limited to single mothers. Our analysis of very large expansions of heavily subsidized child care in Austria shows a precisely estimated zero effect on mothers' careers, in the short run as well as in the long run.

But why are large increases in the take-up of child care not conducive to increasing female labor supply? Uncovering the mechanisms behind this null effect is critical for understanding the ability of family policies to shape the evolution of gender inequality.

Crowding-Out, Frictions & Preferences Two main mechanisms may explain our findings. First, institutional child care may simply crowd-out other child care arrangements, such as care by relatives for instance. Second, mothers may still face large resource and utility costs of increasing labor supply, even when child care becomes freely

available. These high costs may come from the existence of frictions or constraints in the labor market, but may also be due to preferences or gender norms regarding maternal care.

To gain further insights on mechanisms, we use two large external surveys that were conducted with the Austrian Microcensus in 1995 and 2002, and which contain household level information on time use, preferences and child care arrangements. To investigate the correlation between these outcomes and child care provision, we matched this survey information with our measures of institutional child care provision at the political district level.³¹

To increase statistical power, we pool observations from both waves of the survey, and control for differences in the education and age structure across political districts. In practice, we regress the average outcome in a political district on the fraction of the population in three education bins and five age bins. We then plot a binscatter of the residuals, rescaled at the average level of education and age in Austria, against the value of the child care index (index 1-5).

To validate this cross-sectional design, we start by looking at the fraction of children in institutional care, and the employment level of mothers. For these two outcomes, the correlation with the child care index obtained in this cross-sectional design can be directly compared to the causal estimates from our local expansion design. Results, displayed in panel (a) and (b) of Figure 11, suggest that the correlations closely match our quasi-experimental estimates. For take-up, we find that a 1 percentage point increase in the child care index translates into a .36 (.09) percentage point increase in the fraction of children attending institutional care. This is remarkably close to the corresponding estimate from our local expansion design: when pooling all nursery and pre-school expansion episodes together, we indeed find that the take-up rate of an exogenous increase in child care index 1-5 is .35. Similarly, both the cross-sectional and local expansion designs are congruent in finding no significant correlation between child care coverage and mothers' probability of employment.

After having established the internal validity of this cross-sectional design, we then turn to the correlation between institutional care provision and child care arrangements. Panel (c) of Figure 11 displays the average number of hours that mothers spend weekly on child care. It shows first that the prevalence of maternal care is extremely high, each mother spending on average 40 hours a week taking care of their children. Second, we find no sign that the availability of institutional care decreases the pervasiveness of maternal care: mothers spend roughly the same amount of time on child care whether the child care index in their political district is below 30 or above 70. If

³¹The surveys are unfortunately not representative at the municipal level, but are representative of the 95 Austrian political districts ("*politischer Bezirk*"). These districts correspond to the second most granular level of administrative authority after municipalities.

higher child care provision does not affect maternal care, panel (d) shows that it does crowd-out almost one for one alternative care by relatives: the fraction of children that are under the care of relatives decreases by .35 (.11) percentage point for every percentage point increase in the child care index. This large substitution effect clearly suggests that working mothers predominantly use relatives as a child care alternative, and switch to nurseries and pre-schools when institutional child care becomes available.

But what about women who do not work: are they facing constraints preventing them from increasing their labor supply? In panel (e), we correlate child care provision with the fraction of mothers reporting that their career is constrained due to a lack of appropriate child care. Interestingly, only a small fraction of surveyed women say that they do feel constrained by the supply of institutional care. And what is more, the fraction is not larger in political districts with very low levels of child care provision. We find no significant correlation between child care availability and the experience of child care constraints.

So if child care constraints do not refrain mothers from investing in their careers, what does? Evidence in panel (f) strongly points to the crucial role of preferences and norms regarding maternal care: the overwhelming majority (i.e. 70 to 80%) of women report that they do not work because they do want to take care of their children. Furthermore, we find only a marginally significant negative correlation between these reported preferences and the level of child care provision, showing little signs that expanding access to institutional child care is shifting these strong norms around the role of maternal care.

Overall, these results suggest that the combination of cheap available care by relatives and strong preferences for taking care of children has prevented even very large expansions of child care provision from improving the labor market outcomes of Austrian mothers.

Fertility Effects To fully understand the impact of child care on the careers of women, we finally need to account for the possibility that the expansion of child care provision affected fertility behaviors. If the increased availability of child care induced women to have more children, this may have mechanically increased the child penalties on their careers. On the other hand, the development of cheap institutional child care in the last sixty years has been accompanied in Austria by a steady decline in fertility rates. Is there any causality behind this correlation?

Unfortunately, the existing literature offers little guidance to answer this question. There are remarkably few studies investigating the issue and therefore little consensus on the responsiveness of fertility decisions to child care provision. Interestingly, our quasi-experimental set up allows us to identify the effect of municipal child care expansions

on local fertility. To this effect, we measure local fertility rates as the total number of births divided by the total number of women aged 15 to 49 in a municipality. In Appendix Figure C.VI, we plot the evolution of fertility rates in treated vs control municipalities, around the time of large shocks in child care supply. Looking at the long run effects, ten years after the expansion, we find a very small, and insignificant, negative effect of child care on fertility, for both nursery and pre-school care expansions.

5 How Much of the Evolution of Gender Inequality Can Family Policies Explain ?

Equipped with causal estimates of the effects of all family policy reforms in Austria over the last sixty years, we can now measure how the dramatic expansion of family policies has shaped the evolution of gender inequality. To do this, we proceed in two steps. First, we use the estimated child penalties to decompose gender inequality into child-related and residual inequality. We can then, in a second step, leverage our estimates of the effects of family policies on child penalties to simulate gender inequality under alternative policy scenarii. This will uncover how family policies have affected the long-run historical evolution of gender inequality.

5.1 Child-Related Gender Inequality Over the Last 60 Years

The consistency and quality of the Austrian administrative data allows us to reconstruct the evolution of gender gaps in earnings since the beginning of the 1950s. We combined for this purpose the digitized ASSD data, which starts in 1972, and detailed tabulated statistics on the distribution of earnings by gender based on the same underlying ASSD data before it was computerized (*Lohnstufenstatistik*).³² This allows us to measure, for all years s between 1953 and 2017, the total gender gap in raw earnings $\Delta_s \equiv 1 - E[Y_{ist}^w | s] / E[Y_{ist}^m | s]$. Two points should be stressed about our measure of gender gap. First, it does correct for top coding in the data, by interpolating the top tail of the distribution using information from uncensored tax data.³³ Second, we systematically account for the extensive margin of labor supply, by including individuals with zero earnings.³⁴ In that sense, our measure of the gender gap encompasses the broadest possible set of behaviors underlying labor market outcomes.

³²See [Christl and Wagner \(1981\)](#). We provide all the details of the data sources and construction of our gender gap measures in Appendix D.1.

³³In practice, we allow for different Pareto parameters by gender and by age group. Details of our procedure are given in Appendix D.1.

³⁴As explained in Appendix D.1, we made sure that the fractions of men and of women with zero earnings in our data matched the Census data.

For each year, we can then ask how much of the cross-sectional gender gap Δ_s is explained by child penalties, that is, by the differential impact of children on the careers of mothers relative to fathers. To answer this question, we follow a decomposition methodology similar to [Kleven, Landais and Sogaard \(2019\)](#). This strategy relies on using child penalty estimates to get a measure of the counterfactual gender gap $\tilde{\Delta}_s \equiv 1 - E[\tilde{Y}_{ist}^w | s] / E[Y_{ist}^m | s]$. This counterfactual gender gap represents the level of gender inequality in earnings that would prevail if we removed the negative effect of children on the careers of women. Child-related gender inequality is then simply estimated as the difference between the observed and the counterfactual gender gap:

$$\Delta_s - \tilde{\Delta}_s \equiv \frac{1}{E[Y_{ist}^m | s]} \left[\sum_t \psi_{st} \cdot P_{st} \cdot E[\tilde{Y}_{ist}^w | s, t] \right]$$

In the language of Oaxaca-Blinder decompositions, our measure of child-related inequality corresponds to the “unexplained” effects of children and two sets of inputs are necessary to compute it. First, we need a measure of the fraction of women ψ_{st} who are at each event time t relative to the birth of their first child in year s . Second, we need estimates of their child penalty P_{st} and of their corresponding average counterfactual earnings $E[\tilde{Y}_{ist}^w | s, t]$.

We compute the measures ψ_{st} by complementing the ASSD data with information from the birth register and from the Austrian Birth Barometer, to account for women whose first birth is not observed in the ASSD or REV data. To get estimates of child penalties and counterfactual earnings by year, we extend the baseline specification (1) to allow for year-specific coefficients on event time. In practice, we estimate the following model:

$$Y_{ist}^w = \sum_y (\alpha_y^g D_{ist}^{Event}) \cdot \mathbf{I}[y = s] + \beta^g D_{ist}^{Age} + \gamma^g D_{ist}^{Year} + v_{ist}, \quad (7)$$

where we interact the event time dummies with year dummies in order to estimate year-specific event coefficients α_{st}^w . As in our baseline specification (1), we include a full set of age and year dummies in the set of covariates X_k . We can then reconstruct counterfactual earnings for mothers in event time t in year s : they simply correspond to the predicted earnings from the above specification, when omitting the contribution of the α_{st}^w coefficients.³⁵

Figure (12) panel (a) presents the results from our decomposition. The first striking

³⁵Note that estimating event coefficients by calendar year s and event year t amounts to estimating event coefficients by birth cohort $c = s - t$. For births cohorts before the start of our data, we observe only post-event years and are unable to directly estimate child penalties associated with these births. We therefore rely on an extrapolation of the observed penalties in the data that we describe in detail in Appendix D. Overall, results reported in Appendix Figure D.VII show that penalties up to 10 years after birth have been remarkably stable over the past 50 years. Longer run penalties, from 10 to 30 years after birth, to the contrary have declined significantly.

fact is that despite having experienced a significant decline over the last sixty years, of about 30 percentage points, the gender gap in earnings remains very large in Austria: the average woman still earns 40% less than the average man in 2017. Furthermore, the graph shows that the level of the child-related gender gap is also very large: women earn on average around 30 to 35% less than men *because of children*. This naturally reflects the very large child penalties experienced by Austrian mothers, compared to mothers of other developed countries. Interestingly, child-related gender inequality has been quite stable over the past six decades. As a consequence, children went from being responsible for just about half of gender inequality at the beginning of the 1950s, to explaining almost the entirety of the remaining gender gap in Austria today. This also implies that most of the decline in gender gaps in the labor market has been due to factors that are not directly related to children, such as the closing of the educational gender gap.

While the overall level of the gender gap is much higher in Austria than in Scandinavia, we note that results from Figure (12) are strikingly similar to the decomposition evidence of [Kleven, Landais and Sogaard \(2019\)](#) in the context of Denmark. They confirm that child penalties are the main driving force behind gender inequality in developed nations today. We can now turn to exploring how family policies, by affecting these child penalties, have influenced the long run evolution of gender inequality.

5.2 Gender Gaps Under Counterfactual Family Policies

Our focus is on estimating how gender inequality has been affected by the dramatic expansion of three family policies: parental leave duration (τ_{PL}), the coverage of free nursery care (τ_N), and the coverage of free pre-school care (τ_{PS}). This means identifying, for family policy $\tau_k \in \{\tau_{PL}, \tau_N, \tau_{PS}\}$, the causal effect of each policy change $d\tau_k$ on the level of the gender gap: $\frac{d\Delta_s}{d\tau_k}$. Fortunately, we can express this effect as a simple function of the causal effects of the policy change $d\tau_k$ on the full dynamics of the child penalty. Indeed, we have that:

$$\begin{aligned} \frac{d\Delta_s}{d\tau_k} &= -\frac{dE[Y_{ist}^w|s]/d\tau_k}{E[Y_{ist}^m|s]} \\ &= -\frac{1}{E[Y_{ist}^m|s]} \cdot \left\{ \sum_t \psi_{st} \cdot \frac{dP_t}{d\tau_k}(\tau_{PL}, \tau_N, \tau_{PS}) \cdot E[\tilde{Y}_{ist}^w|s, t] \right\} \end{aligned}$$

where we use the empirical fact that earnings of men are affected neither by the arrival of kids, nor by family policies. Motivated by our empirical evidence, the above expression assumes that family policies have no effect on fertility, such that $d\psi_{st}/d\tau_k = 0$. But

we explore below scenarii where we relax this assumption and allow family policies to affect fertility as well. We also note that the effect of any policy change on the child penalty is theoretically a function of *all* family policies. But in practice, we can leverage our empirical evidence which found no interaction effects between parental leave and child care provision. This means that we can consider the impact of a policy change on the child penalty to be solely a function of that particular policy: $\frac{dP_t}{d\tau_k}(\tau_k)$. This, in turn, allows us to easily explore counterfactual scenarii where we change all three policies together.

The scenarii we study involve large, non-marginal, policy changes. Our unique ability to measure their impact on the evolution of aggregate gender inequality relies on two fundamental assets. First, it is made possible by the fact that we have estimated the effect of each policy on the full dynamics of child penalties, that is, we have estimated $\frac{dP_t}{d\tau_k}(\tau_k)$ for all event times t . Second, it relies on the fact that we can account for non-linearity in the effects of these policies. For parental leave policies, we have identified the effect of changing the length of PL benefits at various baseline durations, starting from no parental leave at all. For child care provision, our estimates show that even massive reforms, that increase the level of local coverage by more than 40 percentage points, do not have any significant effects on penalties, even in the long run. We therefore believe that our counterfactual gender gaps are robust to adjustment frictions, and other sources of non-linearities or of delayed timing in the effect of family policies.

Gender Inequality Without Parental Leave & Institutional Child Care Our baseline counterfactual policy scenario consists in getting rid of all family policies. In other words, we simulate the evolution of the gender gap since 1961, when eliminating the effects of all parental leave reforms, and setting the level of nursery care and of pre-school care provision to zero. Results, displayed in panel (b) of Figure (12) convey a clear and striking message: In this hypothetical world, a world that would have experienced absolutely no development in family policies over the past sixty years, the level of gender inequality would be just about the same, or if anything, slightly smaller than what it is today. In 2010 for instance, the gender gap in earnings would have been 41.4% in the absence of any family policy, compared to the observed 43.7%. And the gender gap caused by children would have been 30.1%, instead of 34.7%. This suggests that the massive expansion of family policies has had surprisingly little effect on the long run evolution of gender inequality. If anything, it made things marginally worse.

To offer further insights on this surprising outcome, in Appendix Figure D.VIII, we compute counterfactual gender gaps where we remove each family policy independently. Panel (a) shows what the gender gap would have been in the absence of any parental leave reform. In this scenario, we only undo parental leave expansion, and

assume that child care provision would have followed the same path as actually observed. The graph highlights that parental leave expansions have had a negative, albeit small, impact on gender inequality. In the absence of parental leave, we find that gender inequality in earnings would have been about 3 percentage points smaller today. This result stems from the significant negative impact of parental leave policies on labor supply in the short run. But the absence of long run effects on the careers on women that we documented in section 3 means that the overall negative impact of parental leave policies on gender inequality remains somewhat modest. While parental leave policies do have a small effect on the evolution of the gender gap, we find in panel (b) and (c), to the contrary, that child care policies have none. When undoing these policies, we find that the gender gap would have been absolutely unaffected by the absence of nursery care or of pre-school care.

Of course, the latter results directly stem from our point estimates of the impact of child care on child penalties being close to zero. But how sensitive are our counterfactual gender gaps to the precision with which these null effects are estimated? Can we actually reject that family policies have had large effects on gender inequality? To provide some guidance for inference, we computed two bounding counterfactual gender gaps. The first, “optimistic” one, uses, instead of our point estimates, the upper bound of the 95% confidence interval on $\frac{dP_t}{d\tau_k}(\tau_k)$ for each policy change. The second, or “pessimistic” scenario, uses the lower bound of the 95% confidence interval. The results of these two bounding scenarii are presented in Table 1. The bottom-line is that even in the optimistic scenario, family policies do not cause any substantial reduction in the gender gap. And the difference between the optimistic and pessimistic scenario is relatively minor (.428 vs .400). Our null effects are relatively precise: we can clearly rule out large positive effects of family policies on gender inequality.

Robustness to Equilibrium Effects Our mapping between family policies and the long run evolution of gender inequality rests on our estimates of the effects of these policies on child penalties. It is therefore legitimate to wonder whether these estimates fully account for the mechanisms through which these policies may affect the gender gap.

One potential worry is equilibrium effects in the labor market: do these estimates properly account for all equilibrium effects that may arise from large family policy changes? As far as childcare expansion is concerned, our quasi-experimental strategy is precisely aimed at capturing equilibrium effects by focusing on differential treatment *across* local labor markets. But we found absolutely no evidence of labor supply or earnings responses to childcare provision, and therefore no evidence of equilibrium impacts on local labor markets. For parental leave policies, our quasi-experimental strategy com-

pares treated and non-treated mothers within the same labor market, and therefore does not explicitly capture macro responses. Yet, we find very limited effects of these policies on child penalties, thus greatly limiting the scope for significant equilibrium effects.

Another concern may be the presence of long run effects of these policies operating through slow and gradual changes in preferences and norms. Can our estimates capture them adequately? We note that our difference-in-difference strategy reveals absolutely no long run effects of local child care expansion, up to ten years after they occur. And our investigation in Figure (11) suggests that preferences for maternal care remain extremely strong, and do not react much to the level of child care provision. It is therefore unlikely that our counterfactual exercise misses significant effects of family policies mediated by changes in preferences.

Anticipatory Effects & Fertility Responses Overall, it seems family policies have not shaped the evolution of gender inequality over the last sixty years, because their impact on child penalties has been surprisingly muted. As a consequence most of the significant decline in the gender gap experienced since the 1960s has been due to factors unrelated to children, or to secular changes in choices made before the arrival of kids, such as educational investment. This, in turn, opens two final questions.

