Behind the Child Penalty: Understanding What Contributes to the Labour Market Costs of Motherhood*

Alessandra Casarico
Bocconi University, CESifo, Dondena

Salvatore Lattanzio
University of Cambridge

June 19, 2021

Abstract

We study the short- and long-run impact of motherhood on labour market outcomes and explore the individual and firm-level factors that influence it. Using matched employer-employee data for Italy over 1985-2018, through an event study methodology around childbirth, we show that the long-run child penalty in annual earnings is 57 log points and it largely depends on the change in labour supply along the intensive margin. The birth of a child increases the probability of transition to non-employment, reduces the likelihood of having executive roles and increases that of working in firms with lower productivity, sales, capital and wages, providing evidence of sorting into worse firms after childbirth. In the heterogeneity analysis, we find that the child penalty is higher for young, low-wage mothers and those taking longer leaves. It is larger in firms with less generous pay, worse peers, in more gender-conservative regions and where childcare services are scarcer.

Keywords: Child penalty, motherhood, labour supply, heterogeneous effects, matched employer-employee data

***JEL* codes**: J13, J16, J31

*We would like to thank Augusto Cerqua, Joanna Kopinska and seminar participants at University of Rome La Sapienza for useful comments and suggestions. Salvatore Lattanzio gratefully acknowledges ESRC DTP Studentship (nr. ES/J500033/1) and Cambridge Commonwealth European and International Trust for financial support for PhD studies. Contacts: Casarico, alessandra.casarico@unibocconi.it; Lattanzio, sl828@cam.ac.uk.

1 Introduction

In spite of the large inroads of women in education and employment, gender gaps in participation and wages are persistent labour market phenomena, which characterise all countries in the world, though to a different extent. For instance, the Gender Equality Index computed by EIGE (2020) for the European Union stands at 67.9 – with 100 indicating perfect equality – and at 72.2 for the work domain. Maternity has long been recognised as a key source of gender inequality in the labour market (Becker, 1985, 1991). More recently, Kleven et al. (2019b) have estimated child penalties in female earnings for Denmark, showing that they amount to around 20 per cent. The penalties that mothers experience compared to fathers are even larger for countries like Sweden, Spain, the US and Germany, where the penalty stands at 32, 28, 30 and 60 per cent, respectively (Angelov et al., 2016; Kleven et al., 2019a; de Quinto et al., 2020). Also in view of this evidence, Bertrand (2020) identifies maternity as one of the persisting causes of gender gaps in the labour market, and Cortés and Pan (2020) discuss factors, both at home and at work, that may contribute to amplify the career-family trade off that women face.

In this paper we provide evidence on the short- and long-run effects of having a child on female labour market outcomes and explore the individual and firm-level factors that influence the size of child penalties. We use matched employer-employee data for Italy and in an event study set-up like the one developed in Kleven et al. (2019b) and Angelov et al. (2016), we compare labour market outcomes of mothers and non-mothers to study the different trajectories of the two groups over a fifteen year period post childbirth. In particular, we focus on sorting across firms with different characteristics after childbirth as an adjustment margin which is novel to the literature on child penalties. Furthermore, exploiting the richness of our data, and with the goal of understanding what may drive the costs of motherhood at the individual and firm-level, we investigate how the child penalty varies with the duration of the parental leave, workers' ability, occupation and age at birth and, on the firm side, how it changes with firm wage premia, the quality of peers, firm size and a measure of family friendliness of the firm. Last, we relate child penalty estimates to attitudes towards gender roles and the availability of childcare services to explore whether differences in gender norms or diffusion of childcare centers correlate with the magnitude of the child penalty across regions.

We find that the impact of motherhood on female labour market outcomes is large. The difference in log annual earnings between women without and with children is 57 log points fifteen years after childbirth relative to the year before maternity. Most of the impact on earnings is due to a reduction of the labour supply along the intensive margin: the difference in the number of full-time equivalent weeks worked in a year amounts to 50 log points. The remaining 7 log points are due to weekly wage differences between mothers and non-mothers. Furthermore, women with children tend to have slower career progression and to enter non-employment more frequently. Exploiting data on firms' balance sheet, we show that women with children tend to sort into firms with lower value added, sales, capital and average wages, compared to women without children. This is evidence of an important behavioural response after childbirth.

Heterogeneity analysis reveals that the child penalty is stronger for mothers taking longer parental leaves, for those who are low-wage and blue-collar, and for those who are younger than 30 years old when they have their first child. As to firms' characteristics, working in low-wage firms, with low-quality peers and in firms with high share of female workers increases the penalty in annual earnings. There are no differences, instead, between large and small firms. Finally, child penalties are lower in regions characterised by more progressive views on gender roles and more availability of childcare services, suggesting that not only individual and firm-level factors matter, but also cultural and institutional features.

The paper contributes to the literature by providing evidence on child penalties in Italy, a country displaying one of the highest gender gap in employment in Europe and with very slow progress in closing gender gaps in the labour market. In particular, we complement the short-run evidence presented in Martino (2020), first, by comparing the earnings trajectories of mothers to those of non-mothers (therefore, going beyond a before-after comparison by adding an explicit control group) and, second, by showing the long-run penalty exploiting data up to fifteen years after childbirth. Our paper is also related to the literature that investigates the factors influencing the magnitude and persistence of the child penalty for different groups of workers. Specifically, we add to the evidence provided in Bruns (2019) on how the firm contributes to labour market outcomes of women after childbirth and how mothers tend to shy away from high-pay establishments. We further show that women with children tend to enter non-employment more frequently and – within firm – they have lower career prospects than women without children. At the same time, mothers employed in high-pay firms before maternity experience a lower child penalty compared to those employed in low-pay firms. We also show that sorting into firms with worse financial outcomes may explain why women experience a child penalty following childbirth.

The remainder of the paper is organised as follows. Section 2 describes the data and provides descriptive statistics. Section 3 describes the empirical strategy. Section 4 reports the event study results. Finally, section 5 concludes.

2 Data

We carry out our analyses using three archives on workers, firms and social security contributions from LoSaI (Longitudinal Sample Inps) records, a matched employer-employee dataset that contains a random sample of the universe of workers in the Italian non-agricultural private sector. The data covers approximately 7 percent of the universe of employees over the period 1985-2018. The worker archive contains the entire work and pay history of each sampled individual; in particular, it records information on annual gross earnings,[1] the number of weeks and days worked in a year, the type of contract (full-time or part-time and permanent or temporary) and broad occupation categories (apprentice, blue-collar, white-collar, middle manager, executive). In a separate archive, the dataset contains demographic information on each employee, such as year of birth, gender and region of residence. The firm archive records total firm size in discrete brackets and firm's 2-digit industry, based on NACE Rev. 2 classification. The contribution archive reports for each worker the full history of social security contributions from the first employment spell to the last of his or her career. This archive not only records actual contributions paid by employers, but also imputed contributions related to leaves of absence, sick leaves, unemployment benefit receipt and, crucially, maternity leave. This latter information allows us to identify childbirth episodes based on the first month of maternity leave, which has a mandatory duration of five months and can be taken one to two months before the expected childbirth and lasts until three to four months after.

We perform a number of sample restrictions. First, for workers holding multiple contracts in a year we only retain the information on the main job, i.e. the one with the highest number of weeks worked or with the highest earnings. We deflate earnings using the OECD non-food non-energy consumer price index. Moreover, as we use worker and firm fixed effects from AKM two-way fixed effects regressions (Abowd et al., 1999) to measure workers' ability and firm pay policy, we restrict the sample to the largest connected set of workers and firms.[2]

[1] The measure of earnings is gross of labour income taxes and pension contributions on the side of the employee.

[2] Connected groups contain all the individuals that have ever been employed at one of the firms in the group and all the firms that have ever hired one of the workers in the group. Thus, two groups are not connected if one person of the second group has never been employed by a firm of the first group and a firm in the first group has never employed a person of the second group (or viceversa). Since fixed effects are identified up to a normalising constant, different connected groups give fixed effects estimates that are not comparable across each other. Thus, we keep the largest connected group, only. This group comprises 94.5 percent of the original sample. Note that we estimate AKM worker and firm effects using the full sample of both men and women, although the main analysis will focus on the latter, only. We do this to maximise the size of the largest connected group and to reduce limited mobility bias (Andrews et al., 2008). Specifically, we estimate the following regression:

$$w_{it} = \alpha_i + \psi_{J(i,t)} + \mathbf{x}_{it}'\gamma + \varepsilon_{it}$$

Finally, since we are interested in female outcomes, we focus on the sample of women only and, specifically, on those who have their first child (their first maternity leave episode) between the age of 18 and 40. Therefore, as we follow them for a period of at most 15 years after childbirth, our sample comprises women between 18 and 55 years old. Columns 1 and 2 of Table 1 report descriptive statistics on the full sample of women after such restrictions and show that average annual earnings and full-time equivalent weekly wages are around 16,000 and 432 Euros, respectively. Women work on average 35 full-time equivalent weeks per year and 28 percent of them are employed with part-time contracts. Average age is 36. The majority of women is employed in white-collar occupations, in services and small firms.

3 Empirical Strategy

3.1 Placebo births

We follow closely the methodology discussed in appendix A of Kleven et al. (2019b) to identify a suitable control group of non-mothers. In our sample of 18-55 years old women, first, we focus on those born between 1945 and 1978, who are not yet 40 by 1985 (the first year in our sample) and who turn 40 by 2018. Among these women, we know who took maternity leave and, therefore, had a child during our observation period or even before, as we have the full social security contribution history of workers. Women born between 1945 and 1978 who do not have a child enter the group of "never mothers". Women born after 1978 are subject to right-censoring as they are not yet 40 by the end of the observation period and therefore might have a child after the last year of the sample (2018). We solve this truncation issue by assigning a birth probability to the truncated cohort. Specifically, we estimate a linear probability model in the non-truncated cohorts 1945-1978 by regressing a dummy taking value one for never mothers on the following set of dummy controls: quartiles of the cohort-specific log daily wage distribution, quartiles of the AKM worker fixed effects distribution, region of residence.[3] We then assign to women in the truncated birth cohorts the predicted probability of giving birth, based on the coefficients estimated in the linear

where w_{it} are log weekly wages of worker i in year t. α_i and $\psi_{J(i,t)}$ are worker and firm fixed effects, where the firm is indexed by $J(i,t)$. $\mathbf{x}'_{it}$ contains time-varying observables (cubic polynomials in age and tenure, occupation dummies, part-time dummy, and their interaction with a gender dummy) and year fixed effects. ε_{it} is an error term.

[3] In other words, we estimate the following regression:

$$\text{NeverMother}_{iT} = \alpha + X'_{it}\beta + \epsilon_{it}$$

where NeverMother_{iT} is a dummy equal to 1 for never mothers in birth cohorts 1945-1978 and X_{it} includes the dummy controls indicated in the text.

