

The Child Penalty Atlas

Henrik Kleven

Princeton University, NBER, and CEPR, USA

Camille Landais

London School of Economics and CEPR, UK

and

Gabriel Leite-Mariante

London School of Economics, UK

First version received March 2024; Editorial decision July 2024; Accepted October 2024 (Eds.)

This paper builds a world atlas of child penalties in employment based on microdata from 134 countries. The estimation of child penalties is based on pseudo-event studies of first child birth using cross-sectional data. The pseudo-event studies are validated against true event studies using panel data for a subset of countries. Most countries display clear and sizable child penalties: men and women follow parallel trends before parenthood, but diverge sharply and persistently after parenthood. While this pattern is pervasive, there is enormous variation in the magnitude of the effects across different regions of the world. The fraction of gender inequality explained by child penalties varies systematically with economic development and proxies for structural transformation. At low levels of development, child penalties represent a minuscule fraction of gender inequality. But as economies develop—incomes rise and the labour market transitions from subsistence agriculture to salaried work in industry and services—child penalties take over as the dominant driver of gender inequality. The relationship between child penalties and development is validated using historical data from current high-income countries, back to the 1700s for some countries. Finally, because parenthood is often tied to marriage, we also investigate the existence of marriage penalties in female employment. In general, women experience both marriage and child penalties, but their relative importance depends on the level of development. The development process is associated with a substitution from marriage penalties to child penalties, with the former gradually converging to zero.

Key words: Child penalty, Motherhood, Gender inequality, Labour market outcomes

JEL codes: J13, J16, J22

1. INTRODUCTION

A recent literature shows that child penalties—the effects of parenthood on women relative to men—account for most of the observed gender inequality in labour market outcomes (Kleven *et al.* 2019a, 2019b). This suggests that the study of child penalties should take centre stage in gender inequality research. However, a key qualification is that we have compelling evidence

from relatively few countries, less than a dozen high-income countries in Europe and North America. The scarcity of evidence limits our understanding of how general child penalties are, how they vary with political and cultural institutions, and how they evolve with economic development. This paper addresses these questions by building a global atlas of child penalties. The atlas is based on nationally representative microdata from 134 countries, covering more than 95% of the global population.

To allow for the estimation of child penalties across countries with varying data availability, we relax the data-demanding approach typically used in the literature. Most studies estimate child penalties through event studies of first child birth using high-quality panel data, but such data are unavailable in most parts of the world. We therefore adopt the pseudo-event study approach developed by Kleven (2023), designed to allow for the estimation of child penalties using cross-sectional data alone. The idea of the approach is to use matching techniques to convert cross-sectional data into panel data, thus facilitating event studies around child birth. The approach can be validated against true event studies in settings where panel data are available. We provide validations for a number of countries, showing that the pseudo-event study approach produces results that are virtually identical to those obtained from a true event study approach.

We focus on child penalties in employment. The event study evidence on employment effects is very sharp, even for countries where one might expect such an approach to be infeasible due to data constraints. Most countries display clear and sizable child penalties: men and women follow parallel trends before parenthood, but diverge sharply and persistently after parenthood. The quantitative magnitudes vary greatly, however, and in some countries child penalties are small or even zero. The effects are precisely estimated for the vast majority of countries, a key advantage of working with cross-sectional data which have larger sample sizes than typical panel datasets. For many countries, we use census data that cover large samples of the population.

We estimate child penalties at three levels of aggregation: continents, countries, and cities. At the continent level, Latin America has the largest child penalties globally (38%), while Africa has the smallest (9%). In Latin America, child penalties are consistently large across countries, with the exception of Caribbean countries such as Cuba or Haiti. In Africa, there is starker heterogeneity across countries: the extremely poor countries of Central Africa feature small or zero penalties, whereas the less-poor countries of Northern and Southern Africa feature large penalties. Even greater heterogeneity is found in Asia and Europe. In Asia, child penalties are very small in China and much of Southeast Asia, but they are extremely large in Japan, South Korea, South Asia, and the Middle East. This continent features the largest employment penalties found anywhere in the world—62% in Bangladesh and 64% in Jordan—along with zero penalties in Vietnam, Laos, and Cambodia. Similarly, Europe displays dramatic heterogeneity, with small child penalties in Scandinavia and large child penalties in English-speaking countries, Central Europe, and much of Southern Europe.

Child penalties are positively associated with urbanization. More urban countries tend to have larger penalties, and urban areas within countries tend to have larger penalties. Our data allow us to estimate city-level child penalties in a number of countries. The findings are quite striking. Even in countries where child penalties are small overall, cities display large penalties. Examples include Beijing, Ho Chi Minh City, and Nairobi, all of which feature sizable child penalties even though the countries as a whole feature small penalties.

How much of gender inequality can be explained by child penalties? The fraction of gender inequality explained by child penalties depends positively on the size of child penalties and the fraction of mothers in the population, while it depends negatively on the raw gender gap. Because child penalties and raw gender gaps are positively correlated across countries, it is not always the case that large penalties translate into a large fraction of gender inequality explained. For example, the modest employment penalties in Scandinavia account for the entire

gender gap in employment there, whereas the large employment penalties in the Middle East account for a smaller fraction of the gender gap in employment in that region. Overall, the fraction of gender inequality explained by child penalties is very large in the Americas, Europe, and Oceania—continents at middle and high incomes. The fraction explained is much smaller in Africa and Asia—continents at low incomes. Using the full variation across countries, we decompose the gender gap in employment into a child-related gender gap and a residual gender gap by GDP per capita. The raw gender gap follows an inverted U-shape as a function of GDP per capita, while child penalties are increasing in GDP per capita except at the very highest income levels. The combination of these two relationships implies that, over the course of economic development, child-related gender inequality gradually takes over as the dominant driver of inequality: it initially constitutes a tiny fraction of inequality, but eventually constitutes all of inequality. At the highest levels of development, eliminating gender inequality is virtually synonymous with eliminating child penalties.¹

We argue that the increasing importance of child penalties over the course of development results from the structural transformation of labour markets. We show that child penalties are negatively associated with the fraction of the workforce in agriculture, positively associated with the fraction of the workforce in industry and services, positively associated with the fraction of the workforce in salaried jobs, and positively associated with urbanization. This points to a story in which the transition from subsistence agriculture towards salaried work in industry and services drives up child penalties. This is a macro-development analogue to the argument that child penalties depend on the flexibility and family friendliness of jobs, previously shown using data from high-income countries such as the U.S. and Denmark (Goldin 2014; Goldin and Katz 2016; Kleven *et al.* 2019b). Arguably, working as a self-employed person in subsistence agriculture is as flexible as it gets. Our findings suggest that the concept of the stay-at-home mom is a relatively modern phenomenon, driven by the structural transformation of labour markets over the course of development.