First, have family policies possibly affected the choices made by women before the arrival of children, thus contributing to the reduction in gender inequality? When child care became more widely available for instance, did women respond by investing significantly more in their career prior to having children? Could our decomposition therefore miss a key mechanism through which family policies reduce the gender gap? The answer is no. Our analysis of local child care expansion in Figures (9) and (10) precisely shows that large increases in child care provision have no effect on the labor market outcomes of women observed prior to having children. Even ten years after local child care expansions, we found no significant differences in employment or in earnings of women who were about to have kids in treated compared to control municipalities.

If family policies do not generate significant anticipatory effects prior to motherhood, they can still affect the probability of motherhood, as we discussed in sections 3.3 and 4.4. Which raises our second question: how robust are our counterfactual estimates to accounting for fertility responses? What has been the causal effect of the expansion of family policies on fertility? And what were the consequences on the gender gap? After all, the development of family policies in the last sixty years has been accompanied in Austria by a steady decline in fertility rates. Is there any causality behind this correlation? In which case, family policies may have mechanically reduced the gender gap

by decreasing the incidence of child penalties. On the other hand, supporters of family policies often allege that they encourage fertility. This, in turn, would increase the incidence of child penalties, and contribute to widen the gender gap.

There is unfortunately very limited evidence, and no clear consensus, on the impact of family policies on fertility. But as far as child care provision is concerned, we can use the results from our local child care expansions in Appendix Figure (C.VI). There, we found small and insignificant fertility responses. For parental leave, evidence is even more scant. Most studies in the literature have compared, like we did in Appendix Figure B.VIII, the completed fertility rates of women who are exposed, for their first child, to short vs long parental leave durations. But this does not capture the relevant effect on total fertility of expanding the duration of parental leave. The reason is that in such a quasi-experimental set-up, women face the same parental leave regime for any future child they have. Furthermore, one cannot identify from such variation the effect of parental leave on the probability of having a first child. To get at the effect we are interested in, one would need to compare the total fertility of women who are experiencing two different parental regimes *permanently*. The issue is that parental leave policies usually apply to all women in a given country. So this suggests relying on cross-country variation, as in Olivetti and Petrongolo (2017). Correlating parental leave generosity with fertility rates for 22 countries over the period 1970–2010, and controlling for country and year fixed effects, they find a significant, but small, positive correlation.

Using these estimates from Olivetti and Petrongolo (2017) for parental leave, and our estimates of fertility responses to child care provision, we can now simulate counterfactual gender gaps accounting for fertility responses.³⁶ And results reported in Appendix Figure D.X show that fertility responses to family policies are simply too small to make a substantial difference to gender inequality. We estimate, accounting for fertility responses, that the gender gap would be 0.423 today in the absence of any family policy. This is almost undistinguishable from our baseline estimate of 0.414 without fertility responses.³⁷

³⁶We take estimates from Table 3 column (8) of Olivetti and Petrongolo (2017). Appendix D.3 provides all details on our methodology.

³⁷In Appendix Table D.III, we also explore the sensitivity of results to the precision of our estimates of fertility response. We use a methodology similar to Table 1, and compute the gender gap in two scenarios: one using only the upper bound of the 95% confidence interval of our estimates of fertility responses, and the other using only the lower bound of the confidence interval. We find that the range of possible gender gap estimates between these two bounds is very narrow, going from 0.437 to 0.408. We can therefore credibly rule out the presence of significant effects of family policies on gender inequality.

6 Conclusion

In this paper, we have measured the contribution of family policies to the long-run evolution of gender inequality, leveraging historical administrative data and quasi-experimental estimates over more than sixty years. Our approach builds on causally identified estimates of the impact of policy reforms, making it more credible than cross-country approaches to assess the importance of family policies for gender inequality (Ruhm 1998; Blau and Kahn 2013; Olivetti and Petrongolo 2017). By accounting for the effects of all family policy variation, our approach is robust to policy interactions and other non-linearities that normally provides threats to out-of-sample extrapolations of marginal treatment effects. This general approach could be fruitfully extended to other contexts such as, for example, measuring the contribution of tax and transfer policies to income and wealth inequality.

Our results point to a limited role for public policy in explaining the secular decline in gender inequality. Our setting focuses on one particular country, Austria, but their historical expansion of family subsidies is similar to the policy path of other European countries, and while the U.S. has not gone as far as Europe, proposals to expand family policies are the subject of ongoing debate. We find that parental leave policies have had negative short-term effects on female labor market outcomes, but no long-term effects. The short-term effects are relatively small and largely inconsequential for the evolution of aggregate gender inequality. More surprisingly, we also find that publicly provided and heavily subsidized child care has had no effect on the secular decline in gender inequality. This stands in contrast to the strong cross-country correlation between child care subsidization and gender gaps, while within-country quasi-experimental studies of specific reforms have led to mixed results.

Broadly speaking, our evidence is consistent with the ideas exposed in Goldin (2014), arguing that gender convergence relies less on government intervention than on equilibrium features of the labor market. Our work is also consistent with the notion that gender convergence reflects the evolution of gendered preferences and norms (Bertrand 2011; Kleven, Landais, Posch, Steinhauer and Zweimuller 2019; Boelmann, Raute and Schönberg 2020). This underlines the need for furthering our understanding of the formation of preferences and social norms regarding the family-career choices of men and women, and to account for their potential interaction with public policies. Endogenous preferences may create multiple equilibria and tipping points, whereby small changes in policies can have large effects on gender inequality. Conversely, as seems to be the case in Austria, gendered preferences can be sufficiently strong that gender gaps do not respond to even very large expansions of policies.

References

- BAIERL, A. AND M. KAINDL (2011): "Kinderbetreuung in Österreich," Working Paper 77, Austrian Institute for Family Studies, Vienna.
- BAILEY, M. J., T. S. BYKER, E. PATEL AND S. RAMNATH (2019): "The Long-Term Effects of California's 2004 Paid Family Leave Act on Women's Careers: Evidence from U.S. Tax Data," Working Paper 26416, National Bureau of Economic Research, Series: Working Paper Series.
- BAKER, M., J. GRUBER AND K. MILLIGAN (2008): "Universal Child Care, Maternal Labor Supply, and Family Well-Being," *Journal of Political Economy*, 116(4), 709–745.
- BARUA, R. (2014): "Intertemporal Substitution in Maternal Labor Supply: Evidence Using State School Entrance Age Laws," *Labour Economics*, 31, 129–140.
- BAUM, C. L. AND C. J. RUHM (2016): "The Effects of Paid Family Leave in California on Labor Market Outcomes," *Journal of Policy Analysis and Management*, 35(2), 333–356.
- BERLINSKI, S. AND S. GALIANI (2007): "The Effect of a Large Expansion of Pre-Primary School Facilities on Preschool Attendance and Maternal Employment," *Labour Economics*, 14(3), 665–680.
- BERTRAND, M. (2011): "New Perspectives on Gender," in *Handbook of Labor Economics*, vol.4b, ed. by O. Ashenfelter and D. Card, chap. 17, pp. 1543–1590. North Holland: Elsevier Science Publishers.
- BERTRAND, M., C. GOLDIN AND L. F. KATZ (2010): "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2, 228–255.
- BLAU, F. AND L. KAHN (2013): "Female Labor Supply: Why Is the United States Falling Behind?," *American Economic Review: Papers and Proceedings*, 103(3), 251–256.
- BOELMANN, B., A. RAUTE AND U. SCHÖNBERG (2020): "Wind of Change? Cultural Determinants of Maternal Labor Supply," Working Papers 914, Queen Mary University of London, School of Economics and Finance.
- CASCIO, E. U. (2009): "Maternal Labour Supply and the Introduction of Kindergartens into American Public Schools," *The Journal of Human Resources*, 44(1), 140–170.
- CHRISTL, J. AND M. WAGNER (1981): "Geschlechtsspezifische Einkommensunterschiede: Österreich 1953-1979," Forschungsbericht, Institut für höhere Studien (IHS), Vienna.

- CORTÉS, P. AND J. PAN (2020): “Children and the Remaining Gender Gaps in the Labor Market,” Working Paper 27980, National Bureau of Economic Research.
- DAHL, G. B., K. V. LØKEN, M. MOGSTAD AND K. V. SALVANES (2016): “What Is the Case for Paid Maternity Leave?,” *Review of Economics and Statistics*, 98(4), 655–670.
- DUFLO, E. (2001): “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 91(4), 795–813.
- ECKHOFF ANDRESEN, M. AND T. HAVNES (2018): “Child Care, Parental Labor Supply and Tax Revenue,” Discussion paper, Institute for the Study of Labor (IZA).
- EJRNÆS, M. AND A. KUNZE (2013): “Work and Wage Dynamics around Childbirth,” *The Scandinavian Journal of Economics*, 115(3), 856–877.
- FITZPATRICK, M. D. (2010): “Preschoolers Enrolled and Mothers at Work? The Effects of Universal Prekindergarten,” *Journal of Labor Economics*, 28(1), 51–85.
- (2012): “Revising Our Thinking About the Relationship Between Maternal Labor Supply and Preschool,” *Journal of Human Resources*, 47(3), 583–612, Publisher: University of Wisconsin Press.
- GOLDIN, C. (2014): “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, 104, 1091–1119.
- GOLDIN, C. AND L. F. KATZ (2016): “A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation,” *Journal of Labor Economics*, 34, 705–746.
- GOUX, D. AND E. MAURIN (2010): “Public School Availability for Two-Year Olds and Mothers’ Labour Supply,” *Labour Economics*, 17(6), 951–962.
- HAVNES, T. AND M. MOGSTAD (2011): “Money for Nothing? Universal Child Care and Maternal Employment,” *Journal of Public Economics*, 95(11-12), 1455–1465.
- KLEVEN, H. (2014): “How Can Scandinavians Tax So Much?,” *Journal of Economic Perspectives*, 28, 77–98.
- KLEVEN, H. AND C. LANDAIS (2017): “Gender Inequality and Economic Development: Fertility, Education, and Norms,” *Economica*, 84, 180–209.
- KLEVEN, H., C. LANDAIS, J. POSCH, A. STEINHAEUER AND J. ZWEIMULLER (2019): “Child Penalties Across Countries: Evidence and Explanations,” *AEA Papers and Proceedings*, 109, 122–126.

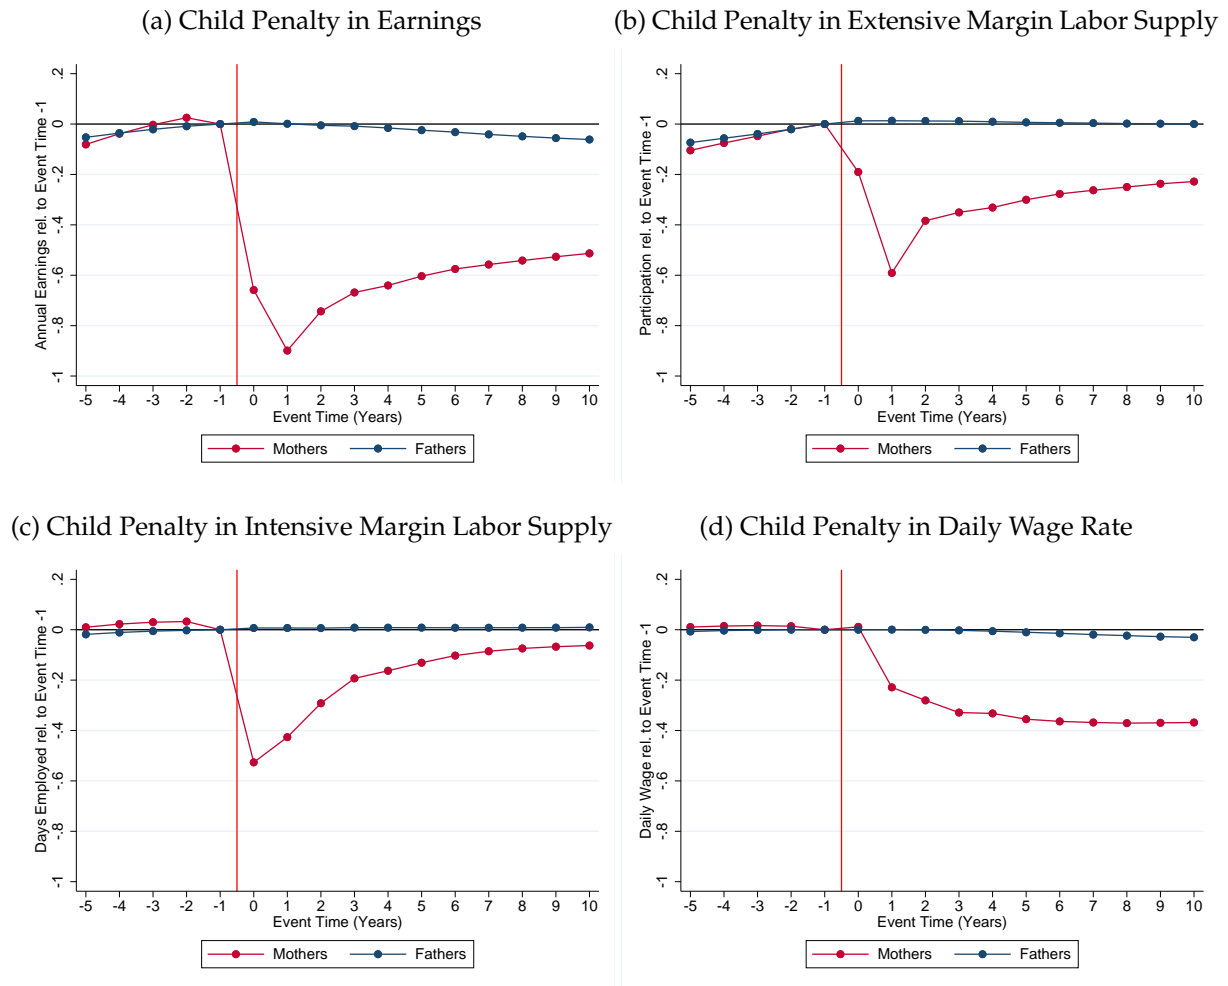
- KLEVEN, H., C. LANDAIS AND J. E. SØGAARD (2019): “Children and Gender Inequality: Evidence from Denmark,” *American Economic Journal: Applied Economics*, 11(4), 181–209.
- (2020): “Does Biology Drive Child Penalties? Evidence from Biological and Adoptive Families,” *American Economic Review: Insights*, p. forthcoming.
- KUZIEMKO, I., J. PAN, J. SHEN AND E. WASHINGTON (2018): “The Mommy Effect: Do Women Anticipate the Employment Effects of Motherhood?,” NBER Working Paper No. 24740.
- LALIVE, R., C. LANDAIS AND J. ZWEIMÜLLER (2015): “Market Externalities of Large Unemployment Insurance Extension Programs,” *American Economic Review*, 105(12), 3564–3596.
- LALIVE, R., A. SCHLOSSER, A. STEINHAEUER AND J. ZWEIMÜLLER (2013): “Parental Leave and Mothers’ Careers: The Relative Importance of Job Protection and Cash Benefits,” *Review of Economic Studies*, 81(1), 219–265.
- LALIVE, R. AND J. ZWEIMÜLLER (2009): “How Does Parental Leave Affect Fertility and Return to Work? Evidence from Two Natural Experiments,” *The Quarterly Journal of Economics*, 124(3), 1363–1402.
- NOLLENBERGER, N. AND N. RODRÍGUEZ-PLANAS (2015): “Full-Time Universal Child-care in a Context of Low Maternal Employment: Quasi-Experimental Evidence from Spain,” *Labour Economics*, 36, 124–136.
- OLIVETTI, C. AND B. PETRONGOLO (2016): “The Evolution of Gender Gaps in Industrialized Countries,” *Annual Review of Economics*, 8, 405–434.
- (2017): “The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries,” *Journal of Economic Perspectives*, 31, 205–230.
- ROSSIN-SLATER, M. (2017): “Maternity and Family Leave Policy,” Discussion paper, National Bureau of Economic Research.
- ROSSIN-SLATER, M., C. J. RUHM AND J. WALDFOGEL (2013): “The Effects of California’s Paid Family Leave Program on Mothers? Leave-Taking and Subsequent Labor Market Outcomes,” *Journal of Policy Analysis and Management*, 32(2), 224–245.
- RUHM, C. J. (1998): “The Economic Consequences of Parental Leave Mandates: Lessons from Europe,” *Quarterly Journal of Economics*, 113(1), 285–317.

SCHÖNBERG, U. AND J. LUDSTECK (2014): "Expansions in Maternity Leave Coverage and Mothers' Labor Market Outcomes after Childbirth," *Journal of Labor Economics*, 32(3), 469–505.

ZWEIMÜLLER, J., R. WINTER-EBMER, R. LALIVE, A. KUHN, J.-P. WUELLRICH, O. RUF AND S. BÜCHI (2009): "Austrian Social Security Database," Working Paper 410, Institute for Empirical Research in Economics (IEW), Zurich.