Table 1: Descriptive statistics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Full sample		Analysis sample							
			Women with children				Women without children			
			Before childbirth		After childbirth		Before childbirth		After childbirth	
	mean	(sd)	mean	(sd)	mean	(sd)	mean	(sd)	mean	(sd)
Annual earnings	16,086.98	(12574.82)	17,459.37	(10331.31)	16,488.68	(12565.66)	13,821.29	(10733.72)	17,517.84	(14301.64)
FTE weekly wage	431.91	(223.30)	424.02	(173.50)	460.84	(236.59)	403.51	(193.23)	458.23	(250.67)
FTE weeks worked	35.08	(17.22)	39.94	(15.16)	34.31	(15.70)	32.35	(18.29)	35.43	(17.62)
Raw weeks worked	39.46	(16.39)	42.92	(13.74)	40.69	(14.99)	34.93	(17.86)	39.12	(16.67)
Part-time	0.28	(0.45)	0.19	(0.40)	0.40	(0.49)	0.19	(0.39)	0.23	(0.42)
Age	35.50	(9.64)	28.09	(4.41)	36.12	(6.03)	28.47	(4.20)	37.82	(6.06)
Blue-collar	0.43	(0.50)	0.37	(0.48)	0.38	(0.49)	0.40	(0.49)	0.42	(0.49)
White-collar	0.51	(0.50)	0.55	(0.50)	0.58	(0.49)	0.56	(0.50)	0.56	(0.50)
Executive	0.01	(0.11)	0.01	(0.09)	0.02	(0.15)	0.00	(0.06)	0.02	(0.13)
ENE transitions	0.05	(0.22)	0.02	(0.14)	0.04	(0.19)	0.09	(0.28)	0.06	(0.23)
Centre-North	0.79	(0.41)	0.83	(0.38)	0.84	(0.37)	0.71	(0.45)	0.72	(0.45)
South	0.21	(0.41)	0.17	(0.38)	0.16	(0.37)	0.29	(0.45)	0.28	(0.45)
Electricity & Construction	0.02	(0.15)	0.03	(0.16)	0.03	(0.17)	0.02	(0.14)	0.02	(0.15)
Manufacturing	0.23	(0.42)	0.27	(0.44)	0.28	(0.45)	0.16	(0.37)	0.18	(0.38)
Wholesale & Retail	0.16	(0.36)	0.20	(0.40)	0.19	(0.39)	0.13	(0.34)	0.12	(0.32)
Transports & Logistics	0.03	(0.16)	0.03	(0.17)	0.03	(0.16)	0.02	(0.15)	0.03	(0.16)
Accomodation & Restaurants	0.10	(0.30)	0.08	(0.27)	0.07	(0.25)	0.12	(0.33)	0.10	(0.30)
Information & Communication	0.03	(0.17)	0.04	(0.20)	0.04	(0.20)	0.03	(0.17)	0.03	(0.17)
Finance & Insurance	0.04	(0.20)	0.05	(0.22)	0.06	(0.23)	0.03	(0.17)	0.04	(0.19)
Professional activities	0.04	(0.20)	0.05	(0.22)	0.05	(0.21)	0.03	(0.18)	0.03	(0.17)
Public Admin. & Education	0.10	(0.30)	0.03	(0.18)	0.04	(0.20)	0.22	(0.41)	0.21	(0.41)
Health	0.08	(0.27)	0.08	(0.26)	0.08	(0.27)	0.07	(0.26)	0.08	(0.27)
Other services	0.16	(0.37)	0.15	(0.35)	0.14	(0.34)	0.16	(0.36)	0.16	(0.37)
Firm size: 1-10 employees	0.44	(0.50)	0.44	(0.50)	0.38	(0.49)	0.45	(0.50)	0.37	(0.48)
Firm size: 11-100 employees	0.18	(0.39)	0.21	(0.41)	0.20	(0.40)	0.18	(0.38)	0.18	(0.39)
Firm size: 100+ employees	0.38	(0.49)	0.35	(0.48)	0.42	(0.49)	0.37	(0.48)	0.44	(0.50)
N. workers	1,106,592		202,921				409,238			
Person-year obs.	11,171,712		2,578,669				2,934,987			

Notes. The table reports means and standard deviations of variables for the full sample of women in the original dataset (columns 1-2) and for the sample of women with children (columns 3-6) and women without children (columns 7-10), before and after childbirth.

probability model. We then sort women born after 1978 based on such predicted probability and, starting from the largest value, we assign them to the control group up to the point in which the fraction of "predicted" never mothers in the truncated cohort post-1978 equals the fraction of actual never mothers in the non-truncated cohorts 1945-1978.[4] Therefore, the final sample consists of three groups of women: actual mothers, actual never mothers from birth cohorts 1945-1978 and predicted never mothers from birth cohorts 1979-2000. The latter two groups constitute the control group.

The second step is to assign a placebo year of birth to the control group of never mothers. We do so by assigning a placebo age at birth to non-mothers, by drawing from the actual distribution of age at birth for mothers. We distinguish again between actual and predicted never mothers. For actual never mothers, we assume that the distribution of age at birth $A_{c,q}$ follows a log-normal distribution within cells of birth cohort c and quartiles of worker fixed effects q, i.e. $A_{c,q} \sim \mathcal{LN}\left(\widehat{\mu}_{c,q}, \widehat{\sigma}_{c,q}\right)$, where mean $\widehat{\mu}_{c,q}$ and variance $\widehat{\sigma}_{c,q}$ are obtained from the actual within-cell distribution for mothers. We assign a random draw from this distribution to actual never mothers. For predicted never mothers, we use random draws from a distribution with same variance $\widehat{\sigma}_{c,q}$ but different mean $\widetilde{\mu}_{c,q}$, which is obtained by predicting age at birth from the estimation of a regression on a quadratic time trend for actual mothers, so to allow women born after 1978 to have their first child at an older age.[5]

Table 1 shows that our final sample consists of 202,921 women with children and 409,238 women without children in the control group. The total number of observations is 5,513,656. The table reports descriptive statistics for women with children in columns 3-6 and for women without children in columns 7-10 before and after (actual or placebo) childbirth. Before maternity, women with children earn 17.5 thousand Euro a year, while women without children earn 13.8 thousand Euro. This difference is determined by both a higher weekly wage and a higher number of weeks worked. After childbirth, the relationship reverses, with women with children earning 16.5 thousand Euro and women without children earning 17.5 thousand Euro. This diverging evolution of annual earnings between mothers and non-mothers already highlights the impact of childbirth on mothers' labour market outcomes. In contrast, there are no differences in terms of occupation between the two groups of women,

[4]By constraining the fraction of never-mothers in the truncated cohorts to be equal to that observed in the non-truncated cohorts, we are assuming that such fraction remains fairly constant across the two groups of cohorts (truncated and non-truncated) over time. This assumption is not unrealistic, as the fraction of never-mothers in the non-truncated cohort is declining, but relatively flat over time and by year of birth of the woman, as Figure A.1 in the Appendix shows.

[5]This adjustment is necessary, because of the truncation issue. As we do not observe completed fertility for truncated cohorts, age at birth would be skewed to the right if we did not make any adjustment to age at birth. By running a regression of age at birth on a quadratic time trend, and using the predicted age at birth as a reference for drawing the mean of the lognormal distribution for predicted never mothers, we allow age at birth to have the "correct" distribution.

but mothers are more likely to be in the Centre-North than non-mothers (83-84 percent vs 71-72 percent) and to work in manufacturing (27-28 percent vs 16-18 percent).

3.2 Event study

The Child Penalty For our analysis, we keep an unbalanced panel and condition on employment to study the impact of childbirth on female labour market outcomes. This choice can likely create a selection bias in our estimates, although the direction of the bias is not clear a priori. On the one hand, women that remain employed after childbirth may be those more in need of income or those more concerned about their career but less needy. To partially address differential selection in the treated group of mothers and in the control group of placebo non-mothers, we will exploit the panel structure of our data and control for individual fixed effects.

We use an event study specification and compare the evolution of outcomes for women with and without children. Specifically, we estimate the following models:

$$y_{its}^{G(i)} = \alpha_i + \sum_{k \neq -1} \beta_k^{G(i)} \cdot \mathbf{1}(k = s) + \sum_{y} \gamma_y^{G(i)} \cdot \mathbf{1}(y = t) + \varepsilon_{its}^{G(i)}, \tag{1}$$

where $y_{its}^{G(i)}$ is the outcome for individual i belonging to group $G(i)$ (women with or without children), year t and event time s (i.e. years relative to first childbirth). We regress these outcomes on individual fixed effects α_i, a full set of event time dummies $\sum_{k \neq -1} \mathbf{1}(k = s)$, with $k = \{-5, ..., 15\}$ and a set of year dummies $\sum_y \mathbf{1}(y = t)$, with $t = \{1985, ..., 2018\}$. Finally, $\varepsilon_{its}^{G(i)}$ is an error component. We plot the coefficients $\beta_k^{G(i)}$ separately for women with and without children. We cluster standard errors at the worker level. We first focus on log annual earnings as main outcome at the individual level and we define the long-run child penalty in log annual earnings as the difference in the event study coefficients fifteen years after childbirth for mothers relative to non-mothers. We then investigate what contributes to the child penalty in annual earnings, by using log weekly wages, log weeks worked, the probability of working part-time, the probability of being manager and the probability to move to non-employment as outcomes in equation (1).

Sorting The negative impact of motherhood on wages can depend on sorting of women with children into low-pay firms. In particular, women may move to firms that offer greater flexibility and better work-life balance in exchange of lower wages, or to less productive firms that require fewer hours of work or that rely less extensively on overtime. Casarico and Lattanzio (2019) show that sorting into low-pay establishments explains around 20 percent of the overall gender wage gap in Italy over the period 1995-2015 and Bruns (2019) provides evidence that moves to low-pay establishments contributes to explaining the child penalty in Germany. We complement these analyses by showing how childbirth impacts the type of firms where mothers work.[6] To do so, we use a subset of our data that allows us to match the administrative archives on workers' pay and employment histories with balance sheet information coming from Cerved. Specifically, we have information on value added, sales and capital for a subset of firms over the period 1994-2012. The sample size reduces from 5.5 million to 900,692. For each firm, we compute mean log value added per worker, mean log sales per worker, mean log capital per worker and mean log annual wages,[7] which we use as outcomes in the following specification:

$$\overline{y}^{G(i)}_{J(i,t),t,s} = \alpha_i + \sum_{k \neq -1} \beta^{G(i)}_k \cdot \mathbf{1}(k = s) + \sum_y \gamma^{G(i)}_y \cdot \mathbf{1}(y = t) + \varepsilon^{G(i)}_{J(i,t),t,s}, \qquad (2)$$

where $\overline{y}^{G(i)}_{J(i,t),t,s}$ is the log average outcome per worker (value added, sales, capital or annual wages) in firm $J(i,t)$, i.e. the firm that employs worker i in year t, in the event time s, for women with and without children ($G(i)$). The remaining variables are defined as in equation (1). Note that, as we use averages over time of outcomes, changes in their evolution after childbirth can happen only if women with and without children move differently between firms or to non-employment.[8] As before, we define the child penalty as the difference $\beta^M_{15} - \beta^{NM}_{15}$.