In many families, the arrival of children follows soon after marriage. This is especially true in low-income societies where parenthood is closely linked to the institution of marriage. It is therefore natural to ask if marriage itself has an impact on gender gaps, and if such marriage penalties confound our estimates of child penalties. For countries where we observe the timing of marriage, there is enough variation in the time between first marriage and first child to separately identify marriage and child penalties. We show that marriage penalties are significant in some countries, but that our child penalty estimates are robust to accounting for them. Marriage penalties are important mainly in low- and middle-income countries, while they are close to zero in high-income countries. We introduce the notion of family penalties that capture the combined impact of marriage and parenthood on gender gaps. These are larger than child penalties, especially in low-income countries, and they increase somewhat less over the course of economic development. In general, accounting for the impact of marriage reinforces our emphasis on family formation as the main driver of gender inequality in labour market outcomes.

In the last part of the paper, we investigate the long-run evolution of child penalties using historical data from current high-income countries. Our data go back to the 1800s for some countries and the 1700s for others. Our empirical framework allows for sharp and statistically

1. Related, Aaronson *et al.* (2021) study the effect of fertility on female labour supply by levels of development. Using cross-country and historical data, they estimate the effect of a second or third child using twin births and sibling sex mix as instrumental variables. They find small effects at low levels of development and large effects at high levels of development. These findings are consistent with ours, although the analyses differ in important respects. The paper by Aaronson *et al.* (2021) does not estimate child penalties or the effect of the first child (parenthood), and it only captures short-run effects.

precise event studies even during the pre-industrialization era. The historical analysis validates the relationship between child penalties and economic development seen in the contemporaneous atlas. In the U.S., for example, child penalties were very small in the late 19th century (around 5pp), increased steeply over the first half of the 20th century (to almost 50pp), declined steadily during the last part of the 20th century (to about 20pp), and then plateaued in recent decades. Consistent with the cross-country patterns, the within-country historical patterns show that child penalties initially account for a tiny fraction of gender inequality—it is instead driven by factors that predate child birth such as education and marriage—but grow in importance over the development process until they become virtually synonymous with gender inequality.

Our paper contributes to a large literature on gender inequality. Extensive reviews have been provided by [Altonji and Blank \(1999\)](#), [Bertrand \(2011\)](#), and [Blau and Kahn \(2017\)](#). Our agenda is most closely related to recent work investigating the impact of parenthood on gender gaps in labour market outcomes, including [Bertrand *et al.* \(2010\)](#), [Angelov *et al.* \(2016\)](#), [Kleven *et al.* \(2019a, 2019b, 2021\)](#), [Cortés and Pan \(2023\)](#), [Andresen and Nix \(2022\)](#), and [Kleven \(2023\)](#). This literature provides evidence on the importance of child penalties and has begun exploring mechanisms, but only for a limited set of high-income countries. We expand the literature to a global scale. Our child penalty atlas elucidates the link between child penalties and gender inequality at different levels of development and under a wide range of institutions. Furthermore, we present novel evidence on marriage penalties and how they interact with child penalties over the course of development. The large scope of our study is made possible by a new pseudo-event study approach that relies only on cross-sectional data ([Kleven 2023](#)), which are widely available and provide more statistical power than typical panel datasets.

Our paper also contributes to the literature on the effects of fertility on female labour supply. This literature has focused on the potential endogeneity of fertility, proposing instruments such as twin births ([Rosenzweig and Wolpin 1980](#); [Bronars and Grogger 1994](#)), sibling sex mix ([Angrist and Evans 1998](#)), infertility shocks ([Aguero and Marks 2008](#)), IVF treatment success ([Lundborg *et al.* 2017](#)), and IUD failure ([Gallen *et al.* 2023](#)).² Most closely related, [Aaronson *et al.* \(2021\)](#) use twins and sibling sex mix instruments to estimate the effect of a second and third child on female labour supply by levels of development. Their findings are consistent with ours. They show that the effect of fertility is small and typically indistinguishable from zero at low levels of development, while the effect is large at higher levels of development. While related, our analyses are based on different methodologies and have different objectives. Our focus is on child penalties and gender gaps, as well as on creating a global atlas. Furthermore, we consider the role of family formation beyond fertility, investigating the interaction between the effects of children and marriage over the course of development.

The paper proceeds as follows. Section 2 describes our data sources. Section 3 describes the empirical methodology. Section 4 presents evidence on child penalties across continents, countries, and cities around the world. Section 5 investigates the existence of marriage penalties and introduces the notion of family penalties, which account for the combined impact of marriage and parenthood on gender gaps. Section 6 analyses the relationship between child penalties (or family penalties) and economic development. Section 7 provides evidence on child penalties over the long run of history. Section 8 concludes.

2. [Kleven *et al.* \(2019b\)](#) validate the event study approach to estimating child penalties against popular IV approaches used in the literature on fertility and labour supply. [Kleven \(2023\)](#) discusses causal identification of child penalties in the pseudo-event study framework.