Figures

Figure 1: Child Penalties



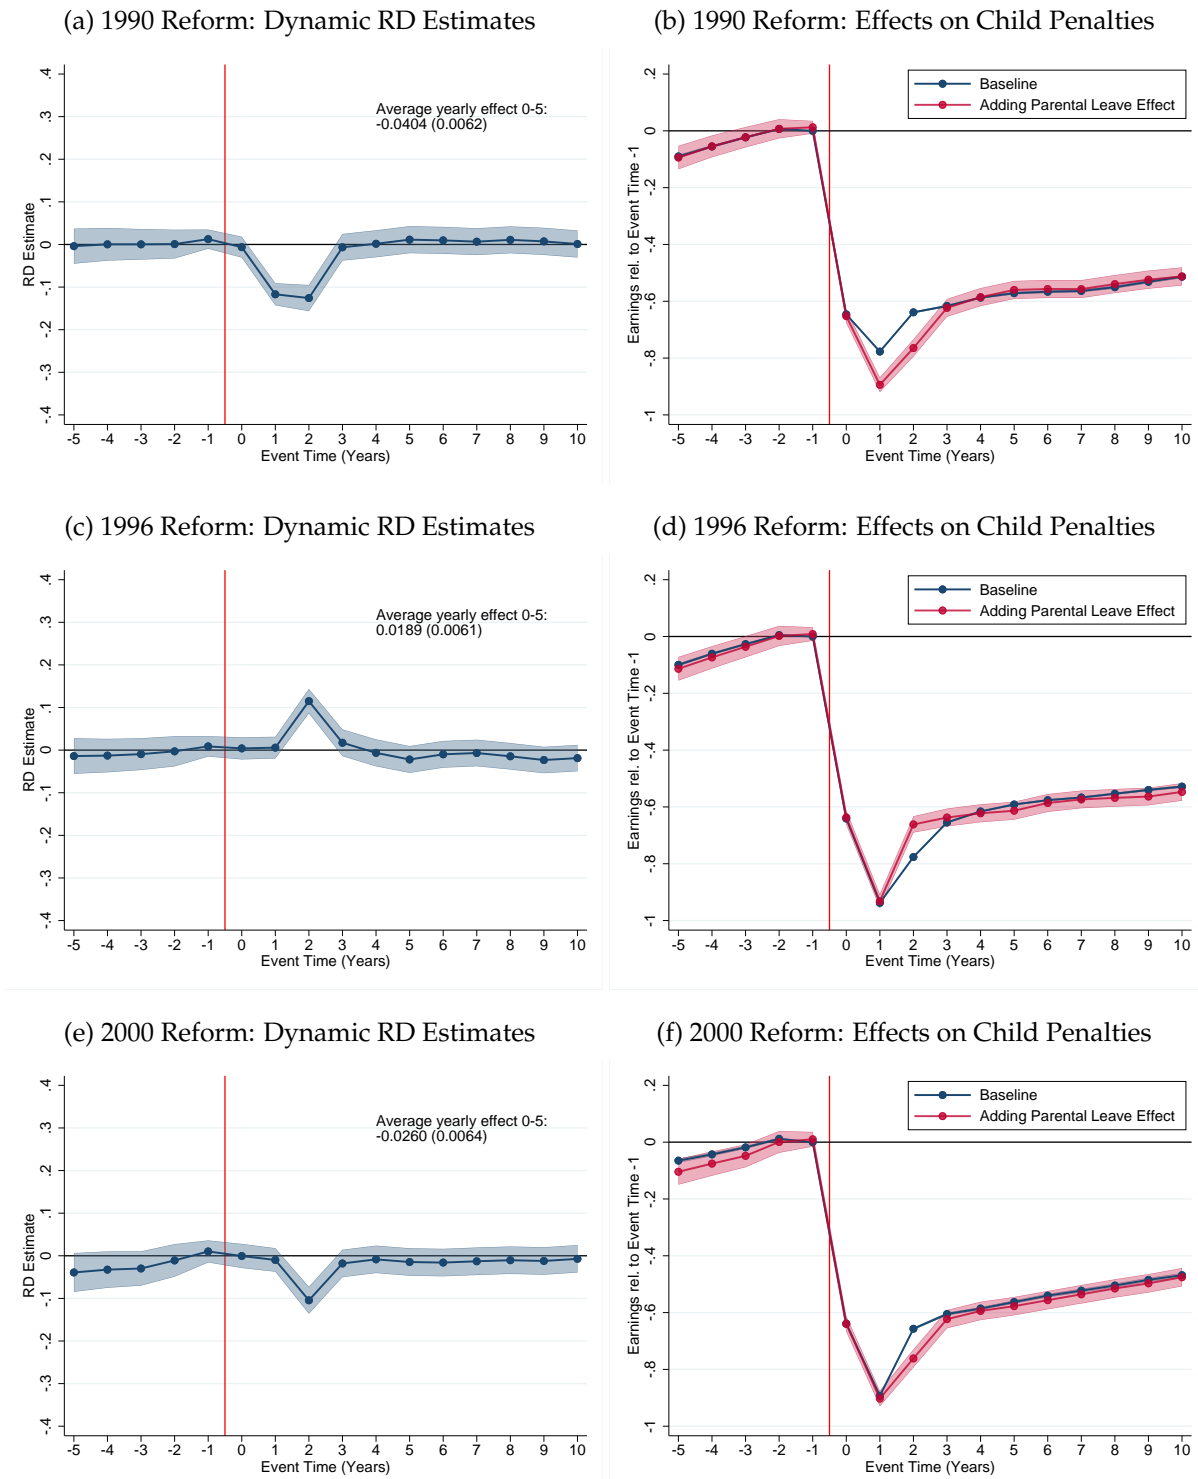
Notes: The figure shows event time coefficients estimated from equation 2 and normalized as explained in Section 2.3 (i.e. we plot $P_t^g \equiv \hat{\alpha}_t^g / E[\tilde{Y}_{ist}^g | t]$) for mothers and fathers and for different outcomes. The sample includes all first births between 1985-2012. Fathers are matched to mothers based on applications for child benefits. In panel (a), the dependent variable is annual earnings (including individuals with zero earnings) converted to real 2000 Euro using the CPI and adjusted for topcoding (see Appendix D.1 for details). In panel (b) the dependent variable is an indicator for whether an individual had any employment spells in a given year. In panel (c), we focus on total days of employment conditional on having at least one employment spell that year. In panel (d), daily wage rate is defined as annual earnings divided by days of employment in that year, and is conditional conditional on having at least one employment spell that year.

Figure 2: The 2008 Parental Leave Regime: Child Penalty by PL Option




Notes: This figure shows normalized event time coefficients ($P_t^g \equiv \hat{a}_t^g / E[\tilde{Y}_{ist}^g | t]$) from child penalty regressions in annual earnings. The sample is mothers and fathers who had their first child in 2008. We separately estimate child penalty regressions for mothers and fathers according to the parental leave option they selected into. The “30+6” option corresponds to the long leave, with up to 30 months for the mother, and 6 months for the father, at low replacement rate. The “20+4” corresponds to the intermediate option of 20 months of PL for the mother and 4 months for the father. The “15+3” option, corresponds to the short parental leave at higher replacement rate. Note that we assign parents to the different options based on the observed length of their effective leave. Shaded areas are 95% confidence intervals on the normalized event time coefficients.

Figure 3: Parental Leave Reforms: Dynamic RD Estimates & Causal Effects on Child Penalties




Notes: The figure reports the estimated effects of the 1990, 1996 and 2000 reforms on child penalties in earnings for women. The left panels plots for each event time the statistic $\hat{\delta}_t / E[\tilde{Y}_{it}^w | t]$ estimated from specification 3, which captures the percentage-point change in the child penalty at event time t caused by the reform. The right panels plot the observed child penalties for women whose first child is born before the reform (baseline), and the counterfactual child penalties that these mothers would have had under the new PL regime, which corresponds to adding the estimates from the corresponding left panel to the baseline penalty. Shaded areas correspond to 95% confidence intervals.

Figure 4: Dynamic Effects of the 1961 Parental Leave Reform




Notes: The figures shows the estimated effects of the introduction of the 1961 reform of parental leave on the dynamics of female earnings (left panel) and female employment (right panel), based on the REV sample of mothers who give birth to their first child in 1959 or 1961 and are eligible for parental leave (52 weeks of employment in the two years before birth). In each panel, we report the estimates $\hat{\alpha}_t^T$ from (4), which correspond to the percentage-point change in the child penalty at event time t for women who become eligible to 12 months of maternal leave. Shaded areas correspond to 95% confidence intervals.

Figure 5: Parental Leave Expansions: Causal Effects by Baseline Duration



Notes: This figure reports the estimated marginal effects of parental leave extensions at different levels of baseline PL duration, on the average yearly child penalty in the first five years following a first child birth. The three estimates on the right of the graph correspond to the effects of 6 months increments in parental leave duration, keeping job protection fixed at 24 months, obtained from our estimated effects of the 1990, 1996 and 2000 reforms in Figure 3, and assuming additivity in the effects of marginal extensions. The estimate on the left corresponds to the estimated effect of the 1961 introduction of parental leave, which captures the effect of a 10 month extension (from 2 to 12 months) of paid parental leave, covered with job protection.


Figure 6: Child Penalty by Level of Child Care Provision
Index 1-5 in 1990




Notes: This figure compares child penalties for women who give birth in municipalities with high vs low child care coverage. We focus on mothers whose first child was born between July 1990 and June 1996. We rank municipalities according to the value of their aggregate child care index (Index 1-5) in 1990 and split the sample between municipalities below vs above the median. We then compute for each subsample the child penalty in earnings following equation 2, as done in Figure 1. To control for differences in pre-birth characteristics of mothers across the two subsamples, we use inverse-probability-weighting based on pre-birth earnings, pre-birth employment, and age. The ITT estimate corresponds to the average difference in penalties in the two groups over the first five year after birth. The TOT estimate corresponds to the ITT estimate scaled by the difference in our index 1-5 of child care coverage across the two groups.

Figure 7: Spatial Variation in Child Care Expansion

(a) Change in Index 1-5 between 1990 and 2000



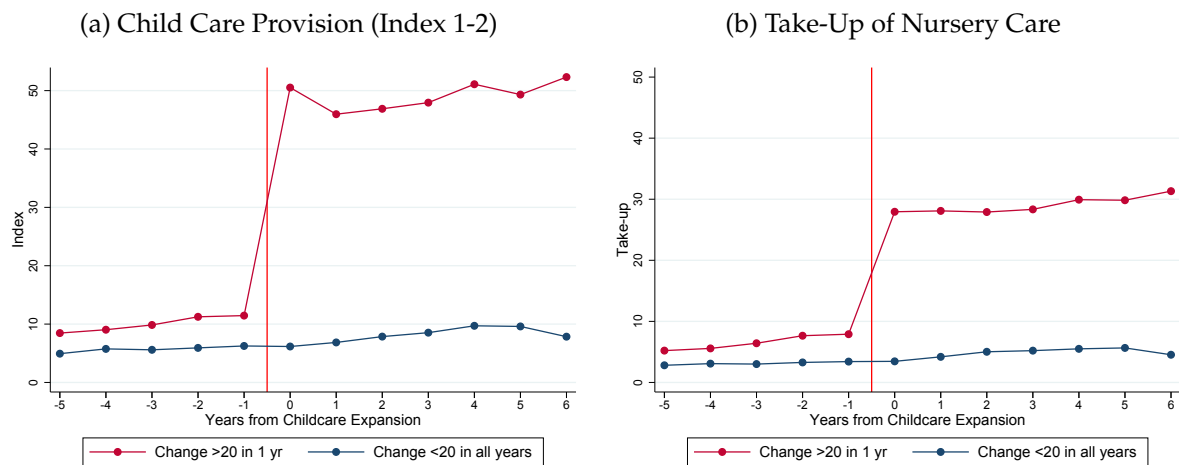
(b) Change in Index 1-5 between 2000 and 2010




Notes: The figure showcases the significant amount of spatial variation in the expansion of local child care provision in Austria. Each panel corresponds to a heatmap that reports for each municipality the change in the value of the child care index 1-5 from 1990 to 2000 (panel (a)) and from 2000 to 2010 (panel (b)). A darker color indicates a larger increase in local child care provision. See section 4.1 for further details on the construction of the child care provision index. GIS files of municipality borders provided by Statistik Austria.

Figure 8: Local Child Care Shocks: Evolution of Child Care Provision and Take-Up for Treated and Control Municipalities

Nursery Care Expansions



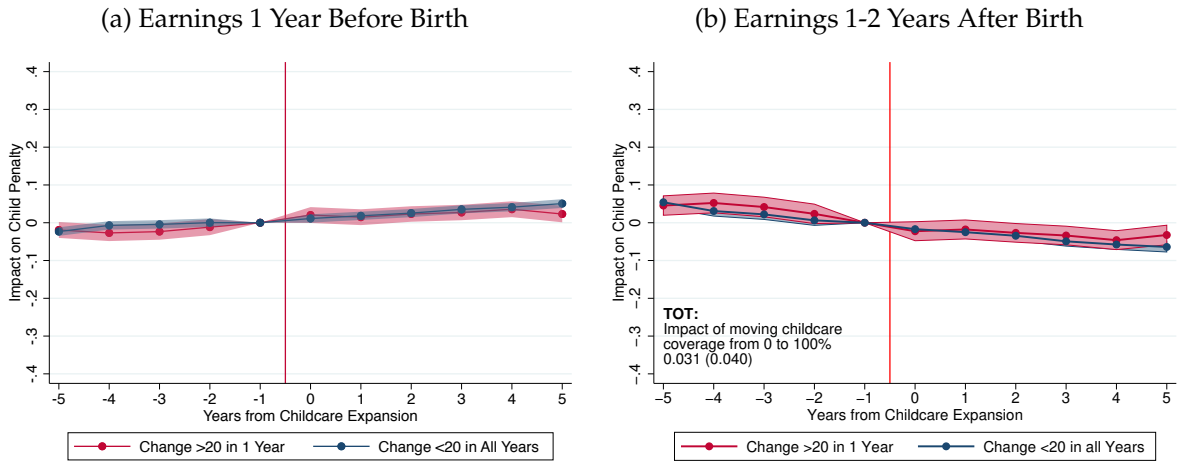
Pre-School Care Expansions




Notes: The figure shows the evolution of child care provision (left panels) and of the take up of institutional care (right panels) for treated and control municipalities around the time of a local shock to child care provision. A nursery care (resp. pre-school care) expansion event is defined as an increase of more than 20 points in the municipal nursery care (resp. pre-school care) index in one single year. Municipalities with an expansion event constitute the treated group, while municipalities that never experience an event constitute the control group. We randomly assign an expansion event year to control municipalities such that the distribution over calendar years is balanced between the two groups.

Figure 9: Estimated Effects of Nursery Care Expansions

Difference-in-Difference Evidence



Long Run Effects




Notes: The figure shows the difference-in-difference estimates of the effects of nursery care expansions on the full dynamics of female earnings. In panel (a) (resp. panel (b)), we plot the evolution of earnings at event time $t=-1$ (resp. $t \in [1,2]$), that is for women who are one year before (resp. between one and two years after) giving birth to their first child, in treated vs control municipalities, around the event of a nursery care expansion. In panel (b) we also report the TOT effect, which corresponds to the diff-in-diff estimate $\hat{\alpha}^{T}_{st}$ for earnings scaled by the first-stage, i.e. the equivalent diff-in-diff estimate for the change in the nursery care index. In panel (c), we plot the diff-in-diff estimate $\hat{\alpha}^{T}_{st}$ for earnings at all event times t w.r.t. birth of the first child, evaluated at $s = 10$, that is 10 years after the local nursery care expansion. Panel (c) is therefore the mirror image, for nursery care reforms, of panels (a), (c) and (e) of Figure 3 in the case of parental leave reforms. Based on the estimates of panel (c), we draw in panel (d) counterfactual child penalties under the pre and post nursery care expansion regimes. Standard errors are clustered at the municipal level, and shaded areas correspond to 95% confidence intervals.

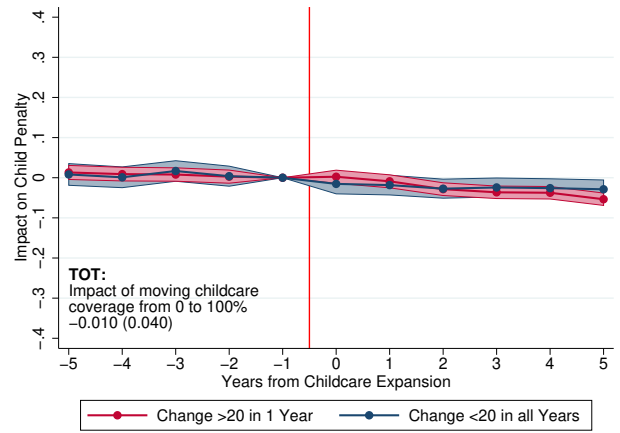
Figure 10: Estimated Effects of Pre-School Care Expansions