Heterogeneity In order to shed light on the potential factors that influence the size of the penalty following childbirth, we investigate how different worker and firm characteristics mediate the impact of motherhood on female annual earnings. Specifically, for each group H – where H identifies different worker or firm types –, we estimate a dynamic difference-in-differences model by estimating one single equation – instead of separate regressions for mothers and non-mothers – and including a dummy for mothers. We estimate, separately

[6] The evidence from France, provided by Wilner (2016) suggests a limited role of worker-firm matching in explaining the motherhood penalty, after accounting for worker's human capital, but does not focus on the type of firms women with and without children match with following childbirth, which is precisely our research question in this section.

[7] Over the same period, we have information on total employment by firm. Therefore, when we compute per worker quantities, we are measuring value added, sales and capital over total workforce, not only the sample of workers for which we have detailed individual-level administrative information.

[8] We also use averages over time in order to reduce noise in financial quantities and to interpolate missing values in cases in which a firm has gaps in the data.

for each group H:

$$w_{its}^H = \alpha_i + \sum_{k \neq -1} \left(\beta_{0,k}^H + \beta_k^H \cdot G_i\right) \cdot \mathbf{1}(k = s) + \sum_y \left(\gamma_{0,y}^H + \gamma_y^H \cdot G_i\right) \cdot \mathbf{1}(y = t) + \varepsilon_{its}^H, \quad (3)$$

where w_{its}^H are log annual earnings for worker i belonging to group H and time t and event time s. We report heterogeneous effects by total duration of the parental leave period (mandatory and optional),[9] by age at birth (below/above 30 years old), by AKM firm, peer and worker effects (above/below median), by occupation (blue- or white-collar), by firm size (1-100 or more than 100 employees) and by the share of female workers employed by the firm (above/below median). All variables in equation (3) are defined as before. G_i is a dummy equal to one for women with children, which we interact with event time dummies and year dummies. Therefore, $\beta_{0,k}^H$ measures average earnings in each event time k for women without children, conditional on individual fixed effects and year fixed effects. $\gamma_{0,y}^H$ and γ_y^H measure group-specific year effects, which we allow to vary between women with and without children. In the results section, we report estimates of β_k^H, which measure the child penalty in log annual earnings – the difference in log earnings for women with and without children – for each subgroup H and event time k, conditional on worker and year effects.

4 The Impact of Motherhood on Labour Market Outcomes

4.1 Estimates of the Long-run Child Penalty

Main Results Figure 1 reports the estimates of $\beta_k^{G(i)}$ separately for women with and without children. Panel A shows results for log annual earnings. Before childbirth, women with children experience slightly steeper earnings growth than women without children. After childbirth, women with children experience a sharp drop in annual earnings, which is particular evident in the year of childbirth and in the following one—a period during which mothers may take up months of optional parental leave, besides the 5 months mandatory

[9] The mandatory maternity leave lasts 5 months. There are no policy changes regarding mandatory maternity leave over the period of analysis. During mandatory leave, mothers receive 80 percent of their average daily wage in the month before the leave. After this period, parents can take an optional parental leave for up to 6 months until the child turns 8 years old before 2015 and 12 years old after 2015. Until the child turns 3 years old (6 years old after 2015), the parent receives a compensation equal to 30 percent of his or her wage for the period the parent is on parental leave. The total duration of the parental leave in a family cannot exceed 10 months, unless the father takes more than 3 months: in this case, the father-specific leave is extended to maximum 7 months and the family's total parental leave duration is extended to 11 months.

maternity leave. Starting from the second year after childbirth earnings of mothers increase as they return to work but remain flat at least until the fourth year from maternity, when they start to gradually increase with a comparable slope to that of non-mothers. After 15 years, earnings of non-mothers are 46 log points higher than in the year before childbirth, whereas earnings of mothers are 11 log points lower: this translates into a long-run child penalty of 57 log points. What contributes to the child penalty in annual earnings? Panel B shows the estimates using full-time equivalent log weekly wages as dependent variable. Immediately after childbirth, wages of mothers drop by 15 log points—a drop consistent with the replacement rate of mandatory maternity leave which is set to be 80 percent of the average daily wage in the month preceding the start of maternity leave: the drop is lower in magnitude as some employers integrate the pay of mothers on leave. In the long-run, the penalty in weekly wages experienced by mothers relative to non-mothers is 7 log points ($24.9 - 17.9$). Panel C shows the impact of maternity on the number of full-time equivalent weeks worked: for mothers, the impact is as large as 60 log points in the first year after childbirth. After 15 years the penalty amounts to 50 log points ($21.5 - (-28.6)$). Panel D further shows that after childbirth, the part-time share among mothers increases relative to non-mothers. Two years after childbirth, the share of part-time contracts among mothers is 14 percentage points larger than among non-mothers, a difference that reaches its maximum level around 7 years after childbirth and, in the long-run, amounts to 20 percentage points.

Figure 2 shows a decomposition of the total child penalty in earnings, i.e. the difference at each point in time between the log earnings of women with children and log earnings of women without children, into the separate contribution of change in weekly wages, in the number of weeks worked and shift to part-time. The contribution of wages is equivalent to the difference in the curves for mothers and non-mothers from Figure 1, panel A. We then compute the contribution of weeks by estimating the child penalty on unadjusted weeks (differently from Figure 1, panel C, where we use full-time equivalent weeks as dependent variable). The difference in the child penalty between full-time equivalent and unadjusted weeks is due to hours reductions related to the switch to part-time contracts, which forms the residual contribution to the overall child penalty. The figure shows that most of the child penalty stems from reductions in labour supply by mothers: the decrease in weeks worked accounts for 67 percent of the child penalty in annual earnings, whereas the decrease in hours worked due to the shift to part-time contracts accounts for 21 percent. Reductions in full-time equivalent wages only account for a smaller fraction, around 12 percent of the total effect on earnings.

The effect on wages, albeit small, may be the consequence of lower career progression, as Figure 3, panel A, suggests. The figure estimates equation (3) using a dummy for executives

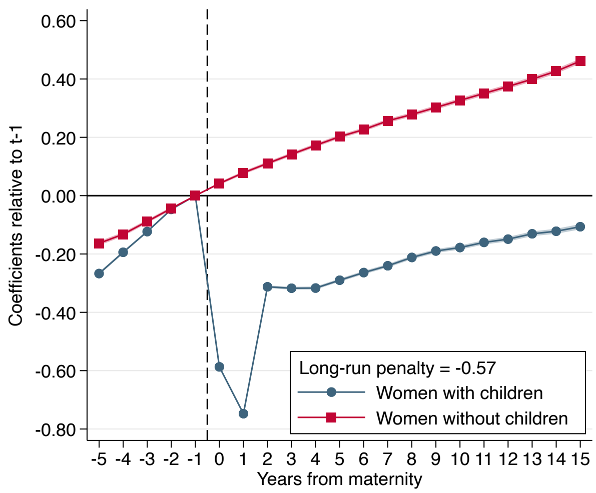

(A) Log annual earnings

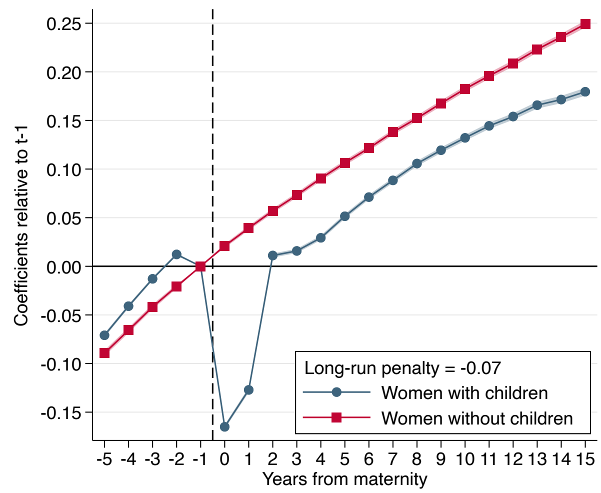

(B) Log weekly wages

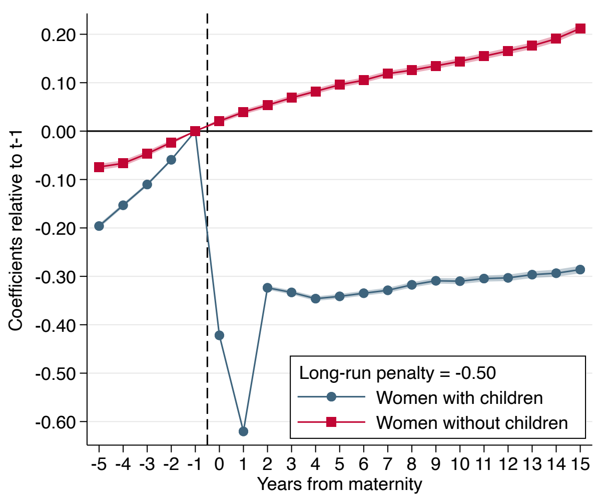

(C) Log FTE weeks worked

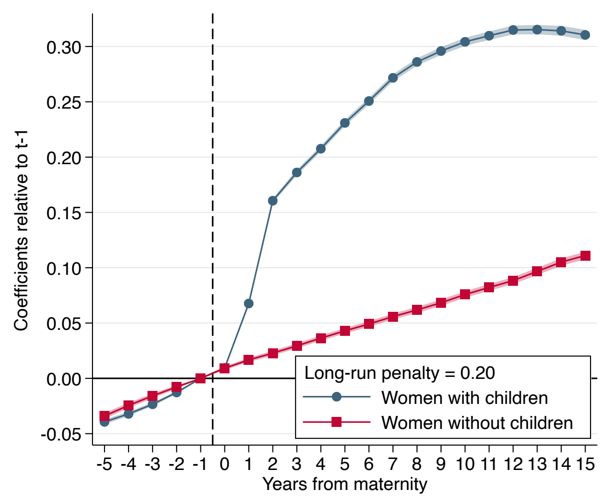

(D) Share part-time

Figure 1: Event study estimates of the impact of first childbirth on female labour market outcomes

Notes. The figures report event study coefficients $\beta_k^{G(i)}$ from equation (1) separately for women with (M) and without (NM) children. The long-run penalty reported in each graph is the difference in coefficients fifteen years after childbirth, $\beta_{15}^{M} - \beta_{15}^{NM}$. Confidence intervals at 95 percent level are obtained from worker-level cluster-robust standard errors.

as dependent variable. Fifteen years after childbirth, mothers have a 0.4 percentage point lower probability of becoming executives relative to non-mothers. Considering that the cross-sectional average share of women employed as managers is 1.5 percent, the effect is quite sizeable in magnitude.