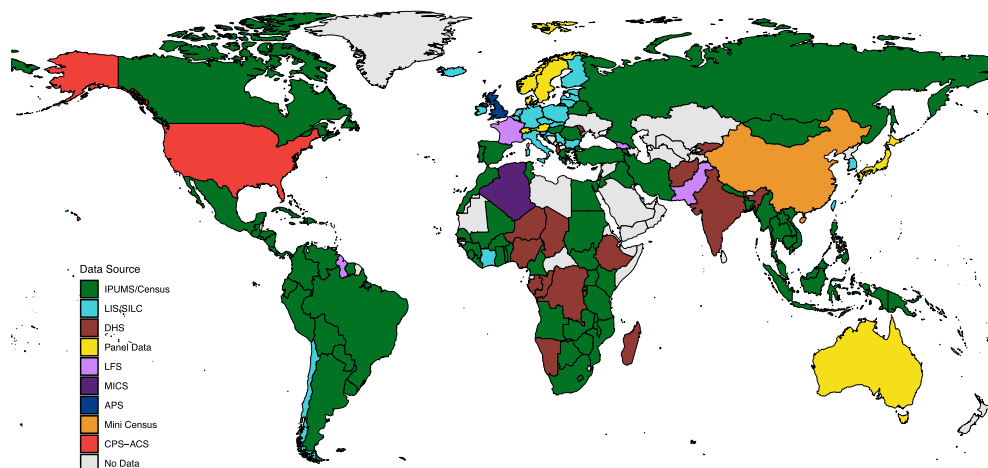


FIGURE 1
Map of data sources

Notes: This map summarizes the data sources used for the 134 countries included in our child penalty atlas. Each colour represents a specific data source. Data for the countries labelled as “Panel Data” refer to either administrative panels or panel surveys. Data for all other countries come from either national censuses or repeated cross-sectional surveys. [Appendix Table A.1](#) provides complete details on the data used for each country.

2. DATA

To construct the child penalty atlas, we harmonize nationally representative microdata from 134 countries. Our sample is representative of ≈ 7.3 billion people or 95.5% of the global population. Figure 1 illustrates the extensive global coverage of our atlas.

For the vast majority of countries (127 out of 134), we use cross-sectional data from a number of different sources, leveraging the power of the pseudo-event study methodology which allows for event studies without panel data. For five of the remaining seven countries, we use administrative panels covering the universe of the population. For two countries, Australia and Japan, we use nationally representative panel surveys. We combine international data sources harmonized across countries and over time with a number of country-specific sources, usually provided by national statistical offices. [Table A.1 in the Appendix](#) details the data source, coverage, employment outcome, and number of observations for every country included in the analysis. We complement this data with country-level information on GDP, population, industrial composition, urbanization, and other variables from the World Bank’s World Development Indicators (WDI).

2.1. Cross-sectional data

For 77 countries, our data come from de-identified samples of national censuses. For 75 of them, we use data harmonized through the IPUMS International Project. For the other two, Tunisia and Angola, we have obtained data directly from their national statistical offices. The samples are stratified to be representative at a subnational level. They typically contain about 10% of the total national population, giving us very large sample sizes as shown in [Table A.1](#). Information on employment is included in the questionnaires of these surveys and refers to a period just before the time of interview, typically the previous week or the previous 15 days. The presence of children is measured using information on own children living in the household. Our census data span somewhat different time periods for different countries such as 1960–2010 for Panama

versus 1970–2015 for Mexico. They cover countries in the Americas, Africa, Asia, and Europe, spanning a wide range of income levels.

For 16 countries, our data come from nationally representative household surveys included in the Luxembourg Income Study (LIS). These surveys have been harmonized across countries and over time. They contain information on employment over a short time horizon (typically at the weekly level) and on annual labour income, allowing us to measure employment at the annual level using an indicator for positive income. They also allow us to measure the presence of children using a household ID combined with information on the relationship between different household members. Most LIS countries are middle- and high-income countries in Europe. There tend to be a large number of yearly survey waves (*e.g.* 35 waves for Germany) spanning different time periods, from as early as 1970 to as recently as 2018.

For another 17 countries, the data come from the Demographic and Health Survey (DHS), conducted by the US Agency for International Development (USAID). These surveys focus on women's health and nutrition, but also contain information on employment—usually at the weekly or biweekly level—as well as direct linkages between parents and children within the household. The data are nationally representative for the population of women of childbearing age and their male spouses. DHS countries are low- and middle-income countries in Africa and Asia. The survey waves span years from as early as 1990 to as late as 2018.

For eight countries, we use data from the EU Survey of Income and Living Conditions (SILC). SILC surveys are nationally representative and contain information on weekly employment as well as direct parent–child linkages within the household. The countries for which we use SILC data are high-income European countries. The survey is conducted yearly, with a varying time span between 2004 and 2020.

For five countries, we use data from Labor Force Surveys (LFS) conducted by national statistical offices. They contain information on weekly employment and household composition, allowing us to measure the presence of children. The countries for which we use LFS data are all low-income countries, and the survey waves span the period between 2010 and 2021.

For the remaining four countries, we use country-specific sources of cross-sectional data. For Algeria, we use four waves of the UNICEF Multiple Indicator Cluster Surveys (MICS), between the years 2012 and 2019. These are mostly focused on the well-being of women and children, but they also contain information on employment and within-household parent–child linkages. For China, we use data from a nationally representative survey of 1% of the population conducted in 2005, often referred to as the “Mini-Census.” It contains information on employment and household composition from which we can infer the presence of children. For the U.K., we use data from the Annual Population Survey (APS), an expanded version of the national LFS covering the period 2012–20. It contains information on weekly employment and household composition. For the U.S., we use pooled data from the Current Population Survey (CPS) and the American Community Survey (ACS), covering the years from 1968 to 2020.³

2.2. Panel data

For five countries, we use data from administrative panels covering the full population. These countries are Austria, Denmark, Norway, Sweden, and Switzerland.⁴ For the remaining two

3. The estimates for the U.S. come from [Kleven \(2023\)](#).

4. The child penalty estimates for Norway are based on [Andresen and Nix \(2022\)](#), who kindly provided data for this project. For Austria, Denmark, and Sweden, we estimate child penalties using the same administrative data sources as in [Kleven *et al.* \(2019a, 2019b\)](#), but the estimates presented here are slightly different as the data have been updated to include more recent years.

countries, we use data from nationally representative panel surveys. For Australia, we use data from the Household, Income and Labour Dynamics in Australia (HILDA) survey, a nationally representative panel covering each year from 2001 to 2019. For Japan, we use data from the Japan Household Panel Survey (JHPS), a nationally representative panel covering each year from 2004 to 2020. Both surveys contain information on weekly employment and within-household parent–child linkages.