Difference-in-Difference Evidence

(a) Earnings 1 Year Before Birth




(b) Earnings 3-5 Years After Birth




Long Run Effects

(c) Effect on Child Penalty 10 Years After Reform

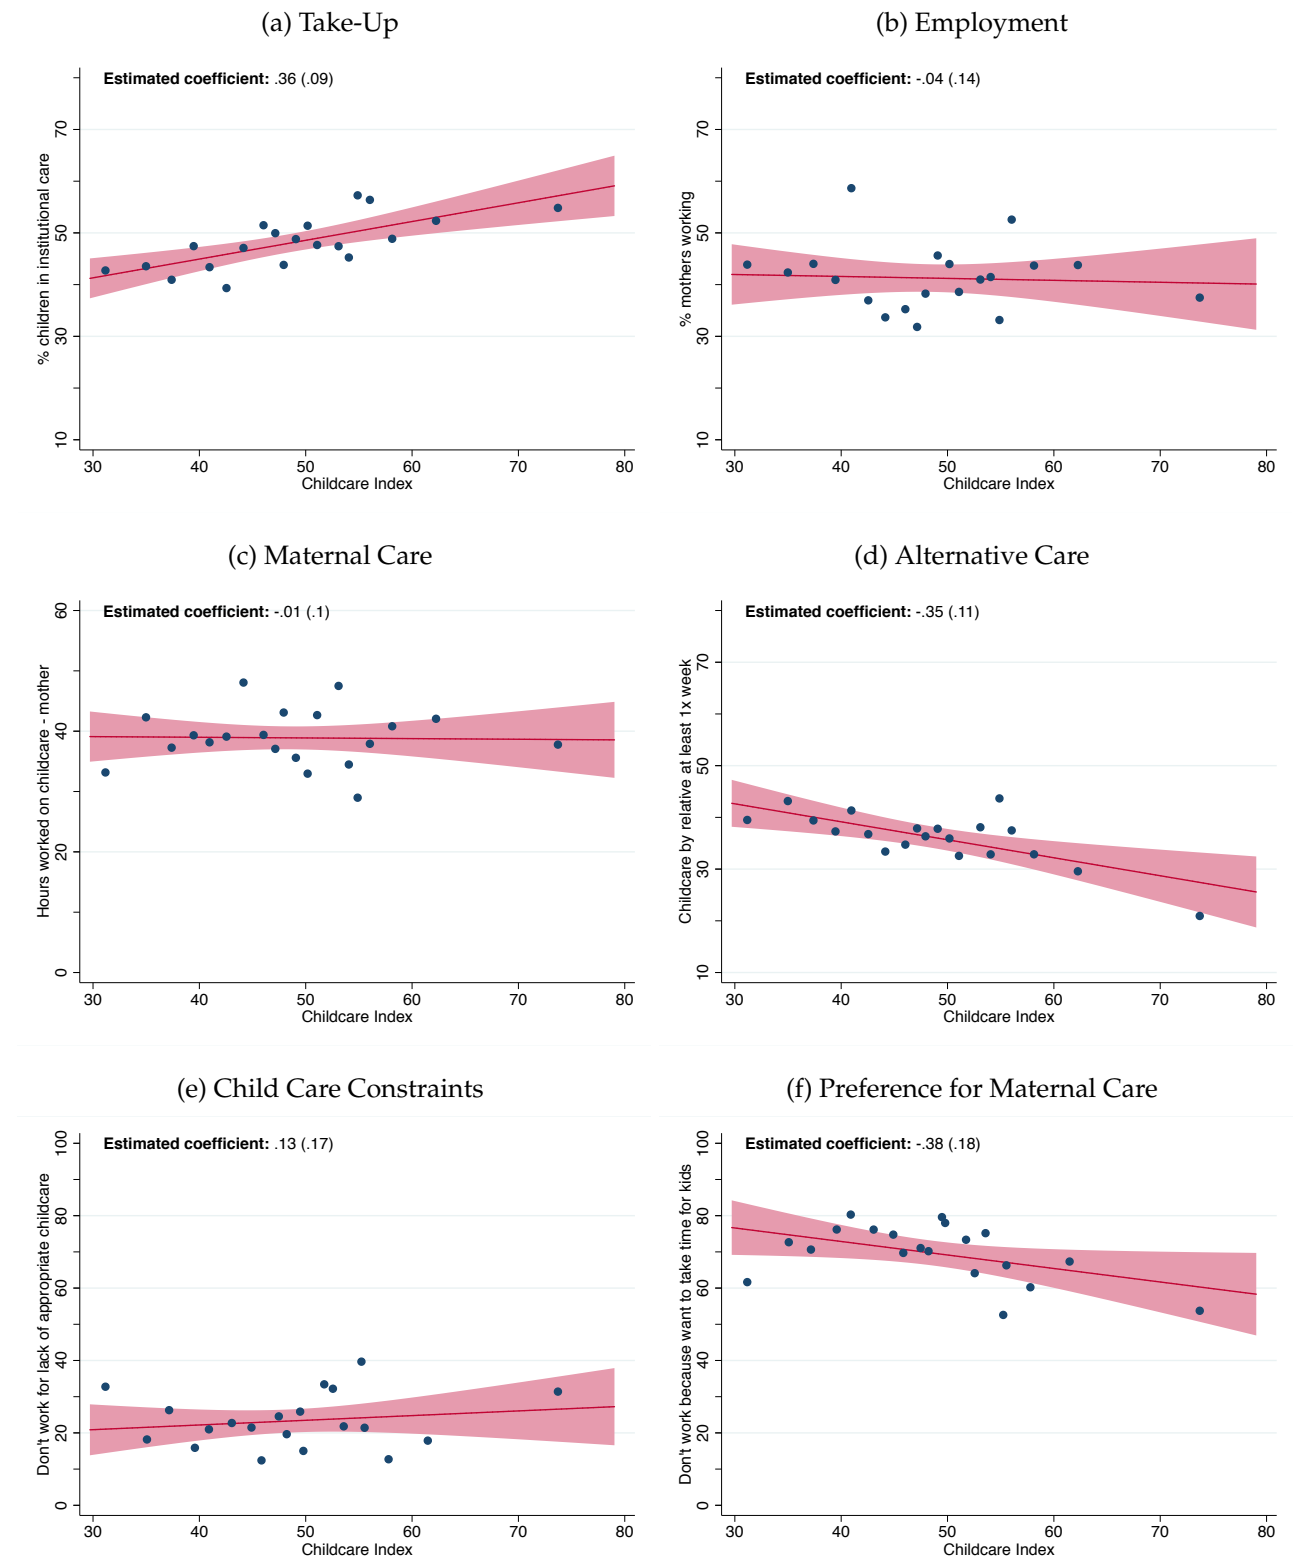


(d) Child Penalty 10 Years After Reform



Notes: The figure replicates the analysis of Figure 9 for pre-school care expansions events. It shows the difference-in-difference estimates of the effects of pre-school care expansions on the full dynamics of female earnings. In panel (a) (resp. panel (b)), we plot the evolution of earnings at event time $t=-1$ (resp. $t \in [1, 2]$), that is for women who are one year before (resp. between one and two years after) giving birth to their first child, in treated vs control municipalities, around the event of a pre-school care expansion. In panel (b) we also report the TOT effect, which corresponds to the diff-in-diff estimate $\hat{\alpha}_{st}^T$ for earnings scaled by the first-stage, i.e. the equivalent diff-in-diff estimate for the change in the pre-school care index. In panel (c), we plot the diff-in-diff estimate $\hat{\alpha}_{st}^T$ for earnings at all event times t w.r.t. birth of the first child, evaluated at $s = 10$, that is 10 years after the local pre-school care expansion. Panel (c) is therefore the mirror image, for pre-school care reforms, of panels (a), (c) and (e) of Figure 3 in the case of parental leave reforms. Based on the estimates of panel (c), we draw in panel (d) counterfactual child penalties under the pre and post pre-school care expansion regimes. Standard errors are clustered at the municipal level, and shaded areas correspond to 95% confidence intervals.


Figure 11: Reported Child Care Arrangements and Preferences: Residual Correlations With Child Care Provision




Notes: The figure reports correlations between measures of institutional child care provision at the political district level and household level information on time use, preferences and child care arrangements from two large external surveys from the Austrian Microcensus in 1995 and 2002. We pool observations from both waves of the survey. Each panel corresponds to a different outcome. We control for differences in the education and age structure across political districts by regressing the average outcome in a political district on the fraction of the population in three education bins and five age bins. Each panel plots a binscatter of the residuals, rescaled at the average level of education and age in Austria, against the value of the child care index (index 1-5). We also report on each panel the residual correlation estimate and its standard error. Panel (a) focuses on the fraction of children aged 2 to 5 in institutional child care. Panel (b) looks at the fraction of women who report working in the survey. Panel (c) focuses on the average number of hours that mothers report spending weekly on child care, while panel (d) looks at the fraction of children that are under the care of relatives. Panel (e) uses as an outcome the fraction of mothers reporting that their career is constrained due to a lack of appropriate child care, while the outcome in panel (f) is the fraction of women who report that they do not work because they do want to take care of their children.

Figure 12: Gender Gap Decomposition & Counterfactual Inequality Absent Family Policies

(a) Decomposition of Gender Gap Between Child-Related and Residual Inequality



(b) Gender Gap Estimates Absent Any Family Policy



Notes: This figure decomposes the gender gap in Austria for the period 1953 to 2017, and provides counterfactual gender inequality estimates absent family policies. Panel (a) plots the results of the decomposition of the earnings gap (corresponding to the sum of the blue and grey shaded areas) into child-related (blue) and residual factors (grey), following the methodology of [Kleven, Landais and Sogaard \(2019\)](#). In Panel (b), we superimpose the result of a counterfactual decomposition where we undo the effects of all family policy reforms since 1953, that is, we set parental leave duration and our child care indices to zero in all years. This decomposition is based on the reduced form estimates of the dynamic effects of all family policy reforms, from sections 3 and 4. See text for details.

Table 1: Gender Gap Decomposition Under Observed & Counterfactual Family Policy Regimes


Year	1953	1964	1970	1980	1990	2000	2010
A. Gender Inequality Decomposition							
Share Childless	0.217	0.180	0.163	0.151	0.151	0.157	0.167
Share Mothers $t < 0$	0.144	0.151	0.153	0.186	0.186	0.177	0.208
Share Mothers $t = 0-10$	0.172	0.173	0.184	0.197	0.217	0.208	0.188
Share Mothers $t = 11-20$	0.166	0.160	0.160	0.163	0.169	0.187	0.187
Share Mothers $t = 21-30$	0.166	0.168	0.166	0.153	0.151	0.156	0.160
Share Mothers $t > 30$	0.134	0.168	0.173	0.151	0.125	0.113	0.083
Earnings Women	307	790	1,308	3,877	6,540	9,705	13,095
Earnings Men	1,011	2,277	3,634	9,128	13,526	18,833	23,258
Employment Rate Women	0.495	0.515	0.486	0.522	0.538	0.599	0.646
Employment Rate Men	0.904	0.887	0.858	0.826	0.775	0.759	0.744
<i>Decomposition Btw Child-Related & Residual Inequality</i>							
Earnings Gap	0.696	0.653	0.640	0.575	0.517	0.485	0.437
Child-Related Gender Gap	0.370	0.356	0.338	0.322	0.367	0.366	0.347
B. Counterfactual Family Policy Scenario							
<i>No Parental Leave and No Child Care (Baseline)</i>							
Earnings Gap	0.696	0.646	0.628	0.562	0.499	0.463	0.414
Child-Related Gender Gap	0.370	0.342	0.314	0.296	0.331	0.323	0.301
<i>No Parental Leave and No Child Care (Optimistic)</i>							
Earnings Gap	0.696	0.648	0.632	0.565	0.503	0.474	0.428
Child-Related Gender Gap	0.370	0.347	0.322	0.302	0.340	0.345	0.329
<i>No Parental Leave and No Child Care (Pessimistic)</i>							
Earnings Gap	0.696	0.644	0.624	0.559	0.495	0.452	0.400
Child-Related Gender Gap	0.370	0.337	0.306	0.289	0.323	0.301	0.272

Notes: Panel A reports inputs and results of our decomposition exercise of gender gaps for selected years. In particular, we start by reporting the fraction of women ψ_{st} who are at each event time t relative to the birth of their first child in year s . We also report average nominal earnings of women and men in each year in Euro (we use the ECB reference rate to convert Austrian Shillings to Euro for years before the introduction of the Euro) estimated from the ASSD earnings distribution (employed only). The employment rates of women and men are obtained from the census. Finally, we report the actual and child-related earnings gaps corresponding to Figure 12 panel (a). In Panel B we report the results of our counterfactual decomposition exercise, where we undo the effects of all family policies since 1953. We first report the results from our baseline scenario, and then the results from two bounding exercises. The first, “optimistic” exercise, uses, instead of our point estimates, the upper bound of the 95% confidence interval on $\frac{dP_t}{d\tau_k}(\tau_k)$ for each policy change. The second, or “pessimistic” scenario, uses the lower bound of the 95% confidence interval. See text for details.


A Supplementary Figures: Context

Figure A.I: Agreement with Statements

(a) "A Woman Should Stay Home when She Has a Child under School Age"




(b) "When a Mother Works for Pay, Her Children Suffer"




Notes: The figure is based on data from the 2012 wave of the International Social Survey Program (ISSP).

Figure A.II: Child Penalties - Other Countries

(a) German-Speaking Countries



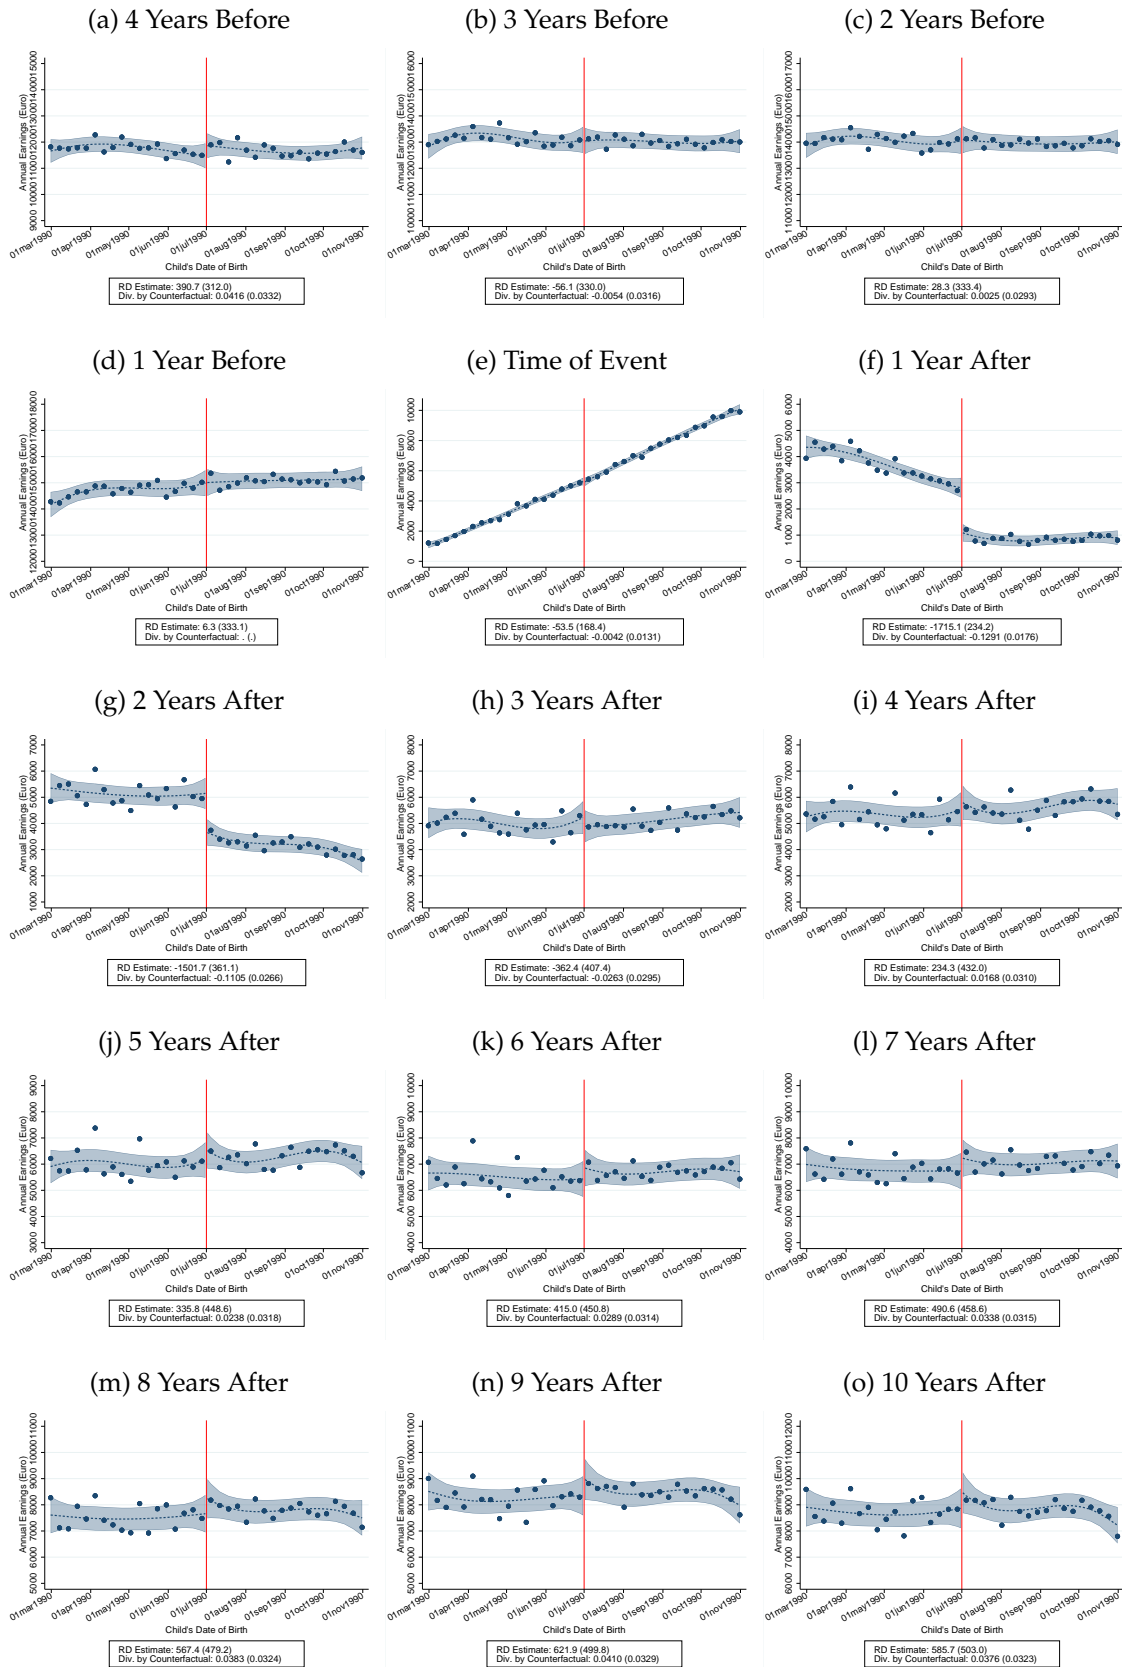
(b) Scandinavian Countries



Notes: Figures reproduced from [Kleven, Landais, Posch, Steinhauer and Zweimuller \(2019\)](#). We plot child penalties P_t^g as defined in equation (2) for each gender g by event time t (years from birth of first child) for a set of countries. Annual earnings are zero if an individual is not working in a given year. Long-run penalties are defined as $P_{10}^m - P_{10}^w$. For details on data sources and methodology see [Kleven, Landais, Posch, Steinhauer and Zweimuller \(2019\)](#).