The labour supply effect captured by the decomposition presented in Figure 2 does not consider women moving out of employment after maternity. Figure 3, panel B, shows that

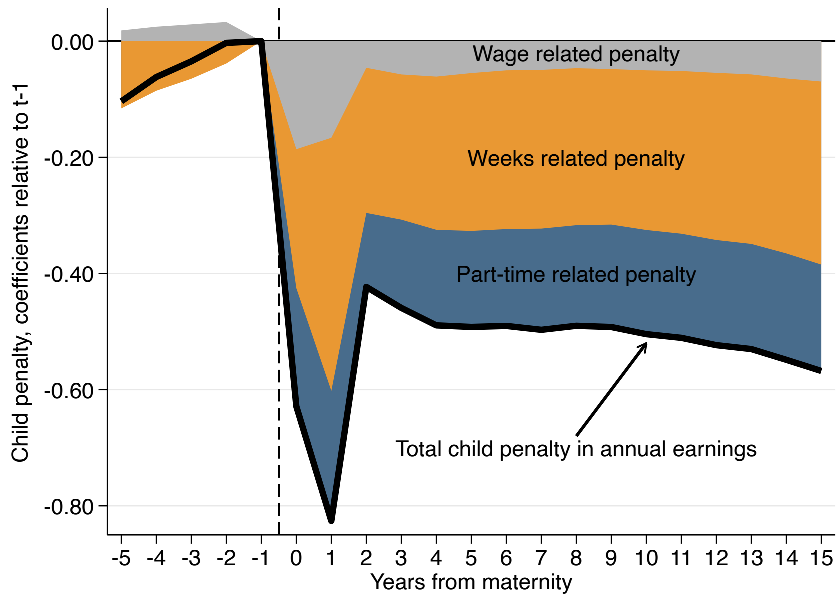

Figure 2: Decomposition of the child penalty into the contribution of reduction in FTE weekly wages, reduction in weeks worked and transition to part-time

Notes. The figure reports the child penalty in annual earnings, and its decomposition into the contribution of reduction in weekly wages, reduction in weeks worked and switch to part-time. The solid line shows the child penalty in annual earnings, i.e. the difference in event study coefficients at each event time $k = \{-5, ..., 15\}$ between mothers and non-mothers from equation (1). The grey and orange areas are obtained by estimating equation (1) using log weekly wages and log raw weeks worked as outcomes, and measure the contribution of reductions in weekly wages and weeks worked to the child penalty. The blue area, i.e. the contribution of part-time, is the difference between the penalty estimated from log adjusted weeks and the one obtained with log raw weeks worked.

women with children are more likely to transition to non-employment[10] following childbirth. The probability of employment to non-employment transitions increases soon after childbirth and it remains higher than the pre-childbirth level until 15 years after maternity. For non-mothers, instead, the probability to move to non-employment follows a decreasing trend and it is 12 percentage points lower in the long-run than for mothers.

Sorting Figure 4 reports the event study coefficients from equation (2), which uses firm-level outcomes to investigate whether women with and without children sort into different firms following childbirth. Panel A shows results for average log value added per worker. Before childbirth, women with and without children display similar trajectories, meaning that they work in firms with parallel evolution of average log value added per worker. After

[10]Given the nature of our data we are unable to tell whether women that exit our sample effectively go into unemployment or move to self-employment or public employment, as we only have information on employment in the private sector. Hence, when we refer to non-employment we refer to non-employment in the private sector.

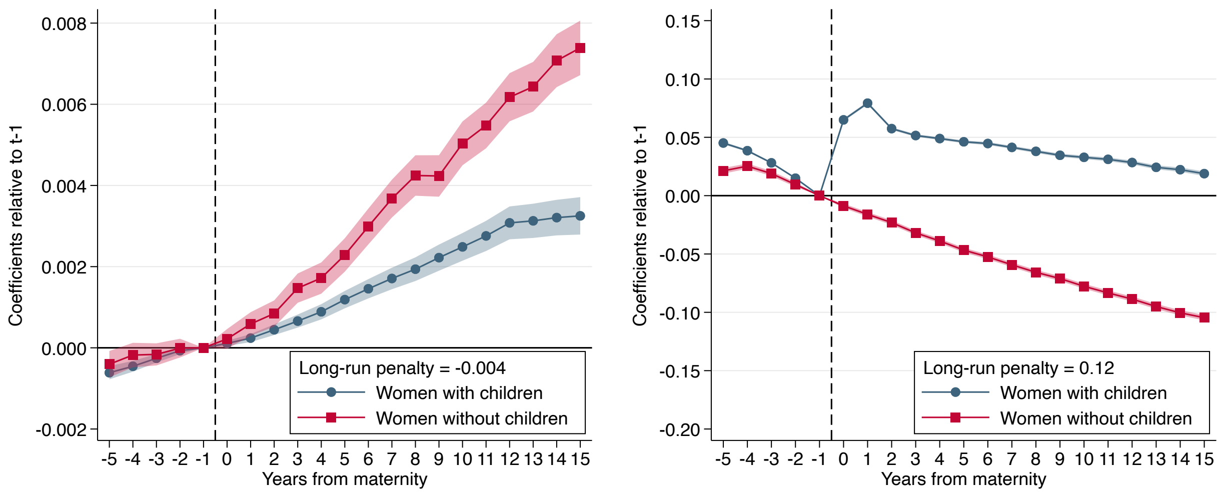

(A) Share executives (B) Employment to non-employment transitions

Figure 3: Event study estimates on the impact of first childbirth on the probability of become executive and on employment to non-employment transitions

Notes. The figures report event study coefficients $\beta_k^{G(i)}$ from equation (1) separately for women with (M) and without (NM) children. The long-run penalty reported in each graph is the difference in coefficients fifteen years after childbirth, $\beta_{15}^{M} - \beta_{15}^{NM}$. The dependent variable in panel A is a dummy variable equal to one for executives and in panel B a dummy equal to one for workers observed in year t but not in year $t+1$. Confidence intervals at 95 percent level are obtained from worker-level cluster-robust standard errors.

childbirth, women without children experience better outcomes with respect to women with children. In other terms, after childbirth, women without children tend to work for relatively better firms – in terms of log value added per worker – than women with children. After 15 years, this difference amounts to 3.2 log points. We reach a similar conclusion if we look at mean log sales per worker in panel B. Also in this case, childbirth represents a shock for mothers, who after their first child tend to work for firms with lower average sales relative to the period before childbirth, while non-mothers sort into relatively better firms. After 15 years, the difference widens by 5.2 log points relative to the year before childbirth. Panel C shows that firms where mothers work have lower log capital per worker compared to firms employing non-mothers. The long-run penalty equals 6.5 log points. Finally, panel D shows that firms employing women with children have lower average annual wages, by as much as 2.9 log points relative to firms employing women without children, 15 years after childbirth. All four panels highlight the presence of a considerable sorting pattern following childbirth, with mothers moving towards firms with lower productivity, revenues, capital and wages compared to non-mothers. Demand-side factors, besides supply-side mechanisms as those outlined above, may be at play as well. In the presence of taste or statistical discrimination, employers from more productive firms (i.e., with higher wages, value added, sales or capital) may discriminate against women with children if they believe they will be less productive after childbirth (see, e.g., Blau and Kahn, 2017, and Altonji and Blank, 1999, for reviews on the topic).

As highlighted above, the estimates for firm-level variables are conducted on a subset of data from the original sample. However, this sample selection makes little difference in terms of the main child penalty estimate, which amounts to 48 log points, as shown in Panel A of Figure A.2 in the Appendix. The lower coefficient is determined by a smaller penalty in terms of full-time equivalent weeks (42 log points), whereas the penalty in log weekly wages is of similar magnitude (7 log points), as is the part-time penalty (24 percentage points).

4.2 Heterogeneity

Figure 5 reports estimates of β_k^H from equation (3) based on subgroups defined by worker characteristics. Panel A reports heterogeneous effects by duration of the parental leave.[11] Longer periods of leave could harm labour market prospects of mothers as staying out of employment for longer periods may destroy firm-specific human capital and slow career progression. The figure plots the child penalty in annual earnings for mothers taking less or more than 6 months of leave relative to non-mothers. The duration is computed as the sum of the mandatory maternity leave and the optional parental leave. The figure shows that the child penalty is considerably larger for women taking longer leaves. The short-run drop is more severe, as women stay out of employment for a longer period of time and this translates into a larger long-run penalty, as it amounts to 48 log points for shorter leave duration and 66 log points for longer leaves after 15 years. Panel B shows heterogeneous effects between low- and high-wage workers. Specifically, we compute the median of the cross-sectional distribution of the AKM worker effect estimated on the full sample of workers and define low- and high-wage workers as women being below and above the median of the female worker fixed effects, respectively. The figure shows that the penalty is larger for low-wage workers (61 log points) than for high-wage workers (45 log points). The former may experience a larger penalty because of lower labour force attachment, worse outside options or because of lower human capital accumulation before childbirth, which could make it easier to experience career interruptions. Panel C shows that similar results are obtained when distinguishing blue- and white-collar workers. To the extent that skills and tasks are correlated, it comes as no surprise that white-collar workers pay a lower penalty after childbirth: as a descriptive evidence on this point, the average worker fixed effect of white-collar women is 16 log points larger than the average worker fixed effect of blue-collar female workers. Finally, panel D

[11] See footnote 9 for details on the legislation of maternity and parental leaves.