We use an additional number of panel surveys to validate our pseudo-event study approach against a true event study approach. These panel datasets include the German Socio-Economic Panel (GSOEP) for Germany, the British Household Panel Survey (BHPS) for the U.K., the Panel Study of Income Dynamics (PSID) for the U.S., the National Longitudinal Survey of Youth (NLSY) also for the U.S., the Encuesta de Protección Social (EPS) for Chile, the Korean Labor and Income Panel Study (KLIPS) for South Korea, and the Russia Longitudinal Monitoring Survey (RLMS) for Russia.

2.3. Harmonization

A large part of our data comes from sources that benefit from careful harmonization efforts (namely DHS, IPUMS, LIS, and SILC). Even so, some additional harmonization is necessary to ensure that our main outcome variable—employment—is as consistently measured as possible across countries. Wherever possible, we focus on employment measures that are contemporaneous to the time of interview and exclude all forms of protected parental leave.

For most countries, our employment measure corresponds to current employment status, defined over a short time period (usually a week) directly preceding the survey. For a few countries, the only available information on employment captures work activity over the previous year or over the current calendar year. The evidence for these countries is therefore based on annual rather than weekly employment. Furthermore, for a number of countries, it is not possible to distinguish being employed *and working* (our preferred outcome) from being employed *and on paid leave*. [Appendix Table A.1](#) shows, for each country in our sample, whether our employment outcome excludes or includes paid maternity leave. The table also provides information on the maximum duration of paid maternity leave in each country. Fortunately, using the subset of countries for which employment can be measured both with and without paid leave, we can verify that this measurement issue has a trivial effect on our results, except for a handful of countries with very long leave durations. This is shown in [Figure A.1 of the Online Appendix](#). It provides a scatterplot of child penalties when excluding paid leave against child penalties when including paid leave for a large set of countries. Light-blue dots depict countries with a leave duration of less than 180 days, while dark-blue dots depict countries with a leave duration of more than 180 days. As can be seen in the graph, the child penalty pairs lie very close to the 45-degree line across almost all countries, except for a small number of countries with very long paid-leave durations.

2.4. Gender gaps in employment around the world

[Figure 2](#) presents a heatmap of raw gender gaps in employment for the 134 countries in our dataset.⁵ The map divides countries into deciles of the gender gap, with darker colours representing larger gaps. There is enormous variation in employment gaps around the world, from

5. Denoting the employment rate by L^g for gender $g = w, m$, the raw gender gap is defined as $(L^m - L^w)/L^m$.

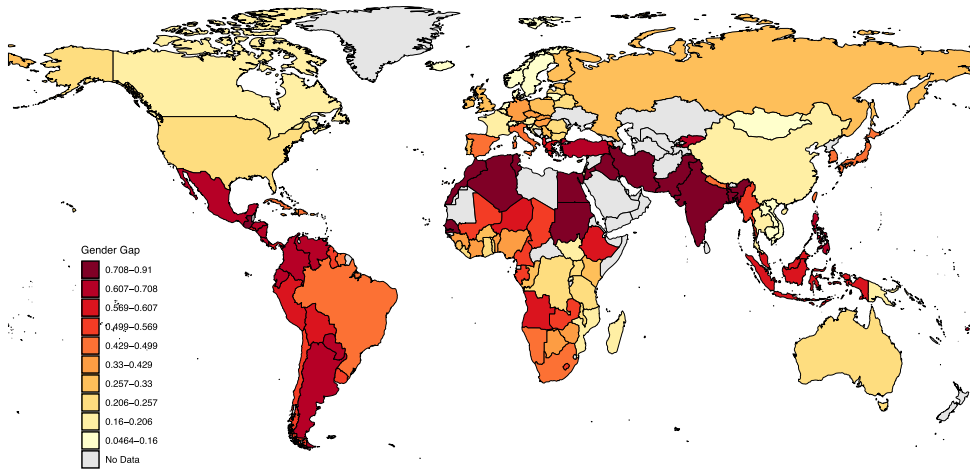


FIGURE 2

Heatmap of gender gaps in employment

Notes: This figure presents a heatmap of raw gender gaps in employment. The gender gap is defined as $(L^m - L^w)/L^m$, where L^g denotes the average employment rate for gender $g = w, m$. The map divides countries into deciles of the gender gap, with darker colours representing larger gaps. The sample includes individuals with and without children between the ages of 15 and 56. For individuals observed with children, we restrict attention to those who had their first child between the ages of 20 and 45.

about 5pp in the most gender-equal countries to about 90pp in the most gender-unequal countries. Gender gaps are relatively small in high-income countries as well as in East Asia and Central Africa. Gender gaps are large in Latin America, North Africa, the Middle East, and South Asia. To illustrate how this variation relates to income levels, [Figure A.2 in the Online Appendix](#) plots gender gaps against log GDP per capita. The relationship between the gender gap and income follows an inverse U-shape. This is akin to a Kuznets Curve, but in the context of gender inequality rather than income inequality.

A key objective of this paper is to investigate if the large differences in gender gaps across countries and by levels of development can be explained by differences in the impact of family formation—marriage and children—on women relative to men. We will show that much of the global variation in gender inequality reflects variation in child penalties (and to a smaller extent marriage penalties as well).

2.5. Historical data

For the historical analysis, we use repeated cross-sectional data from de-identified historical censuses in seven countries, harmonized through the IPUMS International project. For three countries—Denmark, Sweden, and the U.K.—the data contain the full population in all census years. For the remaining four countries—Canada, Iceland, Norway, and the U.S.—the data include the full population in some years and nationally representative subsamples in other years. For the U.S., we create a complete time series between 1880 and 2020 by combining historical census data (1880–1970) with data from CPS and APS (1973–2020).

For consistency, we measure employment in the historical data to resemble the employment measure in contemporary data sources. An individual is considered employed if they report having a gainful occupation at the time of the census. The timing of employment, although not always specified in historical questionnaire texts, is typically a short period directly preceding

the survey (such as the preceding week). An individual is considered non-employed if they report having an occupation that would not be classified as market work today, including, for example, house work. [Table A.2 in the Appendix](#) details the data source, coverage, and number of observations for each country included in the historical analysis.