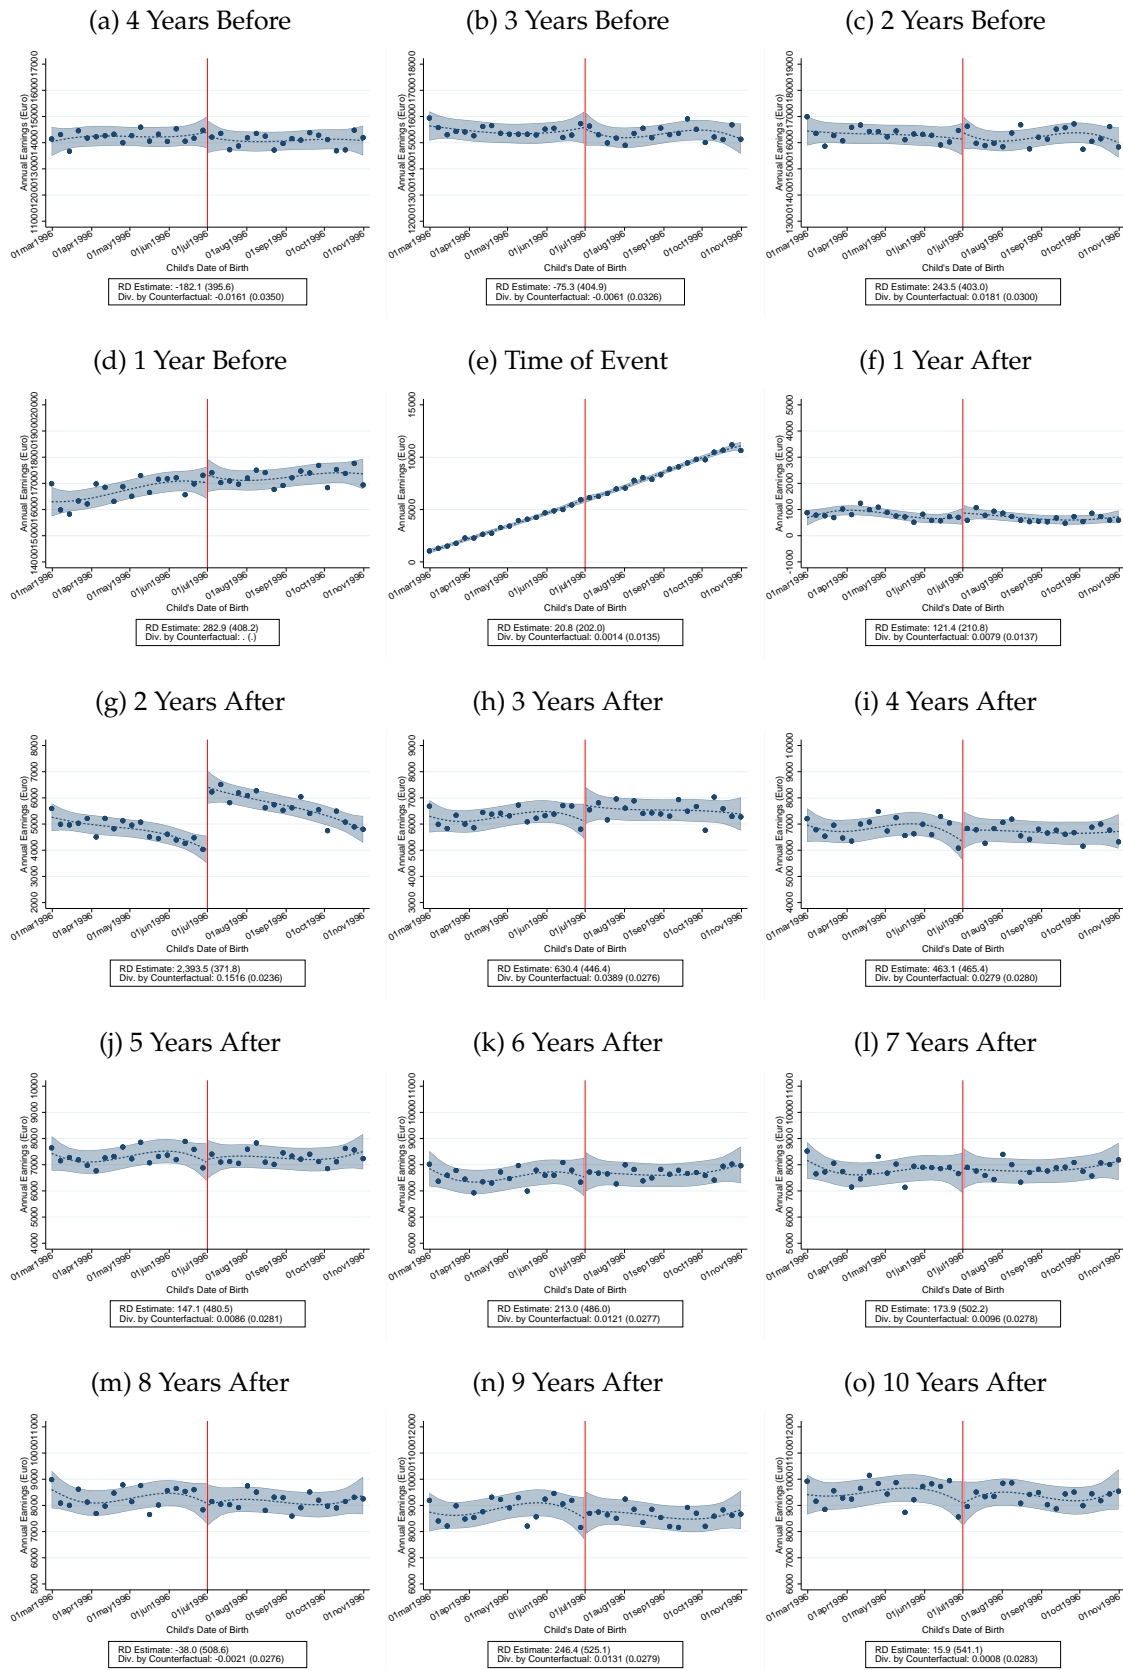
B Supplementary Figures: Parental Leave

Figure B.I: RDD Evidence on Maternal Earnings: 1990 Parental Leave Reform



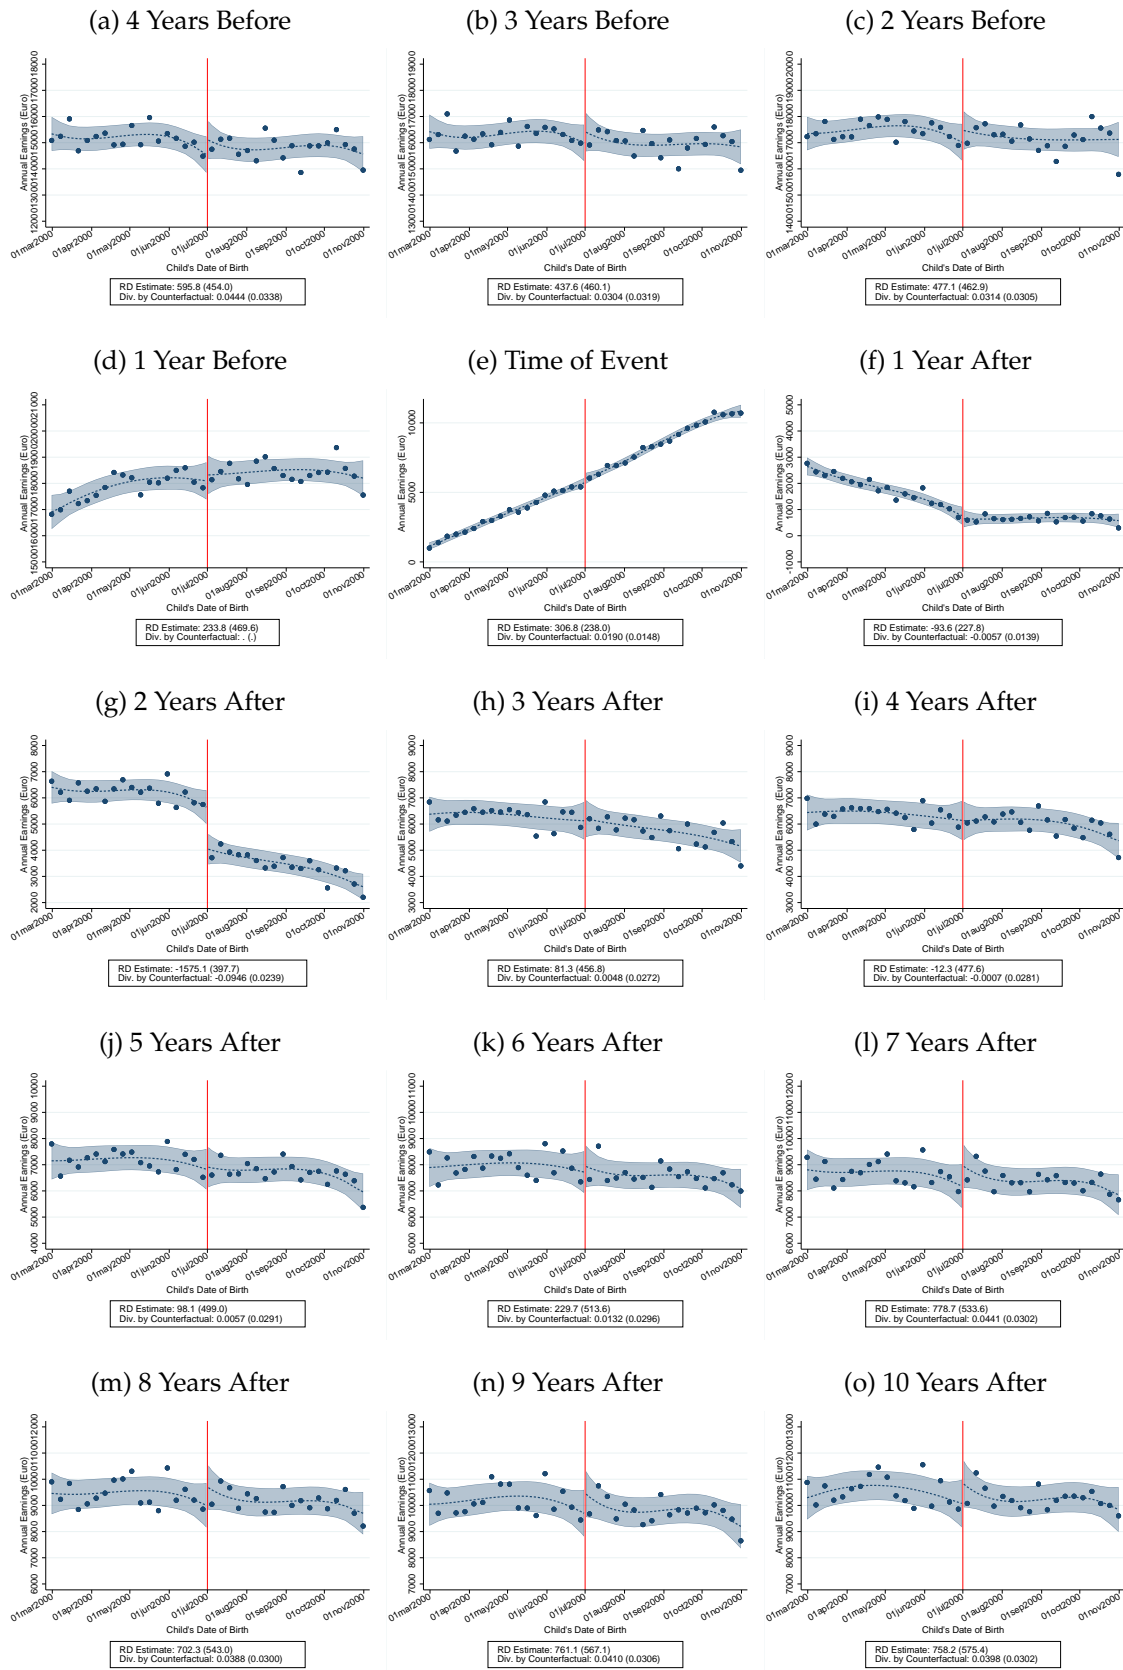
Notes: Figures plot the average annual earnings of mothers by week of first birth for each event time relative to first birth. We adjust earnings by CPI (base year 2000) and correct for topcoding (see Appendix D.1). Earnings are zero if the individual did not work in a particular year. We include first births in a 4-month window around the reform cut-off dates (July 1st 1990, 1996, and 2000). We exclude women ineligible for parental leave (less than 52 weeks of work in the two years preceding birth). For each event time separately we regress annual earnings on age dummies and a cubic in distance to the cut-off, allowing for separate trends on each side. We plot cubic trends from this regression including 95% confidence intervals. We report the discontinuity estimate and standard error as well as the estimate scaled by average counterfactual earnings at the same event time from child penalty regressions estimated on pre-reform samples.

Figure B.II: RDD Evidence on Maternal Earnings: 1996 Parental Leave Reform



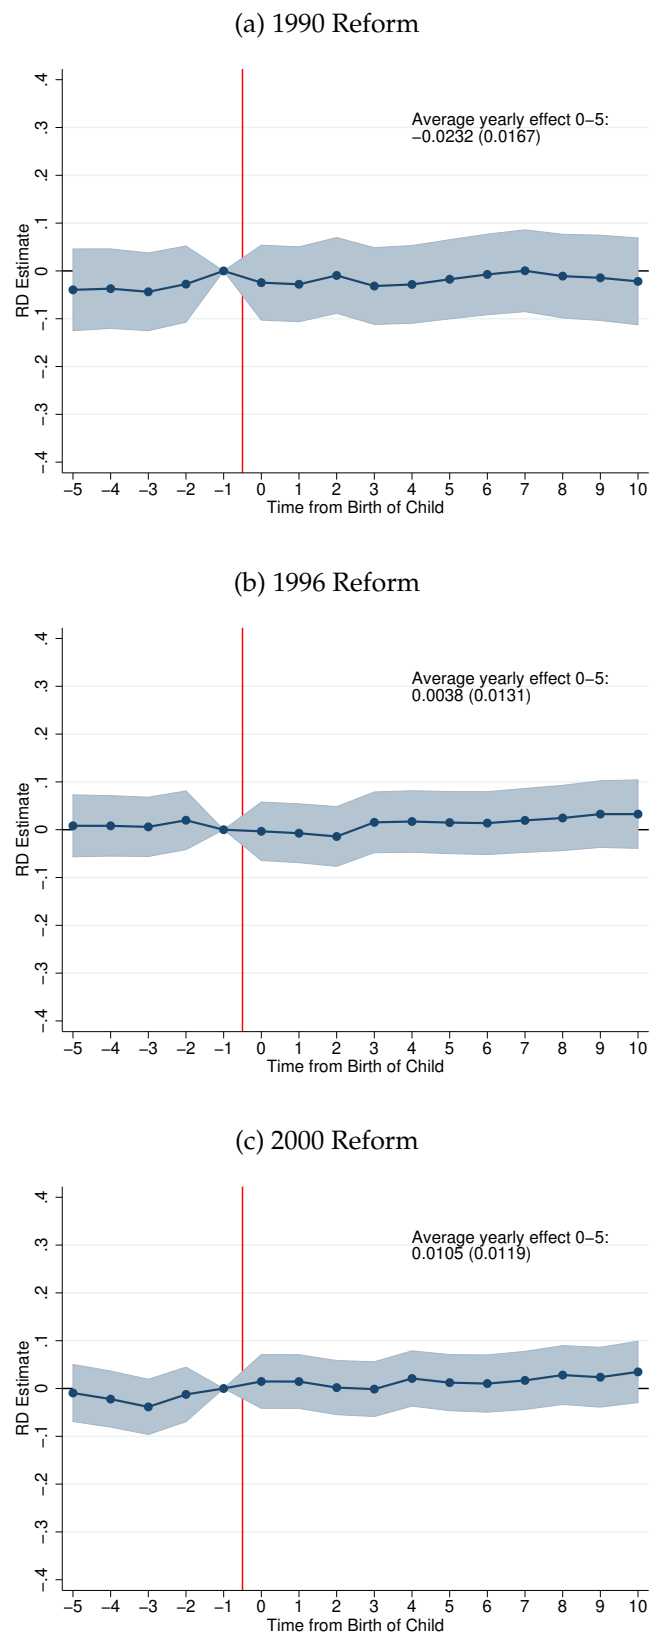
Notes: See notes to Figure B.I.

Figure B.III: RDD Evidence on Maternal Earnings: 2000 Parental Leave Reform




Notes: See notes to Figure B.I.

Figure B.IV: Effect on Paternal Earnings: 1990, 1996 and 2000 Parental Leave Reforms




Notes: The figure replicates Figure 3, but on the sample of first time fathers. It displays estimates of the effects of three parental leave reforms on the dynamic earnings of fathers. Fathers are matched to first births based on administrative data on applications for child benefits. See notes to Figure 3 for further details.

Figure B.V: 1961 Reform: Take-Up




Notes: The figure shows the evolution of maternity leave take up around the 1961 reform which introduced parental leave. We plot average days on parental leave in the year following birth of the first child, by month of first birth (based on REV data). As the reform was grandfathered, women who gave birth in 1960 became eligible, hence the observed gradual increase in take up.