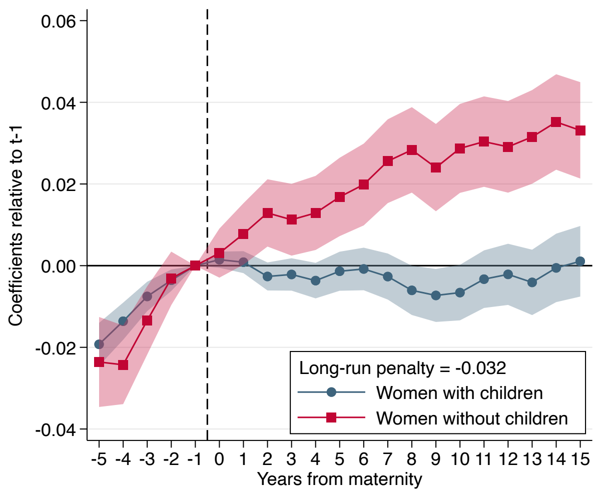

(A) Average log value added per worker

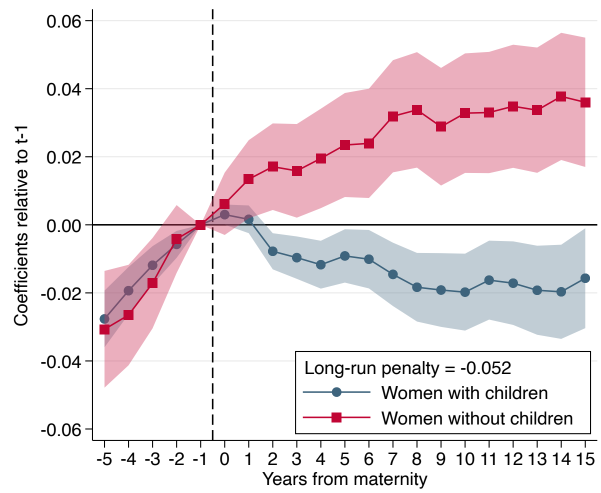

(B) Average log sales per worker

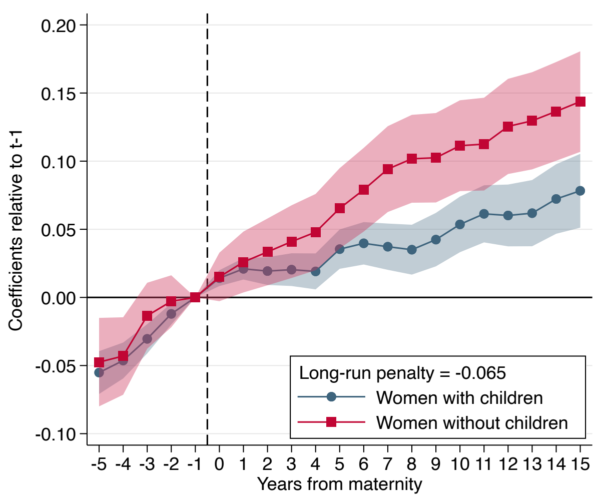

(C) Average log capital per worker

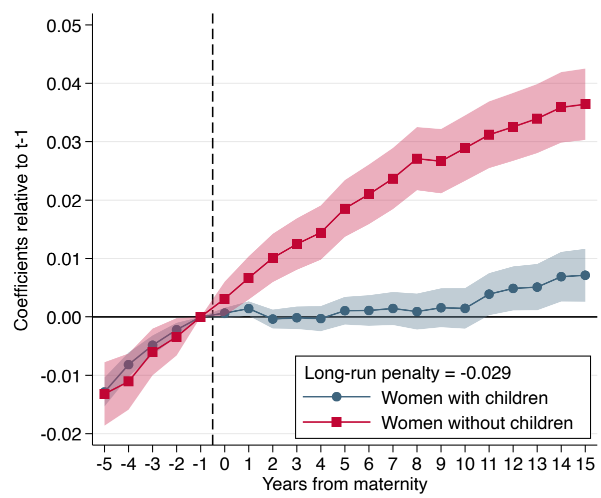

(D) Average log annual earnings

Figure 4: Event study estimates of the impact of first childbirth on firm-level outcomes

Notes. The figures report event study coefficients $\beta_k^{G(i)}$ from equation (2) separately for women with (M) and without (NM) children. The long-run penalty reported in each graph is the difference in coefficients fifteen years after childbirth, $\beta_{15}^{M} - \beta_{15}^{NM}$. The dependent variables are firm-level averages over time. Confidence intervals at 95 percent level are obtained from worker-level cluster-robust standard errors.

shows that women having their first child at a younger age pay a larger penalty than those having their child after 30. Younger mothers may not be able to complete their education and therefore may return to low-pay jobs after childbirth or have lower opportunities of career progression.

Figure 6 reports heterogeneous effects according to firm characteristics. The firm where a woman is employed may impact on the magnitude of the child penalty in a few ways.

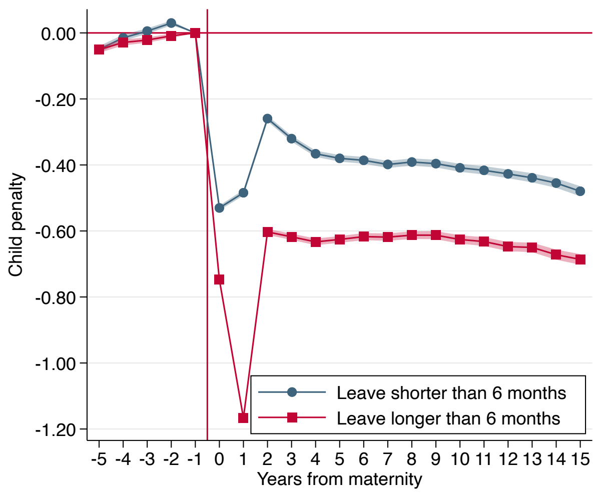

(A) By leave duration

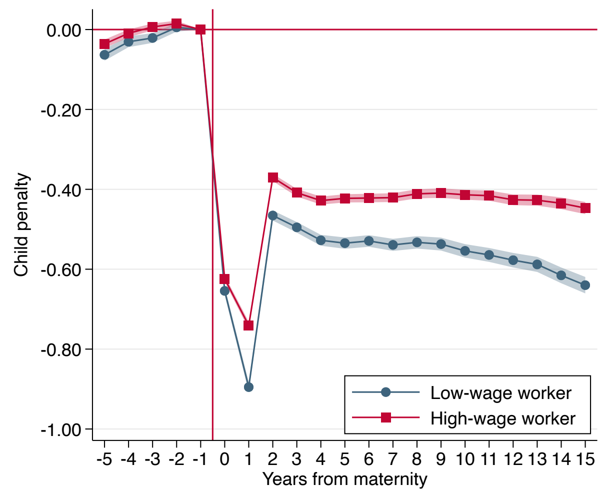

(B) By AKM worker effect

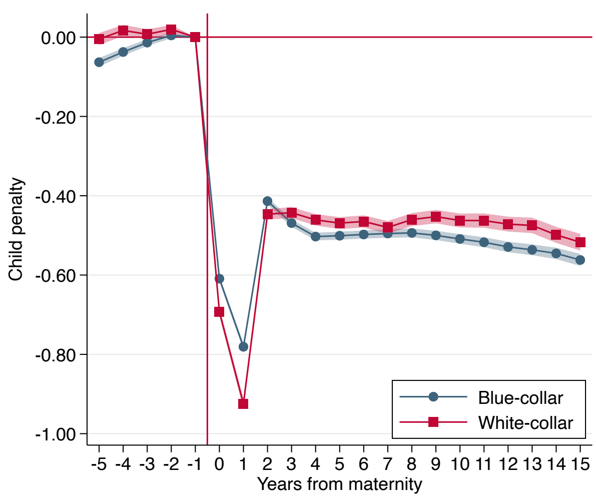

(C) By occupation

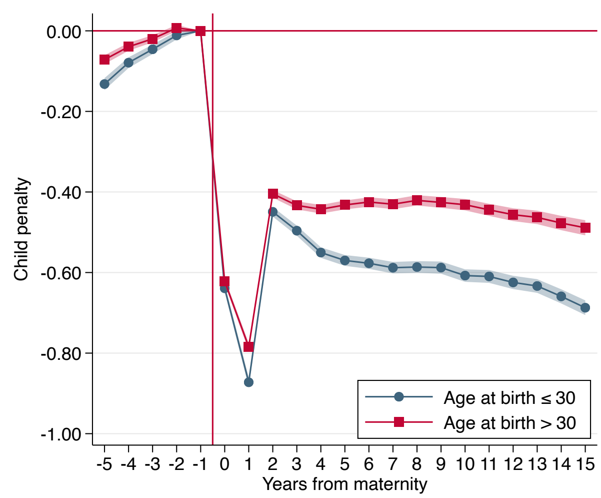

(D) By age at birth

Figure 5: Heterogeneous effects by worker characteristics

Notes. The figures report event study coefficients β_k^H from equation (3), i.e. child penalty estimates at different event time for different group of workers H. The dependent variable is log annual earnings. Confidence intervals at 95 percent level are obtained from worker-level cluster-robust standard errors.

In panels A and B, we investigate how the child penalty varies with the "quality" of the firm, measured by whether the firm is below or above the median AKM firm effect (low- and high-wage firms, respectively), or whether peer quality in the firm is below or above the median of peer quality distribution (where peer quality is measured as the leave one out average AKM worker effect of peers of the focal worker, i.e. the woman with or without children). Both panel A and B indicate that the child penalty is larger in low-quality firms, i.e. in firms that are in the bottom half of the firm effect distribution or in firms with low quality peers. The difference in the child penalty in the two groups of firms can be due to both demand-side and supply-side factors. On the demand-side, low-quality firms may offer fewer non-pecuniary benefits to women with children, e.g. company nursery, flexibility in the work schedule, working from home. On the supply-side, there could be sorting of low-skill mothers in such firms, e.g. because of fewer outside options and lower job search intensity or higher commuting costs, which could explain the larger child penalty (Carta and Rizzica, 2018; Le Barbanchon et al., 2020).

Panel C shows that the child penalty is remarkably similar in small and large firms, defined as those with less and more than 100 employees, respectively. The long-run penalty is only slightly larger (2 log points) in firms employing more than 100 employees. Finally, Panel D shows that the penalty is larger in firms with a female share higher than the median across firms. This result may seem puzzling at first as one could think that a higher female presence is an indicator of a higher women friendliness of the firm and one would expect a lower penalty in these firms. However, it may well be that firms with a higher share of female workers are relatively low performing firms, as women tend to sort into establishments that pay all workers less (Fanfani, 2018; Card et al., 2016; Casarico and Lattanzio, 2019; Coudin et al., 2018). Therefore, the distinction between firms with a high or low female share may partly capture quality and productivity differences across firms, which would explain the stark similarity between child penalties in panels A, B and D.

4.3 Gender Norms and Childcare Services

The child penalty varies considerably across different areas of the country. It tends to be larger in Southern regions, both if one measures it as the gap between mothers and non-mothers in labour market earnings and if one focuses on the gap in employment to non-employment transitions. The variation across Italian regions can be partly correlated with the presence of gender norms and stereotypes that attribute to the mother a greater share of household chores, which itself may be correlated with the presence and capillarity of childcare services. We provide descriptive evidence of these phenomena in Figures 7 and 8, which show scatter plots of the long-run child penalty[12] in annual earnings (panel A) or employment to non-employment transitions (panel B) against gender stereotypes (Figure 7) and the presence of childcare services (Figure 8). We measure gender stereotypes as the share of respondents who agree or strongly agree with the following statement from the 2017 wave of the European Values Study: "A man's job is to earn money; a woman's job

[12]The long-run child penalty is the event study coefficient in the 15th year following childbirth.