3. METHODS

3.1. Event study approach

We start with a brief summary of the event study approach to estimating child penalties, which was first developed in contexts where high-quality panel data are available. The estimation is based on sharp changes in the outcomes of women relative to men around the birth of the first child, indexed to occur at event time $t = 0$. As proposed by [Kleven *et al.* \(2019b\)](#), the following specification is run separately for men and women:

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

where Y_{it}^g is the outcome for individual i of gender $g = w, m$ at event time t . On the right-hand side, boldface is used to denote vectors. The first term includes dummies for each event time t , omitting a base year before child birth. The event time coefficients $\alpha_t^g \in \boldsymbol{\alpha}^g$ measure the impact of child birth on gender g in event year t , relative to the base year.⁶ The second and third terms include a full set of age and year dummies to control non-parametrically for lifecycle trends and time trends. In this framework, the impact of children is identified under the assumption that the precise timing of child birth is orthogonal to non-child dynamics in the labour market outcomes of women relative to men. [Kleven *et al.* \(2019b\)](#) provide further discussion of identification in this event-study setup and validate it against IV-approaches.

The focus in this paper is on employment effects. We estimate the linear probability model of equation (1) and scale the estimated event time coefficients by the counterfactual employment level absent children, *i.e.*

$$P_t^g \equiv \frac{\hat{\alpha}_t^g}{\mathbb{E}[\tilde{Y}_{it}^g | t]}, \quad (2)$$

where \tilde{Y}_{it}^g denotes the predicted outcome from (1) when omitting the contribution of the event time coefficients.

Finally, the child penalty is defined as the average effect of having children on women relative to men over a specified event time horizon, namely

$$\text{Child Penalty} \equiv \mathbb{E}[P_t^m - P_t^w | t \geq 0] - \mathbb{E}[P_t^m - P_t^w | t < 0]. \quad (3)$$

The penalty is specified as the average effect across treated (non-negative) event times net of the average effect across untreated (negative) event times.⁷ A positive child penalty implies that parenthood increases the gender gap.

6. Throughout the paper, the omitted event time dummy is chosen as $t = -2$, the year before pregnancy. The choice of base year hardly impacts the results as there is virtually no pre-trend in the data.

7. Due to having omitted a base year before child birth, the post-birth difference $\mathbb{E}[P_t^m - P_t^w | t \geq 0]$ is already net of any pre-birth difference in the employment rate of men and women. The second term in (3) is therefore not strictly necessary. In practice, its inclusion is inconsequential for our estimated penalties.

3.2. Pseudo-event study approach

3.2.1. Description. The event study approach described above is straightforward to implement given access to high-quality panel data. Unfortunately, such data are not available outside a handful of high-income countries, limiting our understanding of how child penalties vary across geography. But a pseudo-panel approach can be developed to replicate, in cross-sectional data, event study designs like the one described above. Following [Kleven \(2023\)](#), the idea is to use matching techniques to convert cross-sectional data into a pseudo-panel of men and women at different event times, relative to the birth of the first child. This unlocks hundreds of cross-sectional data sources around the globe for the study of child penalties.

Converting cross-sectional data into a pseudo-panel requires allocating individuals to specific event times relative to first child birth. For individuals observed with children, this is easily done: the age of their oldest child determines their position in positive event time, $t \geq 0$. The main challenge pertains to individuals observed without children: we do not observe if and when they will have children and therefore do not know their place in negative event time, $t < 0$. The key idea of our approach is to create, for individuals observed with children, surrogate observations at negative event times through matching on observable characteristics. Specifically, consider parent i observed at event time 0 in calendar year y with age a and characteristics X_i . This parent is matched to a childless individual j observed in year $y - n$ with age $a - n$ and the same characteristics $X_j = X_i$. This gives a surrogate observation for $t = -n$.⁸ By matching each parent at event time 0 to childless individuals for $n = 1, \dots, 5$, a pseudo-panel with 5 years of pre-child data is created.⁹ This procedure implies that age and calendar time are changing along the event time dimension, exactly as in the panel data approach of [Kleven et al. \(2019b\)](#).¹⁰ The effect of lifecycle and time trends are absorbed through the inclusion of age and year dummies in equation (1). The approach is related to the pseudo-panel (or synthetic-cohort) approach originally developed by [Deaton \(1985\)](#).¹¹

Apart from year, age, and gender, the variables we include in our matching algorithm are four categories of education (less than primary education, primary education, secondary education, and post-secondary education), four categories of marital status (single or never married, married or in a union, separated or divorced, and widowed), and a dummy for living in a rural versus urban area. This represents a relatively light set of matching variables, consistently defined and available across all our data sources.

8. A parent will have multiple possible matches whenever there is more than one childless individual in the specified cell of observables (year, age, and other demographics). We match the parent to all childless individuals in the given cell, each of them weighted by $1/k$ where k is the cell size.

9. To clarify, there is a slight difference in the matching protocol depending on the employment measure available in the data source used. The protocol described above is used when considering weekly employment (obtained from a question about work activities last week). The protocol needs to be modified when considering annual employment/earnings (based on a question about earnings last year). To account for the retrospective nature of the annual outcomes, the matching of parents at event time 0 (observed in year y with age a) and non-parents (observed in year $y - n$ with age $a - n$) is done for $n = 0, \dots, 4$ to obtain surrogate observations for $t = -n - 1$. For the same reason, annual outcomes at event times $t = 0, \dots, T$ are obtained from parents observed at event times $t = 1, \dots, T + 1$.

10. For many countries, we do not have repeated cross-sections at an annual frequency. In this case, we create surrogate observations using individuals with age $a - n$ and characteristics $X_j = X_i$, but who are observed in the same cross-section of the data. In practice, this assumes that any calendar-year shock on employment is not cohort-specific.

11. Similar to [Deaton \(1985\)](#), the matching procedure used here ensures that the estimation sample consists of fixed cohorts over time. The procedure is richer than a synthetic-cohort approach by holding other dimensions fixed too. More importantly, the idea to use the pseudo-panel to conduct event studies is new.