Figure B.VI: 1961 Reform: Robustness to Trends



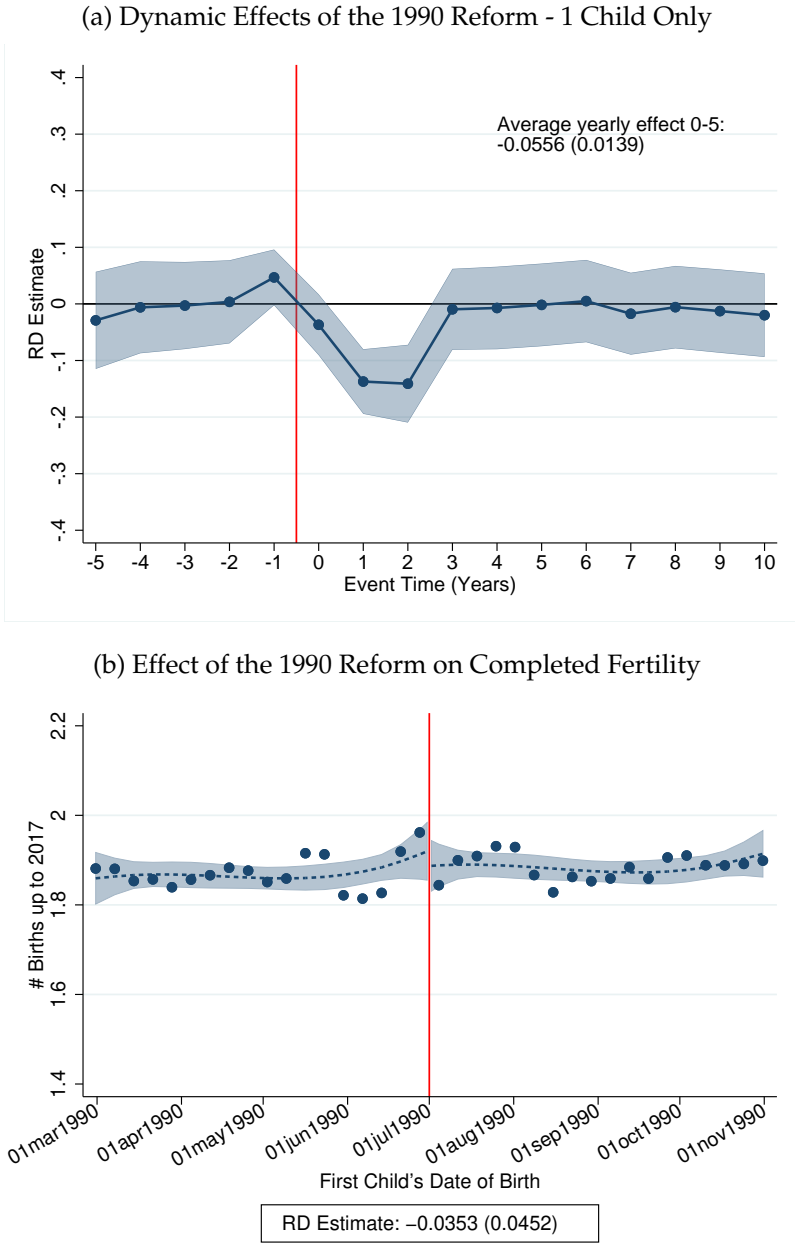
Notes: The figure investigates the presence of trends that may confound the estimated effects of the 1961 reform. The graph shows placebo results from specification 4 estimated on a sample of mothers who gave birth in 1962 ($T = 0$) or 1963 ($T = 1$). See notes to Figure 4.

Figure B.VII: Earnings Effect of Parental Leave Reforms on Mothers in Top Quartile of Pre-Birth Earnings Distribution




Notes: This figure investigates the presence of heterogeneous responses to parental leave reforms by mothers' earnings level. We reproduce estimates from Figure 3 on women at the top end of the earnings distribution. To this effect, we split our reform samples into four quartiles by annual earnings one year before birth. We then run exactly the same procedure as in Figure 3 keeping only mothers in the top pre-birth earnings quartile.

Figure B.VIII: Effects of Parental Leave: Mediating Role of Fertility



Notes: Panel (a) plots the 1990 reform effect (see notes to Figure 3) for the subsample of mothers who only have one child in the 1989-2017 period. Panel (b) plots average number of births by week of birth of the first child for mothers who gave birth to their first child between March-November 1990. Average births includes all observed births up to 2017. Dashed lines are fitted values based on a cubic polynomial separately fitted at each side of the cut off with 95% confidence interval.

Figure B.IX: Effects of Parental Leave on Child Penalties Accounting For Subsequent Births




Notes: Our baseline estimates in Figure 3 identify how child penalties respond to the PL regime applicable to the first child. To compute the total effect on child penalties of moving, in steady state, from one PL regime to another, this figure accounts for the effect of PL extensions for births of higher parity. To this effect, we turn to mothers giving birth to their second, third and fourth child, and replicate, for each parity, our dynamic estimates of the impact of PL reforms on earnings. We then construct the total steady state impact of a reform by adding dynamic estimates at all parities, weighted by the average completed fertility, and average timing between parities, of women who had their first birth around the time of each reform.


C Supplementary Figures: Child Care Provision

Figure C.I: Index of Child Care Provision over Time

(a) Child Care Index 1-2




(b) Child Care Index 3-5



Notes: This figure plots our child care indices over time. Each index is calculated at the municipality level based on child care availability and the number of children in the municipality. We then calculate weighted averages by year weighting by the total number of children age 1-2 (panel a) and children age 3-5 (panel b) in the municipality.

Figure C.II: Spatial Variation in Child Care Provision - Index 1-5 - 1990



Notes: This figure plots our child care index 1-5 by municipality using Austrian GIS maps with municipality borders (administrative boundaries). We shade each municipality according to the level of the index. Darker colors indicate a higher index in 1990.

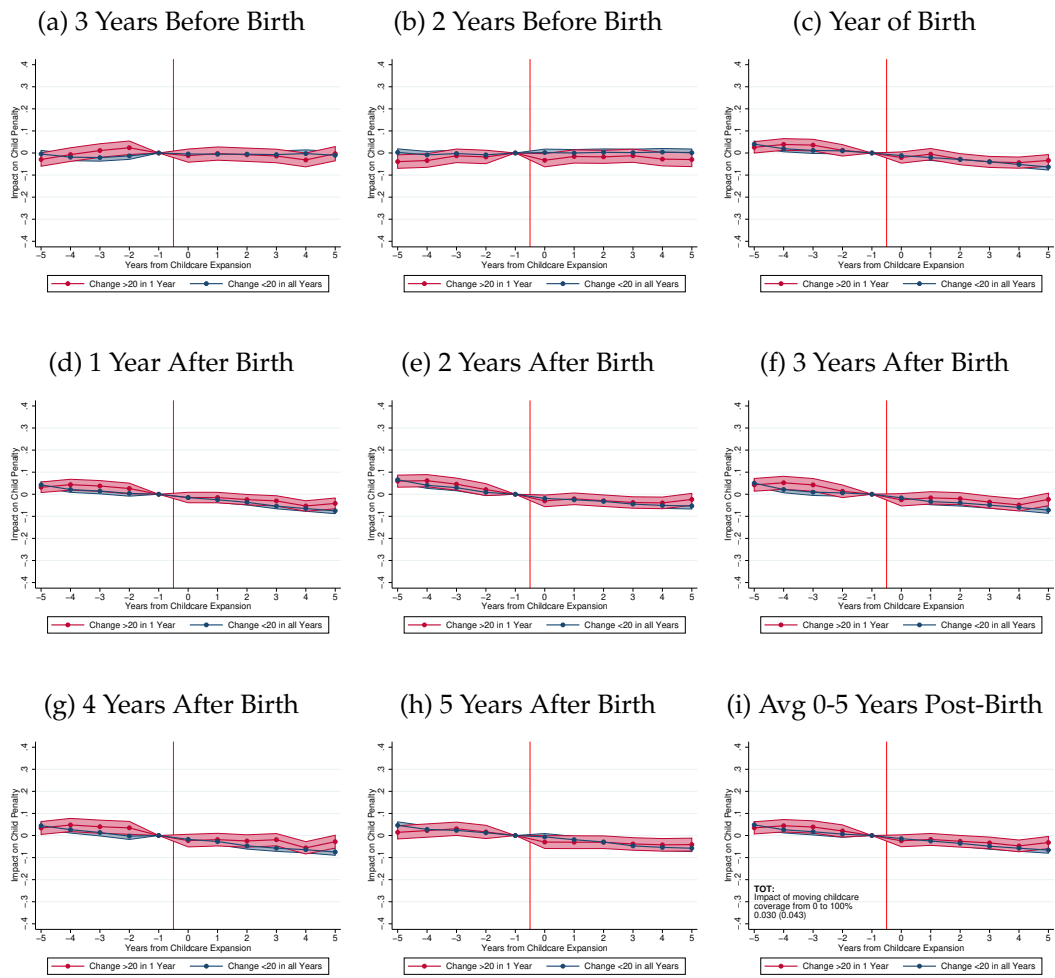
Table C.I: Descriptive Statistics of Municipalities By Treatment Status

Treatment Status	Index 1-2		Index 3-5		All
	Control	Treatment	Control	Treatment	
Panel A. 1991 Census Characteristics					
Number of Municipalities	1,554	284	253	1,507	1,887
Total Population	4,797,414	1,323,339	2,938,329	3,817,840	7,795,786
% Urban	0.39	2.11	9.49	0.07	1.48
% Towns/Suburbs	15.51	32.04	28.46	15.59	18.76
% Rural	84.11	65.85	62.06	84.34	79.76
Average Population	3,087	4,660	11,614	2,533	4,131
Empl. Rate Men	58.63	58.18	57.91	58.68	58.53
Empl. Rate Women	35.18	35.10	37.54	34.88	35.25
Female Earnings at $t=-1$	17,317	17,300	17,896	17,228	17,308
Panel B. Child Care in 1991					
% Municip. with Nursery	0.77	4.93	12.65	0.46	2.33
% Municip. with All-Day Pre-School	40.93	40.85	83.79	33.84	41.92
% of Children (1-2) in Nursery	0.00	0.01	0.02	0.00	0.00
% of Children (3-5) in Pre-School	0.57	0.59	0.82	0.53	0.57
Index 1-2	0.16	0.99	2.95	0.08	0.48
Index 3-5	50.54	49.97	85.86	45.25	51.01
Index 1-5	30.39	30.38	52.69	27.18	30.80

Notes: Panel A shows municipality characteristics in our treatment and control municipalities drawn from the 1991 census and earnings information from the ASSD. For earnings, we focus on women just about to give birth (i.e. in event time -1 with respect to the birth of their first child).

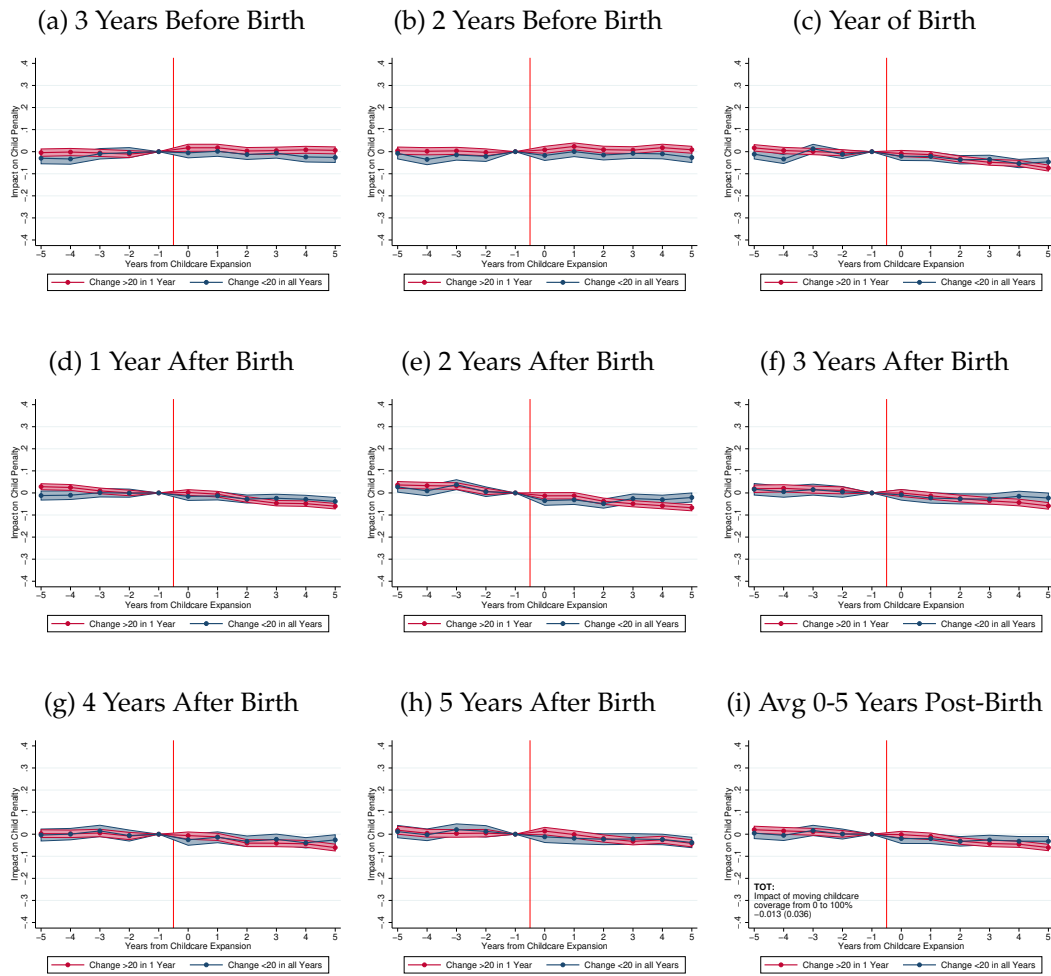
Panel B shows average child care availability according to the child care availability dataset provided by Statistics Austria. We calculate the average share with a nursery (children age 1-2), with all-day pre-school (children age 3-5) and the share of children enrolled in each of these types of child care facilities. We also show the averages of our child care indices age 1-2, 3-5 and 1-5.

Figure C.III: Nursery Care




Notes: The figure plots difference-in-differences estimates from specification 6 for the effects of nursery care expansions on the full dynamics of female earnings. Each panel corresponds to different event time with respect to the birth of the first child. In panel (a) we plot the evolution of earnings at event time $t=-3$, that is for women who are three years before giving birth to their first child, in treated vs control municipalities, around the event of a nursery care expansion. In panel (b), we do the same for women at event time $t=-2$. Etc.

Figure C.IV: Pre-School Care



Notes: The figure plots difference-in-differences estimates from specification 6 for the effects of pre-school care expansions on the full dynamics of female earnings. Each panel corresponds to different event time with respect to the birth of the first child. In panel (a) we plot the evolution of earnings at event time $t=-3$, that is for women who are three years before giving birth to their first child, in treated vs control municipalities, around the event of a pre-school care expansion. In panel (b), we do the same for women at event time $t=-2$. Etc.


Figure C.V: Effects of 1990 Parental Leave Reform by Level of Child Care Provision




Notes: The figure plots our reform impact estimates by whether the municipality of residence at the time of birth was below or above the median child care index at the time of birth. The reform impact estimates use IPW to weight mothers according to their pre-birth characteristics. Specifically, we run a probit regression of the treatment variable (above median index) on logs of annual earnings in years 3-5 before birth (interacted with year of first birth), employment status in the year before birth and age dummies. We then construct adjusted inverse probability weights based on predicted treatment probabilities and perform the same regressions as in Figure 3.

Figure C.VI: Fertility Effects of Local Child Care Expansions

(a) Nursery Care Expansions



(b) Pre-School Care Expansions



Notes: The figure plots event study estimates of a municipality-level regression of the fertility rate on event time relative to child care expansion interacted with treatment dummy and calendar year dummies. The fertility rate is measured as the number of births in a municipality from the birth register divided by the number of women age 15-49 residing in that municipality in the same year. We obtained the number of women age 15-49 by municipality and year from Statistik Austria register data. Shaded areas are 95% confidence intervals.

D Decomposing Gender Gaps: Data & Methodology

D.1 Measuring Gender Gaps in Earnings Over the Long Run

D.1.1 Using ASSD Data: 1972-2017

Our main data come from the Austrian Social Security Database. This is a matched employer-employee data set which contains detailed information on employment spells and earnings by calendar year over the period 1972 to 2017. The ASSD register covers the universe of paid employment spells, with the notable exclusion of civil servants, self-employed and farmers. This means that about 85% of the Austrian population is covered by ASSD, and this coverage rate has been stable over time. Our analysis therefore focuses on the long run evolution of the earnings gap on the ASSD population.