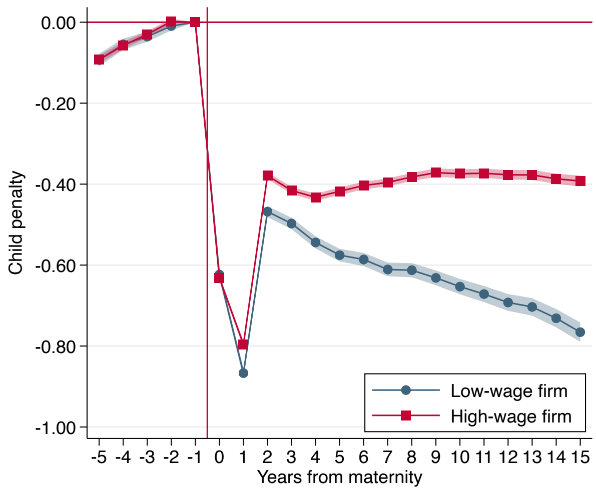

(A) By AKM firm effect

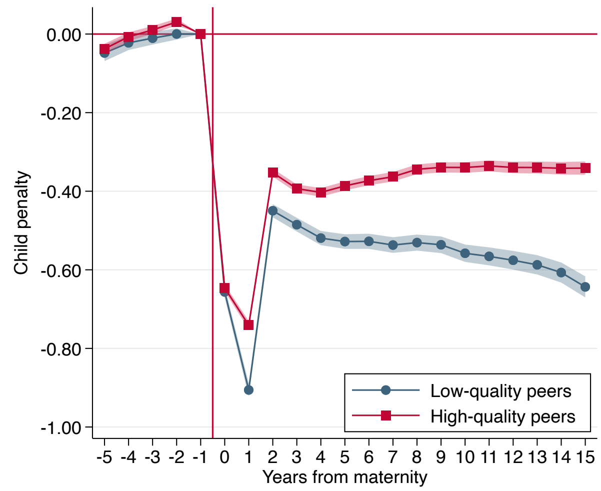

(B) By AKM peer effect

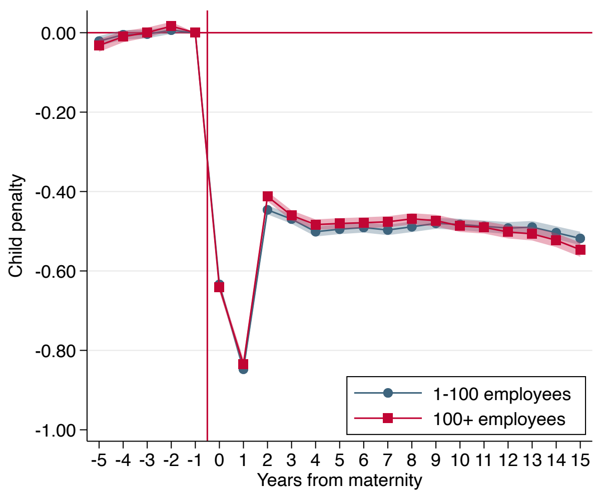

(C) By firm size

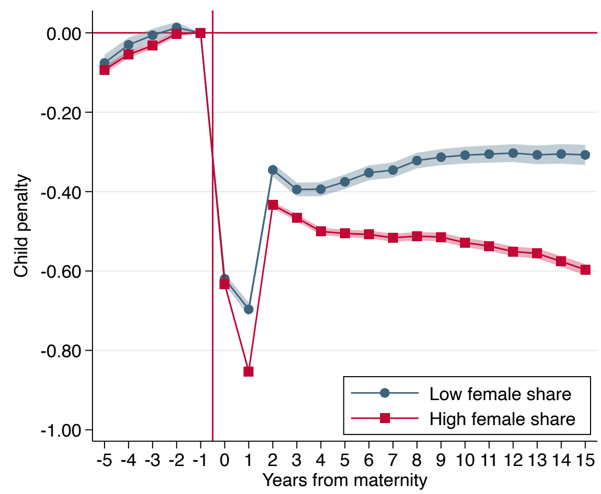

(D) By female share at the firm

Figure 6: Heterogeneous effects by firm characteristics

Notes. The figures report event study coefficients β_k^H from equation (3), i.e. child penalty estimates at different event time for different group of workers H. The dependent variable is log annual earnings. Confidence intervals at 95 percent level are obtained from worker-level cluster-robust standard errors.

is to look after the home and the family".[13] Figure 7 shows in panel A that there exists a negative correlation between such measure of gender stereotypes and the child penalty in annual earnings: in regions where stereotypes are stronger women with children pay a larger penalty in terms of reduced labour market earnings.[14] Panel B shows that there exists a positive correlation with the child penalty in employment to non-employment transitions:

[13] Specifically, we computed weighted averages for each region using the calibration weights provided by the EVS itself.

[14] This is in line with cross-country evidence provided in Kleven et al. (2019b).

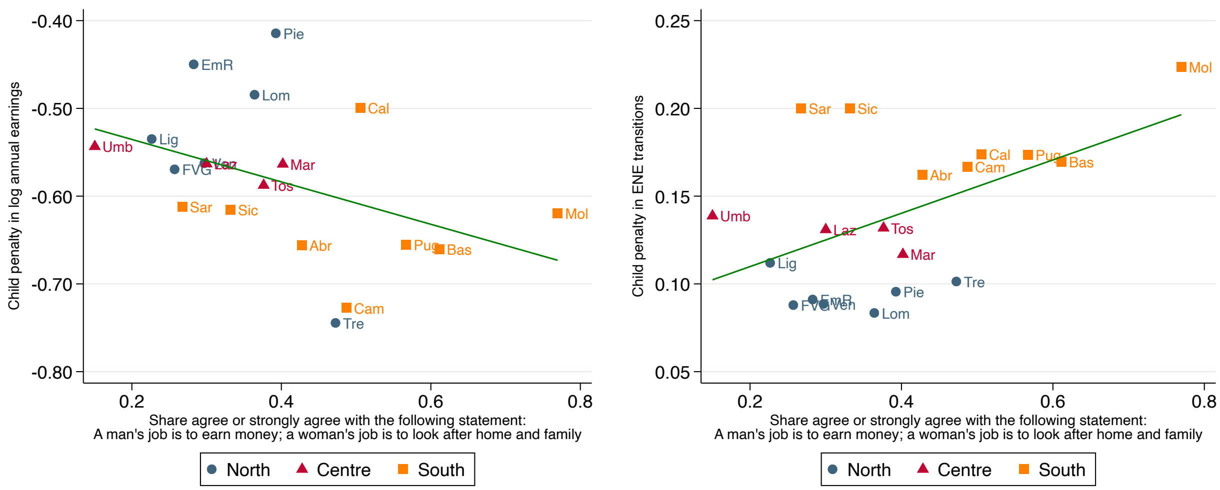

(A) Child penalty in annual earnings

(B) Child penalty in ENE transitions

Figure 7: Correlation between child penalties in annual earnings and ENE transitions and gender norms in Italian regions

Notes. The figures report scatter plots of the relationship between the child penalty in annual earnings (panel A) or in employment to non-employment transitions (panel B) and the share of respondent agreeing or strongly agreeing with the following statement from the 2017 European Values Study: "A man's job is to earn money; a woman's job is to look after the home and the family". The child penalty is measured as the difference in log annual earnings or employment to non-employment transitions 15 years after childbirth, obtained from separate regressions for each Italian region. The share of respondents agreeing or strongly agreeing with the statement is weighted with the calibration weights provided in the survey.

where gender stereotypes are stronger, the likelihood of transitions to non-employment of mothers relative to non-mothers tends to be larger. The said relationship is not specific to this question only from the European Values Study, but holds across a range of different questions, as Figures A.3 and A.4 in the Appendix show for earnings and employment to non-employment transitions, respectively.

Figure 8 shows correlations with the number of places in childcare services per 100 children between 0 and 2 years old. Data on public and private childcare slots are taken from Istat, the National Statistical Institute. The figure shows that the correlation between the child penalty in annual earnings and the availability of childcare services is positive: in regions where childcare facilities are more easily available, the child penalty is lower. On the other hand, the correlation with the child penalty in transitions to non-employment is positive, such that mothers experience lower job churning rates in regions where the capillarity of childcare services is higher. Both figures also show that there exists a North-South gradient: regions in the North tend to have less stereotypical attitudes towards women and smaller child penalties.

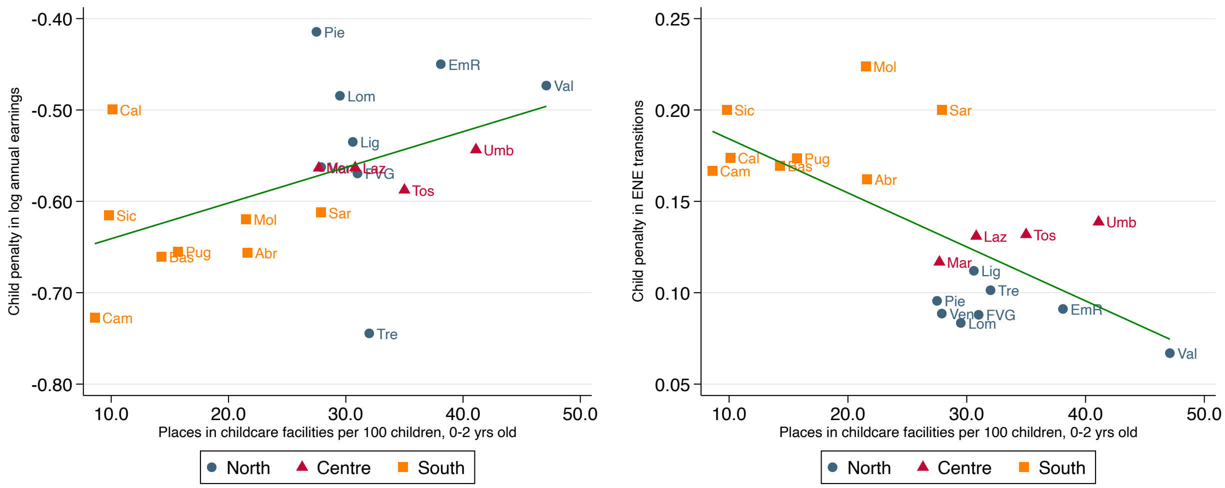

(A) Child penalty in annual earnings

(B) Child penalty in ENE transitions

Figure 8: Correlation between child penalties in annual earnings and ENE transitions and availability of childcare services in Italian regions

Notes. The figures report scatter plots of the relationship between the child penalty in annual earnings (panel A) or in employment to non-employment transitions (panel B) and the number of public and private places per 100 children aged 0-2 in nursery schools (source: Italian Statistical Institute). The child penalty is measured as the difference in log annual earnings or employment to non-employment transitions 15 years after childbirth, obtained from separate regressions for each Italian region.