3.2.2. Validation. We validate the pseudo-event study approach in two ways. First, we investigate how the choice of matching variables affects the results. It is particularly informative to consider the effect of first child birth on men. As shown in previous papers, child birth is a non-event for men. Therefore, if the pseudo-event study delivers a positive jump in the labour market outcomes of men around $t = 0$, this reflects bias from positive selection. [Figures A.3 and A.4 in the Appendix](#) present pseudo-event studies under different matching specifications for a selected set of countries. Each row considers a given country, while each column considers a given matching specification. As we move from the left to the right, an increasingly rich set of matching variables is included: matching only on year, age, and gender (first column), adding education (second column), adding marital status (third column), and adding rural/urban status (fourth column, corresponding to our baseline specification). Three main insights emerge. First, we see that, in some countries, the more parsimonious matching specifications are associated with some selection bias, evidenced by the positive jumps in male employment at event time $t = 0$. Second, this selection bias is generally small and, in many countries, entirely absent. Finally, for countries where selection bias in parsimonious matching specifications seems to be an issue, it is eliminated by the inclusion of education and marital status variables in the matching algorithm. As a result, we see no indication of selection bias in our baseline specification, shown in the last column on the right.

Second, a more direct way of validating the pseudo-event study approach is by comparing it to a true event study approach in countries where panel data are available. Specifically, this can be done by transforming panel data into a matched pseudo-panel, ignoring the panel information on negative event times. Comparing the pseudo-event study estimates obtained from this data transformation with the actual event study estimates obtained from the panel data allows for validating the approach holding the sample fixed. [Figure 3](#) provides such a validation exercise. [Figure 3A–E](#) consider five selected countries (Austria, Chile, Denmark, Japan, and the U.K.), comparing pseudo-event studies of first child birth (solid lines) to actual event studies of first child birth (dashed lines). The five countries differ in their institutional context, their development levels, and in the size and dynamics of their child penalties. The figure shows that, for all five countries, the pseudo-panel estimates and panel estimates align remarkably well. The two methods give virtually identical estimates of child penalties. [Figure 3F](#) confirms this finding in a broader set of countries. This panel provides a scatterplot of pseudo-panel estimates against panel estimates for all twelve countries where we have access to panel data. All countries are located very close to the 45-degree line. The bivariate relationship between pseudo-panel and panel estimates displays remarkably little variance, with an R -squared of 0.96.

4. THE CHILD PENALTY ATLAS

4.1. *A view from above: continent-level child penalties*

[Figure 4](#) presents evidence on child penalties by continent. Each panel shows an event study around first child birth, indexed as year $t = 0$, for men and women separately. The series depict the percentage impact of child birth at each event time, namely P_t^m and P_t^w estimated from equations (1) to (2). The impacts are measured relative to event time $t = -2$, the year before pregnancy. Each panel also displays the average child penalty over event times 0–10 as defined in equation (3). The outcome variable is employment, the measurement of which was described in [Section 2](#).

The event studies look qualitatively similar across continents. They feature parallel trends between men and women before child birth and sharp divergence immediately after child birth. Parenthood is a non-event for men, but it leads to an immediate and persistent drop in