38

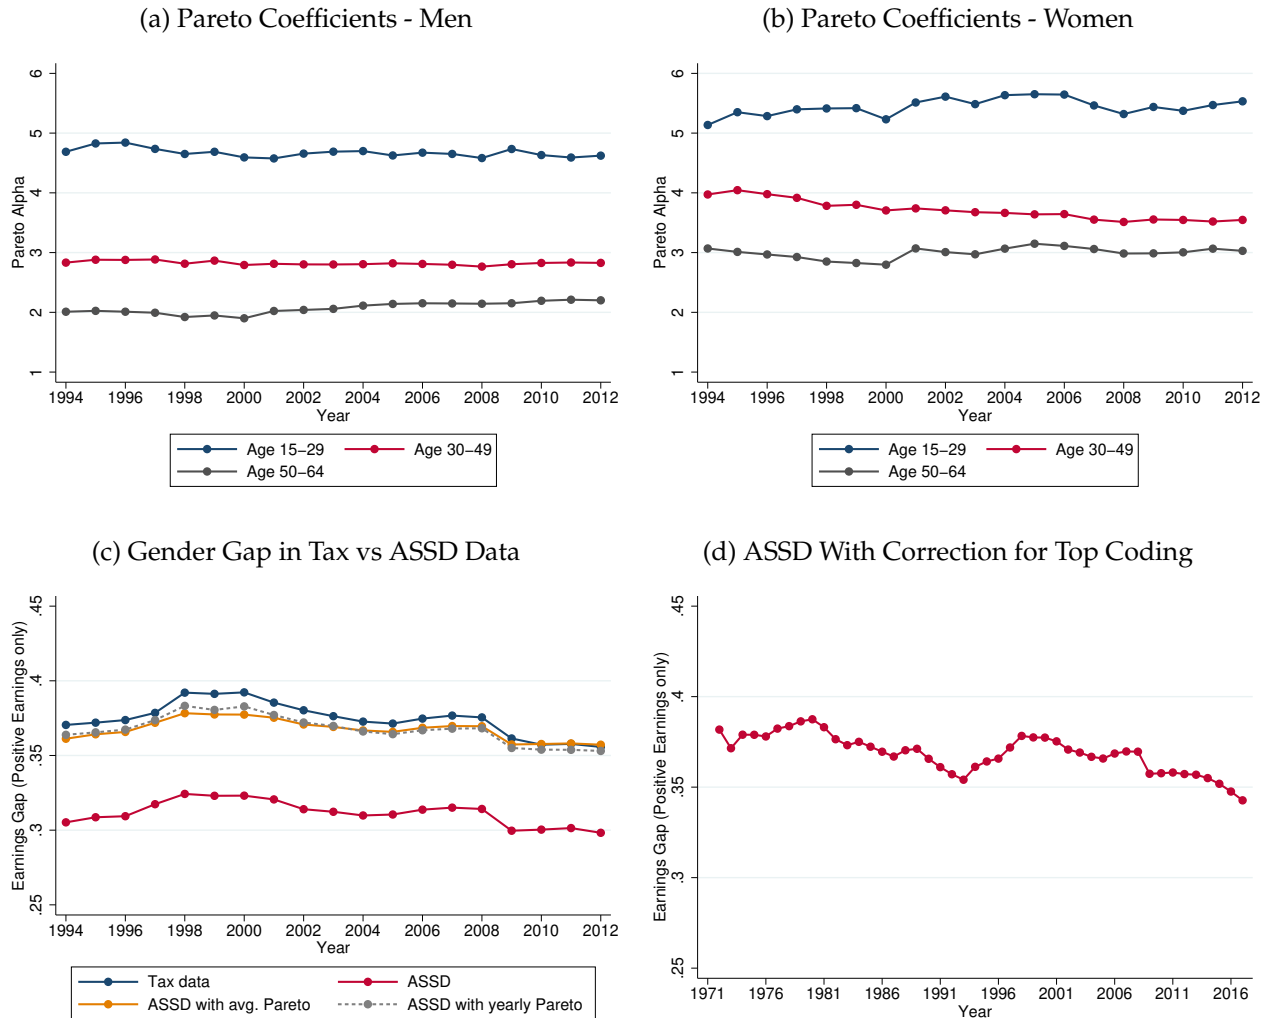
Earnings in the ASSD (and REV) are topcoded at the maximum contribution base, which changes from year to year. The fraction of topcoded individuals is quite stable over the period 1972-2017: around 8% of men and 2% of women have topcoded earnings. During the years 1994-2012 we can check how much topcoding affects earnings gaps by comparing the earnings gaps in the ASSD and in the income tax data (for workers in both datasets). In Figure D.I, panel (c) we plot the earnings gaps calculated from both sources for workers age 15-64, restricting both samples to individual with positive earnings only. Topcoding has mainly a level effect. The ASSD gap is about 12 percentage points below the income tax gap.

To adjust for topcoding in the ASSD data, we impute earnings above the top coding thresholds, using Pareto estimates of the earnings distribution computed using the income tax data. To be precise, we fit, in the tax data, separate Pareto distributions of earnings by gender and age groups. For each age group a and gender g , we compute Pareto coefficients α_a^g . In practice, we use three age groups: 15-29 years old, 30-49 years old, and 50-64 years old. Figure D.I panel (a) reports the Pareto coefficient of these distributions for each year for men, while panel (b) does the same for women. We see that these Pareto coefficients are very stable over time. In panel (c), we impute earnings above the threshold using these estimated Pareto parameters and compute gender gaps adjusted for top coding. The grey dots report estimated gender gaps when we use the yearly Pareto coefficients for each group, while the yellow series shows the gender gaps when we use the average Pareto coefficients over the period 1994-2012. We see that the ASSD series adjusted for top coding match very closely the gender gaps estimated in

³⁸We note that our estimated gender gaps likely understates the earnings gap in the entire Austrian population because men are more often working in the non-ASSD sectors (self employed, civil servants,...) than women and workers in those sectors earn slightly higher wages on average.

the non-top coded tax data, and that using yearly or average Pareto coefficients does not make a significant difference.

Figure D.I: Correcting for Top Coding of Earnings in the ASSD Data



Notes: Panels (a) and (b) plot our estimated Pareto coefficients for men and women by age. This is based on our sample of income tax records (1994-2012) which are not topcoded. We fit a pareto distribution by gender and age groups for gross labour income above the social security limit. The limit is adjusted each year and our procedure only fits the Pareto distribution for income above the relevant cut off in any given year. Panel (c) plots gender earnings gaps calculated from the income tax data, ASSD, ASSD adjusting topcoded earnings based on the average Pareto parameter and ASSD adjusting topcoded earnings based on the relevant Pareto parameter by year. Panel (d) plots the ASSD gender gap in annual earnings (excluding zeroes) based on the ASSD earnings adjusted for topcoding using the 1994-2012 average Pareto parameter.

Our sample includes any individual who has at least one record in the ASSD at any point in time. One issue is that we may miss some individuals who never enter the labor force. In practice, we can compare the fraction of individuals with zero earnings in our

sample and in the census data for years when the census is available. This comparison shows that our sample slightly underrepresents women with zero earnings compared to the census. We therefore added women with zero earnings to our sample to match the share of non-employed in the census.

D.1.2 Measuring Gender Gaps Before 1972

Before 1972, ASSD data exist but has not been digitized. However earnings statistics (*Lohnstufenstatistik*), including detailed tabulations by gender, drawn from the original data, were regularly published, enabling us to consistently compute the evolution of gender inequality on the population covered by the ASSD since 1953.³⁹ We used tabulations of the earnings distributions for women and men in 1953, 1957, 1964, 1967, 1970, 1973, 1976, and 1979 that report average monthly earnings for women and men (in July), deciles of the earnings distribution, and the share of women among all ASSD workers. Importantly, to calculate the mean and deciles of all distributions, the IHS imputed topcoded earnings based on the assumption of a lognormal distribution. To account for individuals not in the labor force, we also added men and women with zero earnings so that the employment rates of men and women before 1972 matches the census.

D.2 Decomposition of Gender Gap Between Child Related and Residual Inequality

D.2.1 Methodology

Our methodology for decomposing the gender gap follows [Kleven, Landais and Sørensen \(2019\)](#). It relies on estimating the effect of children (child penalties) on women and men separately. In the spirit of a Oaxaca-Blinder decomposition, we then ask how much of the average gender gap can differences in these estimates explain. We refer to [Kleven, Landais and Sørensen \(2019\)](#) for further details.

We define the average gender gap in year s as $\Delta_s \equiv 1 - E[Y_{ist}^w | s] / E[Y_{ist}^m | s]$. Based on our empirical evidence, and to simplify notation and exposition, we assume in what follows that there are no effect of children on the earnings dynamics of men. We are interested in the average counterfactual gender gap in year s absent child penalties for women, that is: $\tilde{\Delta}_s \equiv 1 - E[\tilde{Y}_{ist}^w | s] / E[Y_{ist}^m | s]$. This gap captures the level of gender inequality that would happen absent child penalties on women's careers. In other words, $\tilde{\Delta}_s$ represents the residual inequality across men and women in the labor market that is not related to

³⁹IHS Report ("*Geschlechtsspezifische Einkommensunterschiede: Österreich 1953-1979*") by Christl and Wagner (IHS Vienna 1981)

the arrival of children. The child related gender gap is therefore captured by $\Delta_s - \tilde{\Delta}_s \equiv E[\tilde{Y}_{ist}^w - Y_{ist}^w | s, t] / E[Y_{ist}^m | s]$

In practice, we reconstruct the counterfactual earnings gap in year s $\tilde{\Delta}_s$ based on four inputs:

1. the observed average level of earnings of women in year s , $E[Y_{ist}^w | s]$
2. the fraction ψ_{st} of women who are in event time t with respect to the birth of their first child birth,
3. the child penalty for each event time P_{st} ,
4. and the average level of counterfactual earnings at each event time $E[\tilde{Y}_{ist}^w | s, t]$.

To understand the relationship between the counterfactual earnings gap and these four inputs, note that the computation of $\tilde{\Delta}_s$ solely relies on having an estimate of $E[\tilde{Y}_{ist}^w | s]$. We now show that we can estimate $E[\tilde{Y}_{ist}^w | s]$ based on the four inputs mentioned above. For this, we start by noting that the average level of earnings of women in year s $E[Y_{ist}^w | s]$ is the weighted average of the earnings of women who will never have children (childless women) and of the earnings of women who have, or will have, children and are, in year s , at event time t with respect to the birth of their first child.

$$E[Y_{ist}^w | s] = \psi_s^{Childless} \cdot E[Y_{ist}^w | s, childless] + \sum_t \psi_{st} \cdot E[Y_{ist}^w | s, t]$$

where we have that $\psi_s^{Childless} + \sum_t \psi_{st} = 1$. Similarly, we can write the average counterfactual level of earnings of women in year s $E[\tilde{Y}_{ist}^w | s]$ as the weighted average of the earnings of women who will never have children (childless women) and of the counterfactual earnings \tilde{Y}_{ist}^w of women who have, or will have, children and are, in year s , at event time t with respect to the birth of their first child.

$$E[\tilde{Y}_{ist}^w | s] = \psi_s^{Childless} \cdot E[\tilde{Y}_{ist}^w | s, childless] + \sum_t \psi_{st} \cdot E[\tilde{Y}_{ist}^w | s, t]$$

From this, and the definition of child penalties in equation (2), it follows that:


$$\begin{aligned} E[\tilde{Y}_{ist}^w | s] &= E[Y_{ist}^w | s] + \sum_t \psi_{st} \cdot E[\tilde{Y}_{ist}^w - Y_{ist}^w | s, t] \\ &= E[Y_{ist}^w | s] - \sum_t \psi_{st} \cdot E[P_{st} \cdot \tilde{Y}_{ist}^w | s, t] \end{aligned}$$

We show in the following sections how we computed or estimated each of these four inputs.

D.2.2 Estimating Fraction ψ_{st} of Women at Each Event Time

Fraction of Childless Women We estimated the cohort-specific share of childless women from the Human Fertility Database and the Austrian Birth Barometer. Both data sets contain data and estimates by the Vienna Institute of Demography. For older cohorts, we used a simple linear exponential smoothing predictor (Holt's method) to predict childlessness back to cohorts born in 1887 (age 64 in 1951). Figure D.II below shows the evolution of the fraction of childless women, i.e. women whose completed fertility is zero, by birth cohort.


Figure D.II: Evolution of the Fraction of Childless Women by Birth Cohort



Notes: The figure plots the share of childless women (defined as women who will never have children over their lifetime) by birth cohort from the Austrian Birth Barometer (Vienna Institute of Demography) and a backward prediction based on the damped Holt's method ('holt' in R with damped=TRUE and bandwidth 14).

Fraction of Women at Each Event Time To estimate how many mothers are at event time t in a given year s , we need to know the distribution of age at first birth by cohort. Mean age at first birth is available yearly from 1952 to 2017 from the Austrian Birth Barometer (HFD). Figure D.III compares mean age at first birth from the Austrian Birth Barometer and from the Birth Register. We also show age at first birth estimates from IPUMS census microdata (1971-2011).

Figure D.III: Mean Age at First Birth




Notes: The figure plots mean age at first birth from various data sources. Census: From the 1971-2011 decennial census (IPUMS and Statistics Austria). We keep all women whose oldest child is between 0-10 at the census date and calculate the average age of the mother in the year the oldest child was born. Birth register (exact): from 1984 the birth register contains data on parity. We keep first births in each year and calculate average age of the mother. Birth register (first in union): before 1984 parity information was only available by mother-father pairs. We keep first births by union in each year and calculate average age of the mother. Birth register (first in union or out of wedlock): We keep first birth by union or births out of wedlock (no father) each year and calculate average age of the mother. HFD: Estimates from the Human Fertility Database.

To get the fraction of women at each event time, we need the distribution of age at first birth, and not just the mean. To this purpose, we fit a Beta distribution to the Birth Register data (1984-2007) for each year. Then we separately regress the two coefficients α_s and β_s of the distribution on mean age at first birth \bar{A}_s . This gives us: $\hat{\alpha}_s = 6.404 - 0.122 \cdot \bar{A}_s$ and $\hat{\beta}_s = 30.346 - 0.966 \cdot \bar{A}_s$.

To predict out of sample we simply plug in mean age at first birth from the HFD and predict α_s and β_s for all years in HFD (1952-2017). We then use the implied Beta distribution in each year to assign mothers to the different event times.

The distribution of age at first birth by year and the childlessness estimates finally enable us to estimate, for each year s the composition of women ($\psi_s^{Childless}$ and $\psi_{st}, \forall t$). The following figure shows the evolution of the composition of women, splitting event times into 5 bins: $t < 0, t = 0 - 10, t = 11 - 20, t = 21 - 30$, and $t > 30$.

Figure D.IV: Composition of Women



Notes: The figure plots our constructed shares of childless women and mothers by event time relative to first birth in each calendar year. See text for details on construction of these shares.

D.2.3 Estimating Child Penalties P_{st} & Counterfactual Earnings Y_{ist}^w By Birth Cohort

We finally need to estimate child penalties and counterfactual earnings by birth cohort over the long run.

To do this, we extend the baseline specification (1) to allow for year-specific coefficients on event time. Specifically we consider the following specification:

$$Y_{ist}^g = \sum_y \sum_{j \neq -1} \alpha_{yj}^g \cdot \mathbf{I}[j = t] \cdot \mathbf{I}[y = s] + \sum_k \beta_k^g X_{kis}^g + v_{ist}^g \quad (8)$$

where we interact the event time dummies with year dummies in order to estimate year-specific event coefficients α_{yj}^g . Note that estimating event coefficients by calendar year s and event year t amounts to estimating event coefficients by birth cohort $c = s - t$. As in our baseline specification (1), we include a full set of age dummies and year dummies in the set of covariates.


Child Penalties: ASSD Data (1972-2017) and REV data (1960-1972) To compute child penalties by birth cohort since 1960, according to specification (8) we use data from the

ASSD and from the REV before 1972. As explained in section 2, the REV data has a very high coverage rate for women born between the mid-1920s and the 1950s. Panel (a) of Figure D.V below shows that for these birth cohorts, the fraction of women with an employment record represents close to 80% of women with employment records in the census. This number is very similar to the coverage rate of ASSD. But due to the way the REV was created, not all women had their earnings recorded even if the employment and maternity insurance spells are recorded in the REV data. The selection of spells that are recorded with earnings is not random: it is related to the fact that the REV data was set up for pension calculation purposes. When the earnings attached to an employment spell were not relevant for pension computation, i.e. spells with low or no earnings levels (not relevant for top 18 years of earnings variable) the corresponding earnings were less likely to be recorded. This means that spells at younger age with lower earnings level are less likely to have recorded level of earnings attached. As a consequence, the number of first births identified with earnings data in the REV data is lower than the true number of first births, as shown in panel (b) of Figure D.V.


To account for this differential selection of observable first births in the REV data, we use an Inverse Probability Weighting correction. We know that first births are more likely to be recorded with earnings if earnings are relatively high. We therefore reweight the earnings of women who have a first birth in the REV data so that the average earnings of women giving birth relative to average female earnings in Austria in that year matches that observed in the ASSD data. Figure D.VI below shows the evolution of earnings for women observed prior to having children in the REV data up to 1977, and in the ASSD data afterwards. Panel (a) shows the raw data: we see a clear downbreak in the series in 1977, showing that the women who have a first birth spell recorded with earnings information are on average positively selected in the REV data. Panel (b) shows the results when applying our IPW correction.

Figure D.V: Coverage of REV Data

(a) Fraction of Women With Employment Records in REV

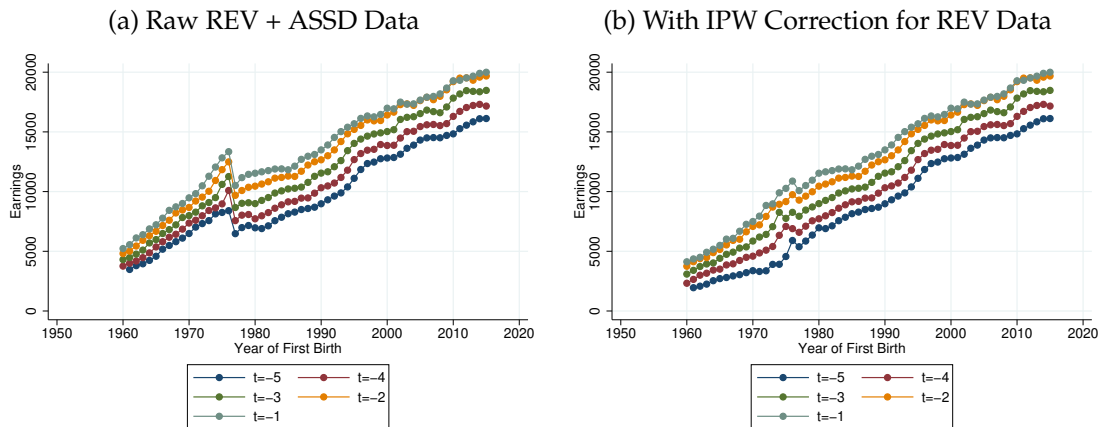


(b) Number of First Births by Year: Comparison Across Data Sets



Notes: Panel (a) plots the share of women with employment records in the REV by year of birth. We select all employment spells from the REV around the 1971 census date (December 5th 1971) and calculate the number of REV employed women by year of birth. We then divide this number by the number of women who are employed according to the 1971 census (again by year of birth). Panel (b) plots the raw number of first birth from various data sources. Census: We keep women with oldest child age 0-10 from each census 1971-2011 and calculate the total number of first births by year. REV+ASSD: We merge the REV and ASSD data and keep the first birth ever observed for each woman. In the REV data this is based on observing a maternity insurance spell, in the ASSD data a live birth is indicated separately. Birth register (various): See notes to Figure D.III for definition of exact and first in union.