5 Conclusions

We study the short- and long-run impact of motherhood on female labour market outcomes and provide evidence on individual, firm-level, institutional and cultural factors, that influence it, highlighting the multidimensionality of the child penalty. In particular, the analysis of firm-level characteristics of child penalties is a novel aspect we bring to the literature on the costs of motherhood. We show that the long-run child penalty in annual earnings is 57 log points and it is mainly determined by a reduction in the labour supply along the intensive margin. Also, the birth of a child increases the probability of transition to non employment. For mothers who stay on the labour market, there is a marked sorting in firms with lower productivity, sales, capital and wages. This complements evidence that sorting is an important component of gender wage gaps, by identifying childbirth as a trigger of changes in labour market matching. There is also evidence of heterogeneous effects according to worker and firm characteristics: child penalties in annual earnings are larger for young, low-wage mothers and those taking longer parental leaves. They are also larger in firms with less generous pay policy, worse peers and a higher share of female workers. Also cultural and institutional factors influence the size of the child penalty: in more gender-conservative regions or where childcare services are scarcer, the child penalty is larger. This evidence, albeit descriptively, seems to highlight the role of gender norms and the scarcity of childcare services as factors reinforcing child penalties in the labour market. Bertrand (2020) highlights the prescriptive role of gender stereotypes which determine different educational choices and labour market careers of men and women that favour the presence and persistence of gender gaps in earnings and employment. The evidence on the impact of childcare provision is, instead, mixed. Olivetti and Petrongolo (2017) highlights the key role of childcare in taming gender gaps in employment and earnings. This evidence is confirmed by some micro-level studies, such as Carta and Rizzica (2018) and Baker et al. (2008), but others find that increased availability of childcare services has mainly a crowding out effect on alternative informal childcare, leaving female labour supply almost unaffected (Fitzpatrick, 2010; Goux and Maurin, 2010). The evidence on the relationship between child penalty in earnings and availability of childcare is more scant. A study on Austria by Kleven et al. (2020) highlights that in municipalities with more childcare facilities the child penalty in annual earnings is significantly smaller than in those with lower childcare availability, but causal estimates of the impact of childcare expansion – that control for non-random selection of mothers into places with better childcare availability – suggest a negligible impact on the child penalty. In contrast, Nix and Andresen (2019) find that childcare expansions in Norway have a sizeable impact on the child penalty. Results may therefore differ depending on the overall institutional and cultural context, and this observation opens avenues for future research on this topic.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.

Altonji, J. G. and Blank, R. M. (1999). Race and Gender in the Labor Market. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3 of *Handbook of Labor Economics*, chapter 48, pages 3143–3259. Elsevier.

Andrews, M. J., Gill, L., Schank, T., and Upward, R. (2008). High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias? *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171(3):673–697.

Angelov, N., Johansson, P., and Lindahl, E. (2016). Parenthood and the Gender Gap in Pay. *Journal of Labor Economics*, 34(3):545–579.

Baker, M., Gruber, J., and Milligan, K. (2008). Universal child care, maternal labor supply, and family well-being. *Journal of Political Economy*, 116(4):709–745.

Becker, G. (1991). *A Treatise on the Family*. Harvard University Press.

Becker, G. S. (1985). Human Capital, Effort, and the Sexual Division of Labor. *Journal of Labor Economics*, 3(1):S33–S58.

Bertrand, M. (2020). Gender in the Twenty-First Century. *AEA Papers and Proceedings*, 110:1–24.

Blau, F. D. and Kahn, L. M. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55(3):789–865.

Bruns, B. (2019). Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap. *American Economic Journal: Applied Economics*, 11(2):74–113.

Card, D., Cardoso, A. R., and Kline, P. (2016). Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *Quarterly Journal of Economics*, 131(2):633–686.

Carta, F. and Rizzica, L. (2018). Early kindergarten, maternal labor supply and children's outcomes: Evidence from italy. *Journal of Public Economics*, 158:79–102.

Casarico, A. and Lattanzio, S. (2019). What Firms Do: Gender Inequality in Linked Employer-Employee Data. Cambridge Working Papers in Economics 1966, Faculty of Economics, University of Cambridge.

Cortés, P. and Pan, J. (2020). Children and the Remaining Gender Gaps in the Labor Market. Working Paper 27980, National Bureau of Economic Research.

Coudin, E., Maillard, S., and To, M. (2018). Family, Firms and the Gender Wage Gap in France. IFS Working Papers W18/01, Institute for Fiscal Studies.

de Quinto, A., Hospido, L., and Sanz, C. (2020). The child penalty in Spain. Occasional Papers 2017, Banco de España.

EIGE (2020). Gender Equality Index 2020. Technical report, European Institute for Gender Equality.

Fanfani, B. (2018). Tastes for Discrimination in Monopsonistic Labour Markets. Department of Economics and Statistic Working Paper Series 54, University of Torino.

Fitzpatrick, M. D. (2010). Preschoolers Enrolled and Mothers at Work? The Effects of Universal Prekindergarten. *Journal of Labor Economics*, 28(1):51–85.

Goux, D. and Maurin, E. (2010). Public school availability for two-year olds and mothers' labour supply. *Labour Economics*, 17(6):951–962.

Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimuller, J. (2019a). Child penalties across countries: Evidence and explanations. *AEA Papers and Proceedings*, 109:122–26.

Kleven, H., Landais, C., Posch, J., Steinhauer, A., and Zweimüller, J. (2020). Do Family Policies Reduce Gender Inequality? Evidence from 60 Years of Policy Experimentation. Working Paper 28082, National Bureau of Economic Research.

Kleven, H., Landais, C., and Søgaard, J. E. (2019b). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2020). Gender Differences in Job Search: Trading off Commute against Wage. *The Quarterly Journal of Economics*, 136(1):381–426.

Martino, E. M. (2020). The Labor Cost of Motherhood: Is a Shorter Leave Helpful? Working paper, Mimeo.

Nix, E. and Andresen, M. E. (2019). What Causes the Child Penalty? Evidence from Same Sex Couples and Policy Reforms. Discussion Papers 902, Statistics Norway, Research Department.

Olivetti, C. and Petrongolo, B. (2017). The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries. *Journal of Economic Perspectives*, 31(1):205–30.

Wilner, L. (2016). Worker-firm matching and the parenthood pay gap: Evidence from linked employer-employee data. *Journal of Population Economics*, 29(4):991–1023.

A Additional Figures

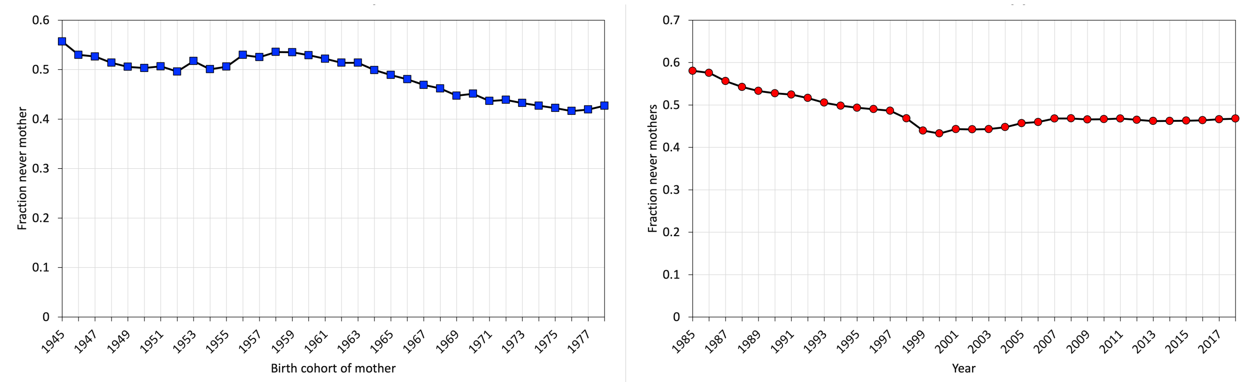

(A) By cohort (B) Over time

Figure A.1: Fraction of never-mothers in the non-truncated cohorts

Notes. The figure shows the share of never-mothers among female employees in the non-truncated cohorts by year of birth of the woman (panel A) and over time (panel B).

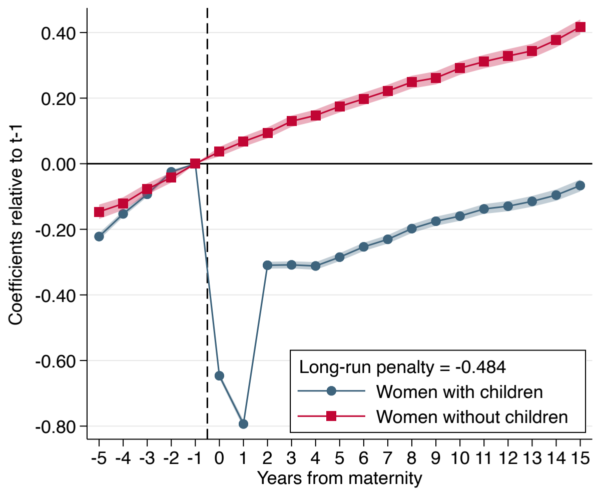

(A) Log annual earnings

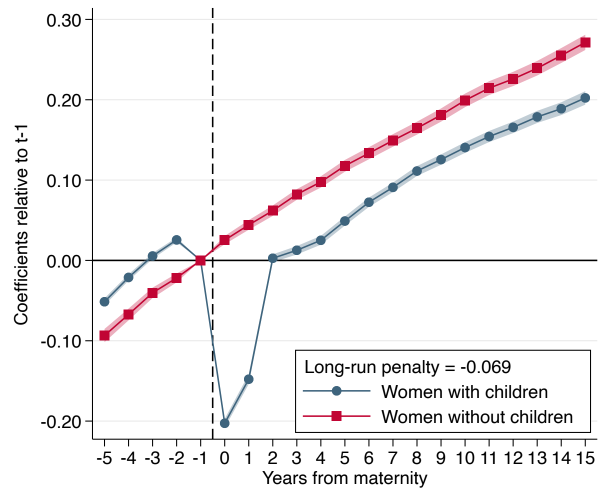

(B) Log weekly wages

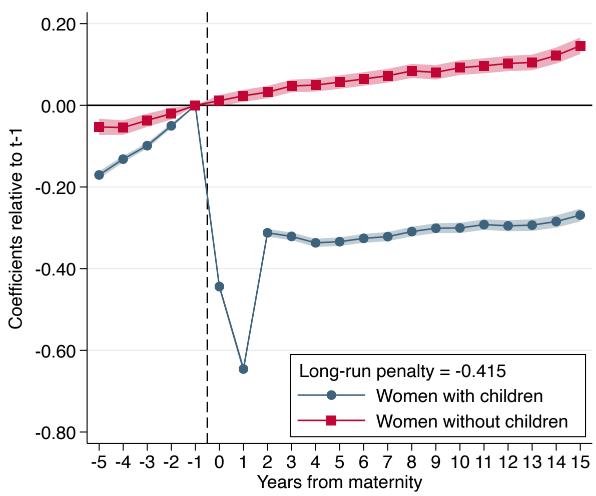

(C) Log FTE weeks worked

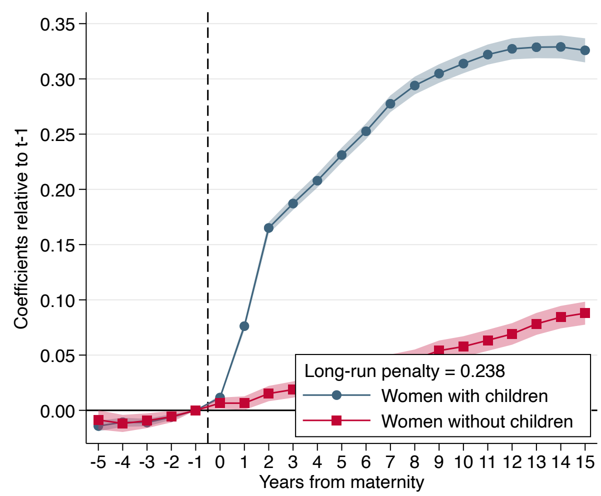

(D) Share part-time

Figure A.2: Event study estimates of the impact of first childbirth on female labour market outcomes, sample matched with firms' balance sheets

Notes. The figures report event study coefficients $\beta_k^{G(i)}$ from equation (1) separately for women with (M) and without (NM) children, focusing on the subsample of the data with non-missing balance sheet information (see section 3.2 for details). The long-run penalty reported in each graph is the difference in coefficients fifteen years after childbirth, $\beta_{15}^{M} - \beta_{15}^{NM}$. Confidence intervals at 95 percent level are obtained from worker-level cluster-robust standard errors.