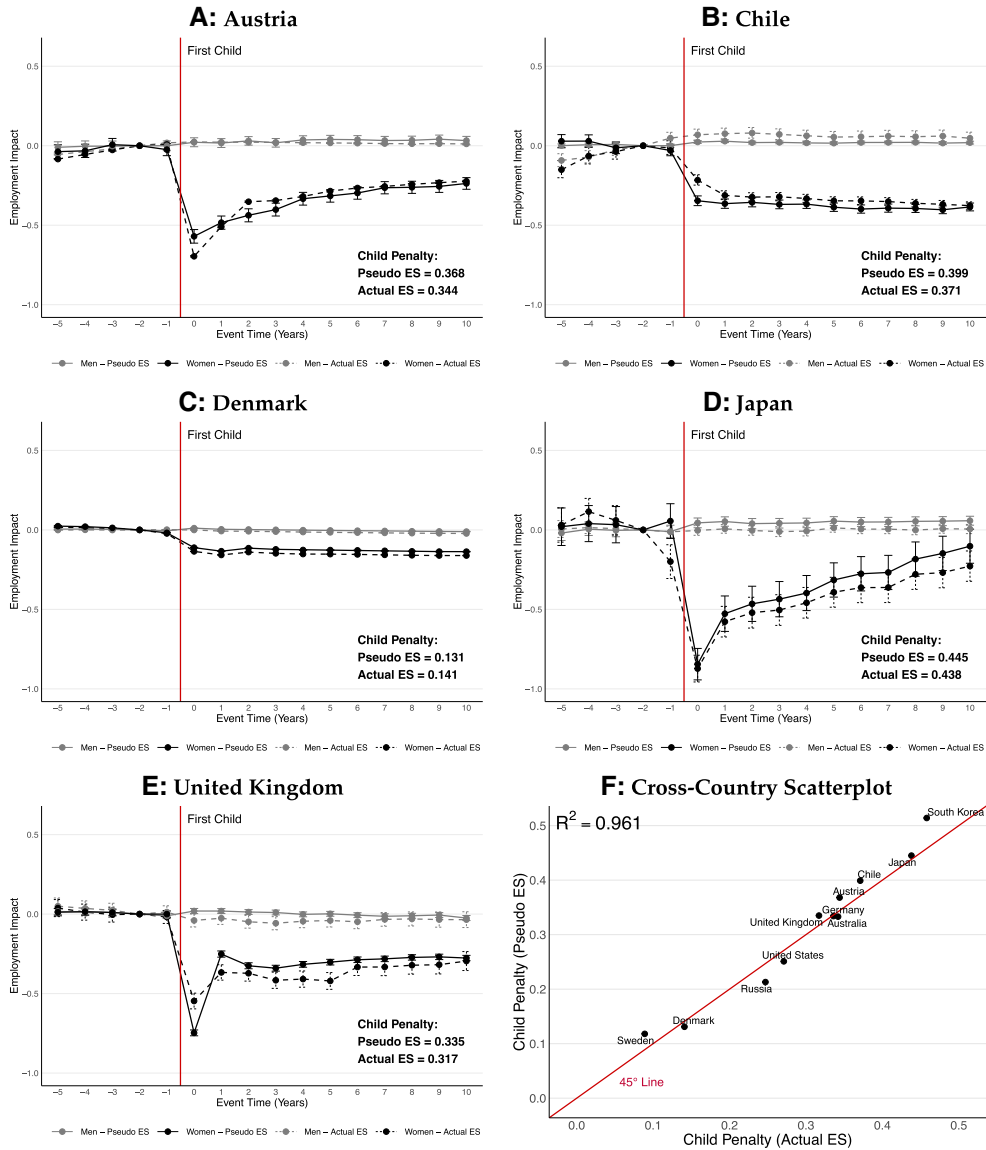


FIGURE 3
 Validation of pseudo-event study approach

Notes: This figure validates the pseudo-event study approach against an actual event study approach using panel data. Figure 3A–E consider selected countries, comparing pseudo-event studies of first child birth (solid lines) to actual event studies of first child birth (dashed lines). Each series show, for men and women, respectively, the percentage impact of child birth on the employment rate at each event time t , i.e. P_t^m and P_t^w defined in equation (2). Each panel also displays average child penalties over event times 0–10 as defined in equation (3). Age at first birth is restricted to 20–45, and the error bars depict 95% confidence intervals based on robust standard errors. Figure 3F presents a scatterplot of child penalties estimated from the two approaches for all countries where the validation is feasible.

employment for women. While qualitatively similar, the quantitative magnitudes vary substantially across continents. The largest child penalties are found in Latin America: the average impact of parenthood on the employment rate of women relative to men equals 38% in this region. Child penalties are more modest in Africa (9%) and Asia (18%), and at intermediate levels in Europe, North America, and Oceania (25–30%). In the next section, we investigate the variation across countries within each continent.

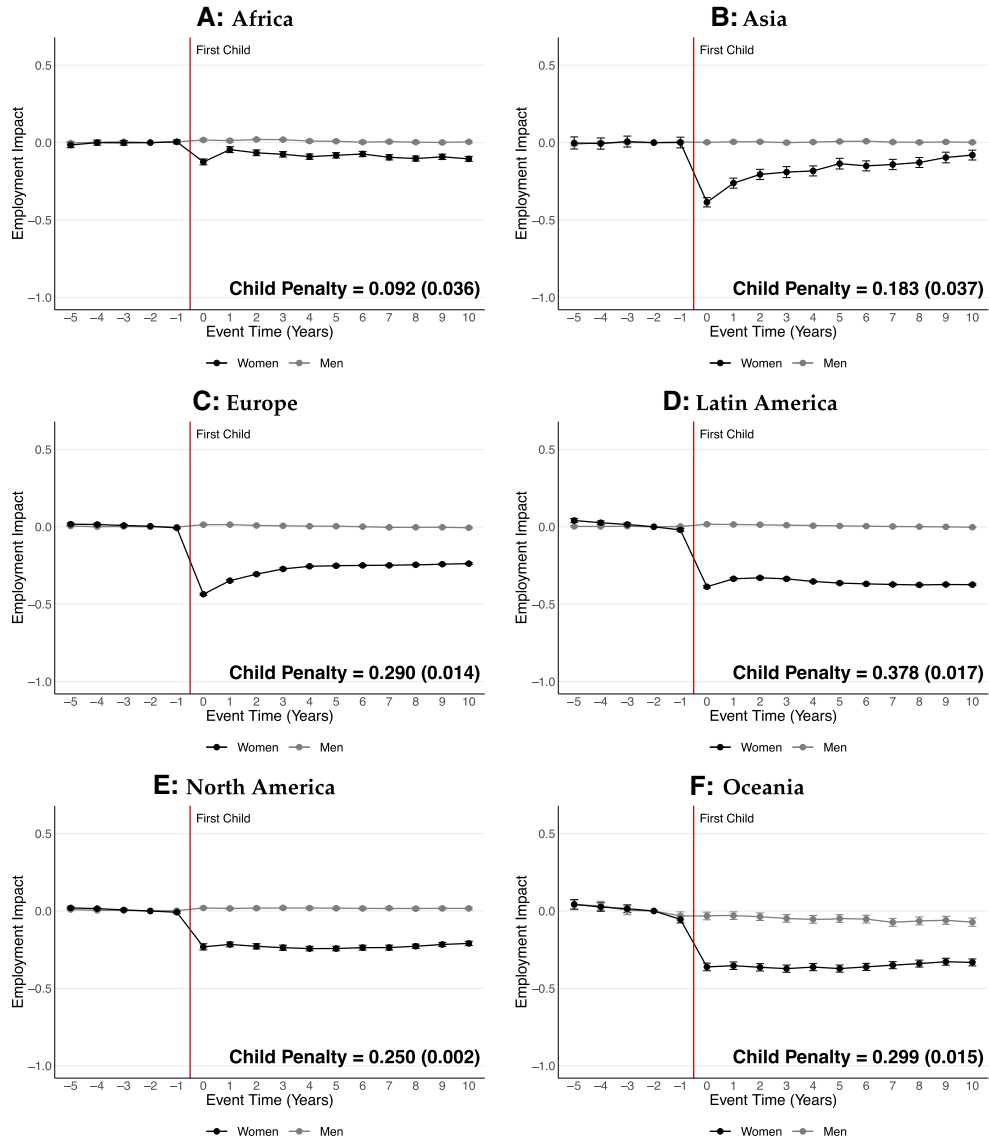


FIGURE 4
Child penalties by continent

Notes: This figure presents event studies of first child birth by continent. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of child birth at each event time t , *i.e.* P_t^m and P_t^w defined in equation (2). These coefficients are population-weighted averages of the coefficients estimated for each country in the given continent (among countries included in our data). Each panel also displays the average child penalty over event times 0–10 as defined in equation (3). Age at first birth is restricted to 20–45. The error bars depict 95% confidence intervals based on robust standard errors.

4.2. A tale of great variation: country-level child penalties

Figure A.5 in the Appendix provides event studies of first child birth for all 134 countries in our analysis. Here, we consider a subset of these countries, selected to highlight key patterns in the data. The discussion is organized by continent.

4.2.1. Europe. Figure 5 shows results for eight countries located in different parts of Europe. Each country features clear and persistent child penalties on women, but the magnitudes vary greatly. The child penalty is smallest in Denmark (14%) and largest in the Czech Republic (50%). The sharp difference between these two countries reflects a general difference between Nordic and Central European countries: child penalties are equally modest in the rest of Scandinavia, while they are very large in other central European countries such as Austria, Germany, Hungary, Slovakia, and Switzerland. Child penalties tend to be large in Southern Europe as well, but with some interesting heterogeneity. As shown in the figure, there is a sharp difference between the neighbouring countries of Spain and Portugal. Child penalties are more than twice as large in Spain, and raw gender gaps are also much larger in Spain.