Figure D.VI: Earnings by Year of Birth and Event Time in REV and ASSD Data, With and Without IPW Correction



Notes: Panel (a) plots average annual earnings from the REV and ASSD by year of first birth at various event times. We select all first births over the REV+ASSD period. We exclude mothers who never have positive earnings in the REV if their birth was before 1977 to ensure we have 5 years of pre-birth data for everyone in the sample. Panel (b) plots average earnings weighted by IPW. To construct the weights we run yearly probit regressions of 'after' (first birth in 1980 vs e.g. 1970) on the logs of earnings divided by average earnings in the same year from our decomposition procedure at event times -4 to -1 (we include zeroes by adding 1 to earnings) and whether earnings are zero at any of these event times. Based on the predicted probabilities we then construct inverse probability weights.

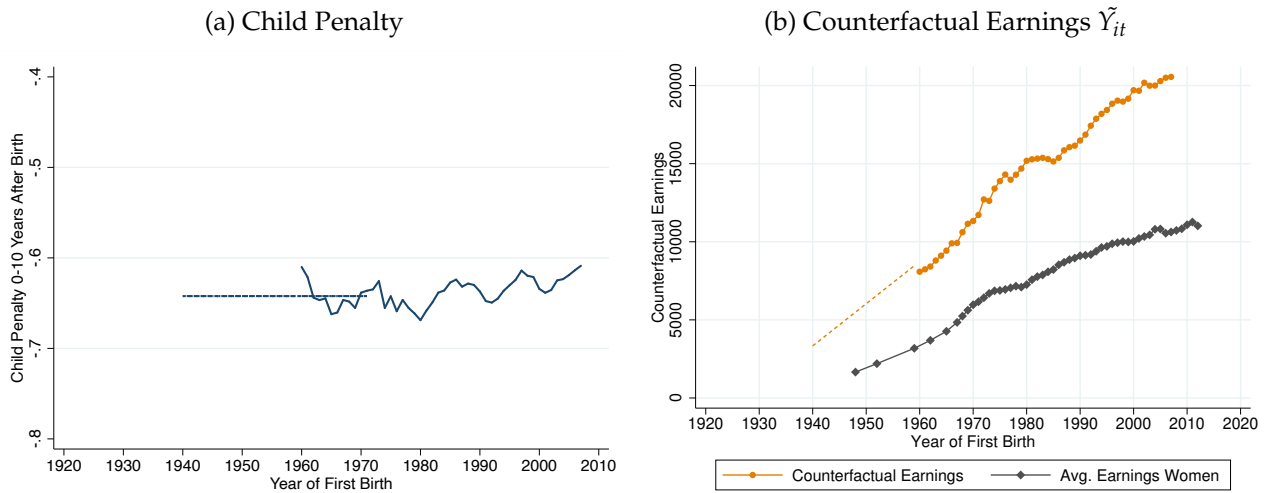
Estimated Child Penalties By Birth Cohorts Figure D.VII shows the estimated child penalties by birth cohort, following specification 8, on the ASSD sample and the REV sample with IPW correction. We also plot the estimated counterfactual earnings $E[\tilde{Y}_{ist}^w | s, t]$ corresponding to these specifications. We see that child penalties in the first ten years after birth have remained remarkably stable over the last 50 years. But child penalties at longer horizons (11 to 20 years after birth, and 21 to 30 years after birth), after being quite stable in the 1960s and early 1970s, have started to decrease significantly.

For first births before 1960, we need to extrapolate child penalties: we assume that child penalties for first births before 1960 were equal to the average penalty observed for first births between 1960 and 1972. For counterfactual earnings corresponding to births before 1960, we extrapolate them linearly based on the observed trend in counterfactual earnings for first births between 1960 and 1972.


We estimate specification 8 using event time dummies up to $t = 30$. As a result, we also need to extrapolate penalties after event time 30: given the extreme stability of penalties after event time 10, we fix them at the estimated child penalty level in event time 30: P_{s30} .

Figure D.VII: Reforms: Dynamic RD Estimates & Effects on Child Penalties


0 to 10 years After Birth



11 to 20 years After Birth



21 to 30 years After Birth



Notes: The figure plots average child penalties and predicted counterfactual earnings over event times 0-10, 11-20 and 21-30 by year of first birth (see text for details). Dashed lines indicate the extrapolations we use based on simple averages for child penalties (1960-72 average) and linear backward prediction for counterfactual earnings. We also add our constructed average earnings of all women (including zeroes) in the figures on the right.

D.3 Historical Gender Gaps With Counterfactual Family Policies

We now turn to explaining how we can use the previous decomposition, and our estimates of the effects of policies on child penalties to measure gender gaps under counterfactual policy scenarii.

Our focus is on the three following family policies: parental leave duration (τ_{PL}), the coverage of free nursery care (τ_N), and the coverage of free pre-school care (τ_{PS}). What we are trying to identify, for family policy $\tau_k \in \{\tau_{PL}, \tau_N, \tau_{PS}\}$, is the causal effect of any policy change $d\tau_k$ on the level of the gender gap: $\frac{d\Delta_s}{d\tau_k}$. Based on our previous decomposition exercise, we can express this effect as a simple function of the causal effects of the policy change $d\tau_k$ on the full dynamics of the child penalty ($dP_t/d\tau_k, \forall t$). Indeed, we have that:

$$\begin{aligned} \frac{d\Delta_s}{d\tau_k} &= -\frac{dE[Y_{ist}^w|s]/d\tau_k}{E[Y_{ist}^m|s]} \\ &= -\frac{1}{E[Y_{ist}^m|s]} \cdot \left\{ \sum_t \psi_{st} \cdot \frac{dP_t}{d\tau_k}(\tau_{PL}, \tau_N, \tau_{PS}) \cdot E[\tilde{Y}_{ist}^w|s, t] \right\} \end{aligned}$$

To get the above expression, we used the empirical fact that earnings of men are affected neither by the arrival of kids, nor by family policies. The above expression also assumes that family policies have no effect on fertility, such that $d\psi_{st}/d\tau_k = 0$. But we explore below scenarii where we relax this assumption and allow family policies to affect fertility as well.

Note that the effect of any policy change on the child penalty is theoretically a function of *all* family policies. But in practice, our empirical evidence found no interaction effects between parental leave and child care provision. This means that we can consider the impact of a policy change on the child penalty to be solely a function of that particular policy: $\frac{dP_t}{d\tau_k}(\tau_k)$.

As a consequence, for all family policy reforms since 1961, we can compute the causal effect of that reform on the gender gap in each year s by simply combining:

1. our estimates of ψ_{st} , the fraction of women at each event time t in year s
2. our estimates of average counterfactual earnings $E[\tilde{Y}_{ist}^w|s, t]$ for women absent child penalties, at each event time t in year s
3. and our estimates of $\frac{dP_t}{d\tau_k}(\tau_k)$, the effect of each the reform on child penalty at each event time t


Baseline Results: Undoing All Family Policies Our baseline counterfactual policy scenario consists in getting rid of all family policies. In other words, we simulate the evolution of the gender gap since 1961, when eliminating the effects of all parental leave reforms, and setting the level of nursery care and of pre-school care provision to zero. The table below summarizes our estimates of the effect of each policy on child penalties that we use to compute the counterfactual gender gaps. Note that for parental leave reforms, we account for the steady state effect of the reform on child penalties, that is we account for the effect on child penalties of getting access to a new regime PL regime for all children, and not just the first one. This issue is discussed at length in section 3.3 and results of these steady state effects on child penalties are presented in Figure B.IX.

	Avg. effect on 0-10 Penalty (first birth)	Avg. effect on 0-10 Penalty (higher order births)
1961 Reform	-0.0281	-0.0459
1990 Reform	-0.0227	-0.0234
1996 Reform	0.0109	0.0251
2000 Reform	-0.0149	-0.0145
Index 1-2 Expansion (for Index from 0 to 100)		0.022
Index 3-5 Expansion (for Index from 0 to 100)		0.027

Note also that given we systematically found no long run effects of any of these policies on child penalties, we systematically set the effect of each policy reform on penalties to be zero for event time superior to 10: $\frac{dP_t}{d\tau_k} = 0, \forall t > 10$.

Figure D.VIII below shows the effect on the gender gap of removing each policy separately. Panel (a) shows what the gender gap would have been in the absence of any parental leave reform. In this scenario, we only undo parental leave expansion, and assume that child care provision would have followed the same path as actually observed. The graph highlights that parental leave expansions have had a negative, albeit small, impact on gender inequality. This result stems from the significant negative impact of parental leave policies on labor supply in the short run. But the absence of long run effects of PL on the careers of women means that the overall negative impact of parental leave policies on gender inequality remains somewhat modest. While parental leave policies do have a small effect on the evolution of the gender gap, we find in panel (b) and (c), to the contrary, that child care policies have none. When undoing these policies, we find that the gender gap would have been absolutely unaffected by the absence of nursery care or of pre-school care.

Figure D.VIII: Gender Gap in Earnings and Counterfactual Gender Gaps Under Different Policy Scenarios



Notes: The figures plot actual earnings gaps and counterfactual earnings gaps based on setting various policy parameters to zero in isolation. See text and notes to Table 1 for details.

Sensitivity Analysis It is critical to assess how sensitive our counterfactual gender gaps estimates are to the precision with which we estimate the effects of policies on child penalties. To this effect, and to provide some guidance for inference, we computed two bounding counterfactual gender gaps. The first, “optimistic” one, uses, instead of our point estimates, the upper bound of the 95% confidence interval on $\frac{dP_t}{d\tau_k}(\tau_k)$ for each policy change. The second, or “pessimistic” scenario, uses the lower bound of the 95% confidence interval. The results of these two bounding scenariii are presented in Table 1. The bottom-line is that our null effects are relatively precise: we can clearly rule out large positive effects of family policies on gender inequality.

Accounting for Fertility Effects We assumed, for our baseline counterfactual policy scenario that family policies had no effect on fertility. We now explore scenariii allowing for potential effects on fertility. In practice, this means accounting for potential effects of each policy change on the fraction of women observed at each event time t : $\frac{d\psi_{st}}{d\tau_k}$. The causal effect of the policy change on the gender gap can now be written as:

$$\begin{aligned} \frac{d\Delta_s}{d\tau_k} &= -\frac{dE[Y_{ist}^w|s]/d\tau_k}{E[Y_{ist}^m|s]} \\ &= -\frac{1}{E[Y_{ist}^m|s]} \cdot \left\{ \sum_t \psi_{st} \cdot \frac{dP_t}{d\tau_k} \cdot E[\tilde{Y}_{ist}^w|s, t] + \sum_t \frac{d\psi_{st}}{d\tau_k} \cdot P_{st} \cdot E[\tilde{Y}_{ist}^w|s, t] \right\} \end{aligned}$$

where the second part in the bracketed term accounts for fertility responses. Intuitively, if a policy reform induces more couples to have children, it will increase the fraction of women observed at each event time t after birth ($d\psi_{st}/d\tau_k > 0$). As these women are facing penalty P_{st} , this will increase the gender gap.

We use two different sources to account for fertility effects. To account for fertility effects of parental leave we use the estimates from [Olivetti and Petrongolo \(2017\)](#) Table 3 column (8). Their formula for total fertility rate is (from the replication data set)

$$TFR = \dots - 0.0007448 \times \text{leave_weeks} + 0.0013043 \times \frac{\text{leave_weeks}^2}{100} + 0.0019767 \times \% \text{weeks_paid} + 0.0003374 \times \% \text{payrate}$$

The following table shows what these estimates imply when we undo parental leave. Note that in their data set `leave_weeks` is equal to weeks of maternity insurance plus weeks of parental leave. So undoing parental leave but keeping maternity insurance means setting leave duration to 8 weeks.

To account for fertility effects of **child care expansions** we use our own estimates of the effect of these roll-outs on total fertility rate. We constructed total fertility rate at the mu-

Table D.I: Summary of Fertility Impact Computed Based on Estimates from [Olivetti and Petrongolo \(2017\)](#)

Regime	leave (wks)	%wks paid	%payrate	Counterfactual			TFR Effect of PL
				leave (wks)	%wks paid	%payrate	
1961-1990	60	100	52	8	100	100	-0.0088
1990-1996	112	100	46	8	100	100	0.0671
1996-2000	112	75	37	8	100	100	0.0146
2000-2008	112	123	30	8	100	100	0.1072

Notes: The table shows the policy parameters we use for the different periods to calculate the total fertility effect of parental leave policies. We use the formulas in the online dofiles provided by [Olivetti and Petrongolo \(2017\)](#).

municipality level as follows. First, we get the number of births at the municipality level from the birth register (1988-2007) and Statistik Austria Register Data (2002-2018).⁴⁰ Then we get the number of women age 15-49 from the census (1991, 2001, 2011) and the Population Register (2011-2017). We interpolate linearly to get yearly data. Dividing the number of births by the number of women is the crude fertility rate. We then use the Austrian aggregate yearly relationship between crude fertility rate and total fertility rate to convert crude to total fertility rate.

We use the effects 10 years after expansion as an estimate for the TFR impact of child care expansion, scaled by the difference in the index during event times 0-5.


	Effect at $t = 10$	Avg. change in index	Scaled Effect
Nursery Care Expansion	-0.049645	0.363941	-0.1364
Pre-School Care Expansion	-0.046957	0.369499	-0.1271

Figure [D.IX](#) shows actual and counterfactual total fertility rates. Counterfactual means fertility when there is no parental leave and no child care.

⁴⁰And we take the average of those two numbers for years 2002-2007 where the data overlaps, to account for a few small discrepancies in the data.

Figure D.IX: Actual and Counterfactual Total Fertility Rates

(a) By Year



Notes: The figure plots the total fertility rate (TFR) calculated by multiplying the number of children by the shares of women who have 0-4 children as constructed for the main decomposition exercise. The counterfactual TFR is based on setting all family policy parameters to zero taking into account our estimated child care expansion fertility effects and the parental leave expansion effects based on [Olivetti and Petrongolo \(2017\)](#).

To convert our estimated fertility effects into estimates for $\frac{d\psi_{st}}{d\tau_k}$, we proceed as follows. Our fertility estimates are, theoretically, measuring the effect of policy reforms on total fertility $dN/d\tau_k$, where N is the average number of children by women. Note that $N = (1 - \psi^{childless}) \cdot \bar{n}$, where \bar{n} is the average number of children among women with children and $\psi^{childless}$ is the fraction of childless women.

What we do is to assume that policy reforms affect the total fertility by changing the fraction of women with kids rather than the number of kids per mother. In other words, we assume that $dN/d\tau_k = d(1 - \psi^{childless})/d\tau_k \cdot \bar{n}$


As a consequence, we easily obtain an estimate of the effect of the reform on the fraction of women at event time t in year s :

$$\begin{aligned} \frac{d\psi_{st}}{d\tau_k} &= -\frac{d(\psi^{childless})}{d\tau_k} \cdot \psi_{st} \\ &= \frac{dN}{d\tau_k} \cdot \frac{\psi_{st}}{\bar{n}} \end{aligned}$$


Figure D.X shows the results of our counterfactual policy scenario of removing all family policies, when accounting for fertility responses. For almost the entire period the policy reforms slightly decreased fertility overall, so undoing the reforms slightly increases the gender earnings gap. In Table D.III we also explore the sensitivity of our results to the precision with which the fertility effects are estimated. We computed two bounding counterfactual gender gaps. The first, “optimistic” one, uses, instead of our point estimates, the upper bound of the 95% confidence interval on $\frac{dN}{d\tau_k}$ and $\frac{dP_t}{d\tau_k}(\tau_k)$ for each policy change. The second, or “pessimistic” scenario, uses the lower bound of the 95% confidence interval.

Figure D.X: Counterfactual Gender Gaps Under No Family Policy Accounting for Fertility Responses

(a) Without Fertility Responses



(b) With Fertility Responses



Notes: The figures plot the results of our decomposition exercise removing all family policies when we do not take fertility effects into account (panel (a)) and when we incorporate fertility effects (panel (b)). See text for details.

Table D.III: Sensitivity Of Counterfactual Gender Gap Estimates

Year	1953	1964	1970	1980	1990	2000	2010
Panel A. Decomposition							
Actual Earnings Gap	0.696	0.653	0.640	0.575	0.517	0.485	0.437
Counterfactual Earnings Gap	0.326	0.297	0.302	0.253	0.149	0.119	0.090
Gap Due to Kids	0.370	0.356	0.338	0.322	0.367	0.366	0.347
Panel B. Undoing the Reforms Without Fertility Response							
<i>No Parental Leave and No Child Care (Baseline)</i>							
Earnings Gap	0.696	0.646	0.628	0.562	0.499	0.463	0.414
Child-Related Gender Gap	0.370	0.349	0.326	0.309	0.349	0.344	0.324
<i>No Parental Leave and No Child Care (Optimistic)</i>							
Actual Earnings Gap	0.696	0.648	0.632	0.565	0.503	0.474	0.428
Child-Related Gender Gaps	0.370	0.351	0.330	0.312	0.353	0.355	0.338
<i>No Parental Leave and No Child Care (Pessimistic)</i>							
Earnings Gap	0.696	0.644	0.624	0.559	0.495	0.452	0.400
Child-Related Gender Gap	0.370	0.346	0.322	0.306	0.345	0.333	0.309
Panel C. Undoing the Reforms With Fertility Responses							
<i>No Parental Leave and No Child Care (Baseline)</i>							
Earnings Gap	0.696	0.646	0.629	0.565	0.505	0.472	0.423
Child-Related Gender Gap	0.370	0.349	0.327	0.312	0.356	0.353	0.332
<i>No Parental Leave and No Child Care (Optimistic)</i>							
Earnings Gap	0.696	0.649	0.633	0.568	0.510	0.483	0.437
Child-Related Gender Gap	0.370	0.352	0.331	0.315	0.360	0.364	0.347
<i>No Parental Leave and No Child Care (Pessimistic)</i>							
Earnings Gap	0.696	0.644	0.625	0.561	0.501	0.460	0.408
Child-Related Gender Gap	0.370	0.347	0.323	0.308	0.352	0.341	0.318

Notes: The table reports inputs and results of our decomposition exercise of gender gaps for selected years, when accounting for fertility responses to family policies. See notes to Table 1 for additional details. The “optimistic” scenario is based on the upper bound of the 95% confidence interval for all policy treatment effect estimates and the “pessimistic” scenario is based on the lower bound of the 95% confidence interval.