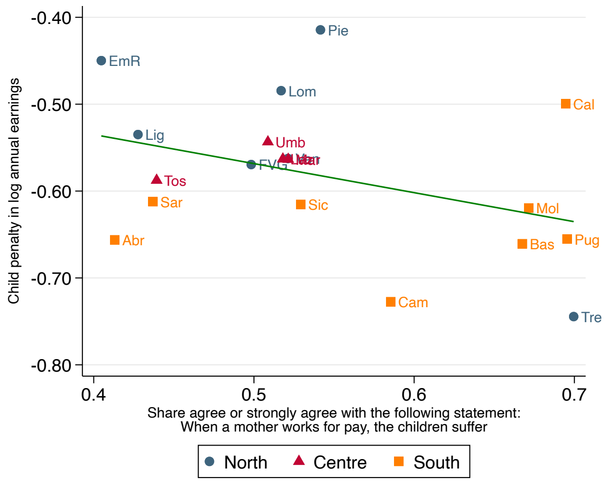

(A) When a mother works for pay the children suffer

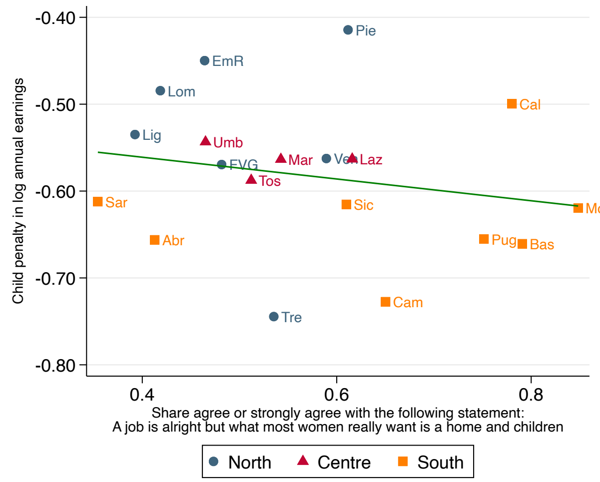

(B) A job is alright but what most women really want is home and children

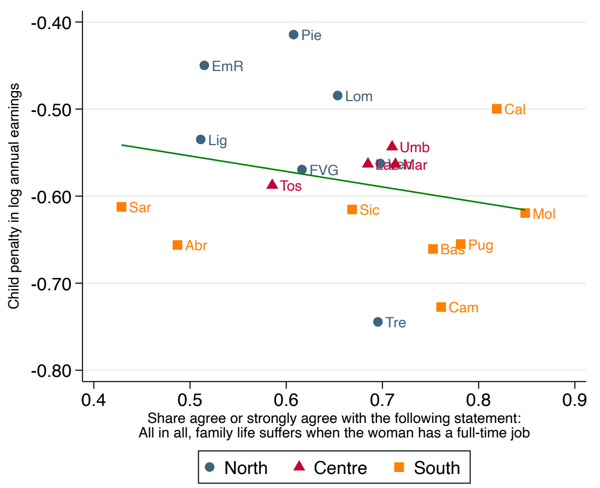

(C) All in all, family life suffers when the woman has a full-time job

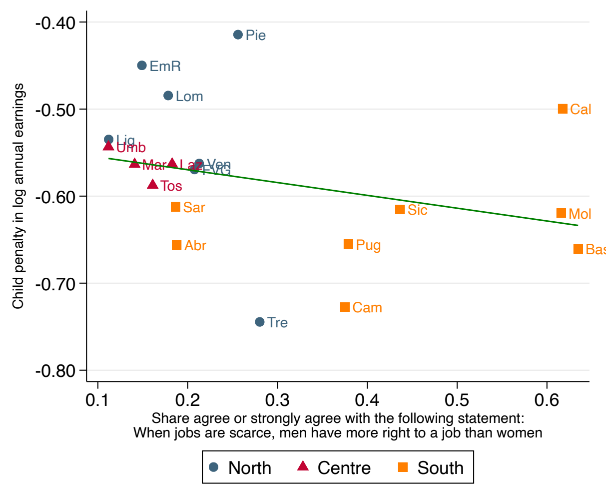

(D) When jobs are scarce, mean have more right to a job than women

Figure A.3: Correlation between child penalties in annual earnings and gender norms, from different EVS questions, in Italian regions

Notes. The figures report scatter plots of the relationship between the child penalty in annual earnings and the share of respondents agreeing or strongly agreeing with four alternative statements from the European Values Study 2017 (reported on the horizontal axis). The child penalty is measured as the difference in log annual earnings or employment to non-employment transitions 15 years after childbirth, obtained from separate regressions for each Italian region. The share of respondents agreeing or strongly agreeing with the statements is weighted with the calibration weights provided in the survey.

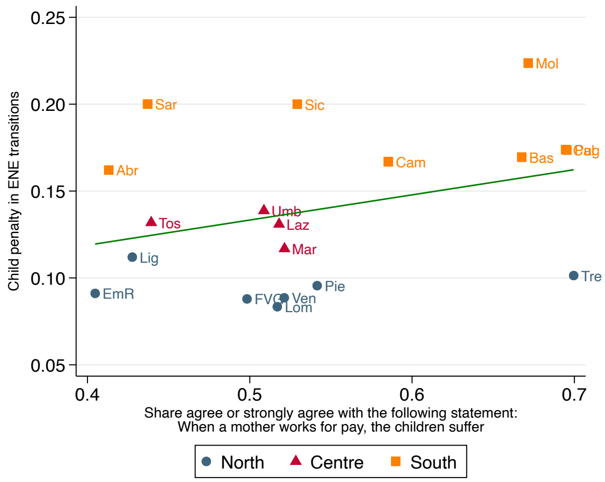

(A) When a mother works for pay the children suffer

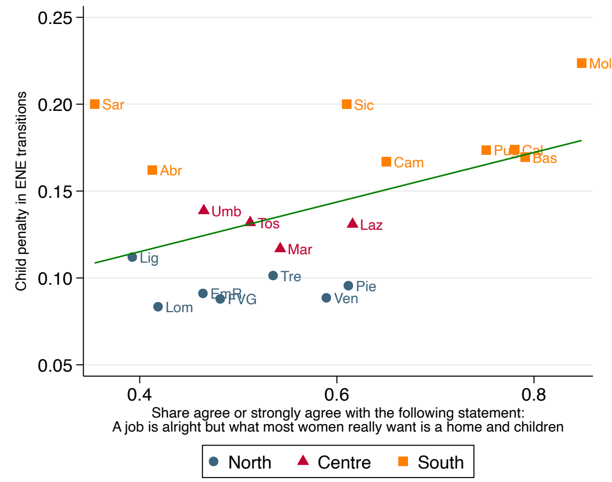

(B) A job is alright but what most women really want is home and children

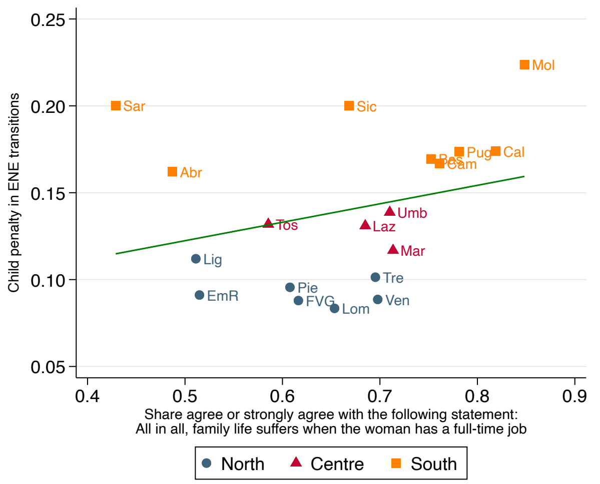

(C) All in all, family life suffers when the woman has a full-time job

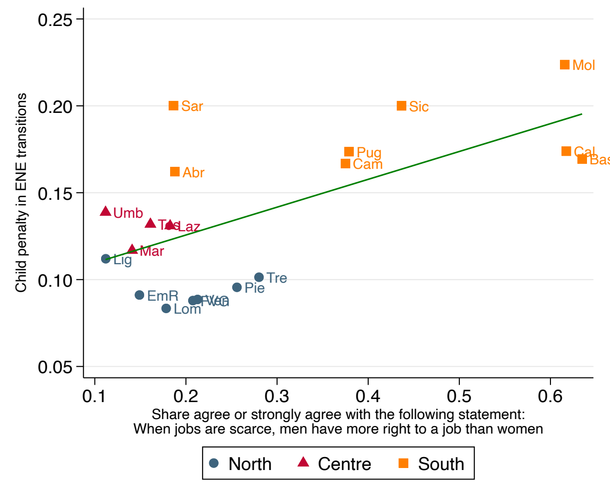

(D) When jobs are scarce, mean have more right to a job than women

Figure A.4: Correlation between child penalties in ENE transitions and gender norms, from different EVS questions, in Italian regions

Notes. The figures report scatter plots of the relationship between the child penalty in annual earnings and the share of respondents agreeing or strongly agreeing with four alternative statements from the European Values Study 2017 (reported on the horizontal axis). The child penalty is measured as the difference in log annual earnings or employment to non-employment transitions 15 years after childbirth, obtained from separate regressions for each Italian region. The share of respondents agreeing or strongly agreeing with the statements is weighted with the calibration weights provided in the survey.