Besides differences in average magnitudes, Figure 5 also highlights interesting differences in the particular shape of the event study series. All countries see a sharp drop for women just after motherhood, but the persistence of the initial drop varies greatly across countries. In some countries, the effect is fully persistent over time, while in other countries the career trajectory of women shows partial recovery. Consider again the contrast between Denmark and the Czech Republic. The modest 14% penalty in Denmark represents a virtually constant effect over time, whereas the large 50% penalty in Czech Republic is strongly declining over time (close to 100% during infant child care and about 20% after ten years). The patterns may represent a combination of factors, including the generosity and time profile of family policies such as maternity leave and child care provision (see Kleven *et al.* 2024) as well as the timing and number of subsequent child births.

4.2.2. Asia. Figure 6 shows results for eight Asian countries. In Asia, the variation in child penalties is even more dramatic than in Europe. For example, Bangladesh features a child penalty of 62%, whereas Vietnam features a child penalty of just 1%. The enormous differences across countries do not appear to be driven by imperfections in the empirical design or by statistical noise: all of the event studies look compelling and are precisely estimated. Looking across the entire Asian continent (see Appendix Figure A.5), child penalties are small in China and in most of Southeast Asia, especially countries from the former French Indochina (Vietnam, Laos, and Cambodia). Conversely, child penalties are extremely large in Japan, South Korea, South Asia, and the Middle East. Large differences are often observed between neighbouring countries such as China versus South Korea (4% versus 49%), Thailand versus Malaysia (5% versus 45%), and Myanmar versus Bangladesh (19% versus 62%). Of course, while these countries are geographically close, they vary greatly in terms of economic, political, and cultural/religious institutions. A key objective of building a global atlas of child penalties is precisely to be able to study the role of such institutions for child penalties and their impact on gender inequality.

4.2.3. Latin America. Figure 7 considers countries in Latin America. In this region, there is much less heterogeneity in child penalties. They are consistently large, typically between 35% and 50%, except for a handful of countries in the Caribbean. What is more, the dynamics of the event study series are strikingly similar across countries: a sharp drop for women at event time zero (without any drop for men) and almost perfect persistence thereafter. There are hardly any dynamics in the data, apart from what happens to women immediately after child birth. Among the countries in Figure 7, this type of pattern is seen in Argentina, Brazil, Chile, Colombia, Costa Rica, and Mexico. A similar pattern is seen in a number of other Latin American countries shown in the Appendix. Only Caribbean countries deviate: their child penalties are smaller in magnitude and somewhat less persistent over time. This is exemplified by the event studies for Cuba and Haiti in the figure.

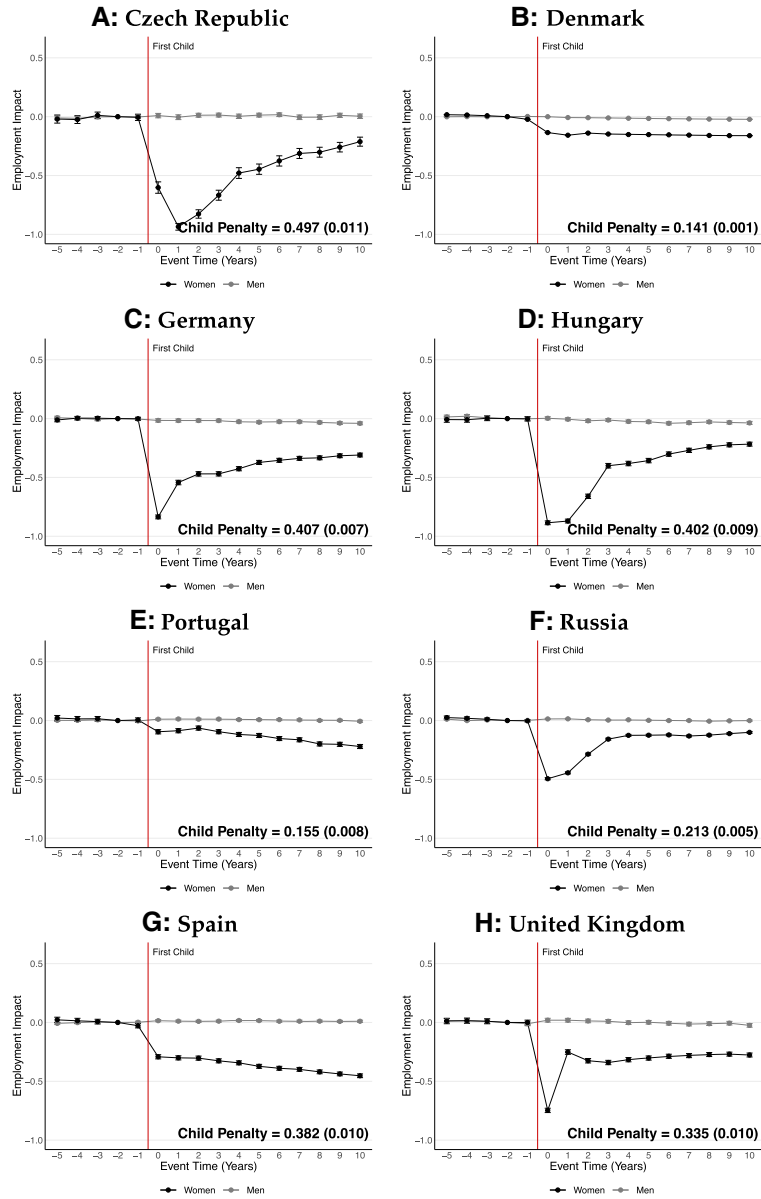


FIGURE 5
Child penalties in European countries

Notes: This figure presents event studies of first child birth for a selection of European countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of child birth at each event time t , i.e. P_t^m and P_t^w defined in equation (2). Each panel also displays the average child penalty over event times 0–10 as defined in equation (3). Age at first birth is restricted to 20–45. The data used for each country are described in Table A.1. The error bars depict 95% confidence intervals based on robust standard errors.

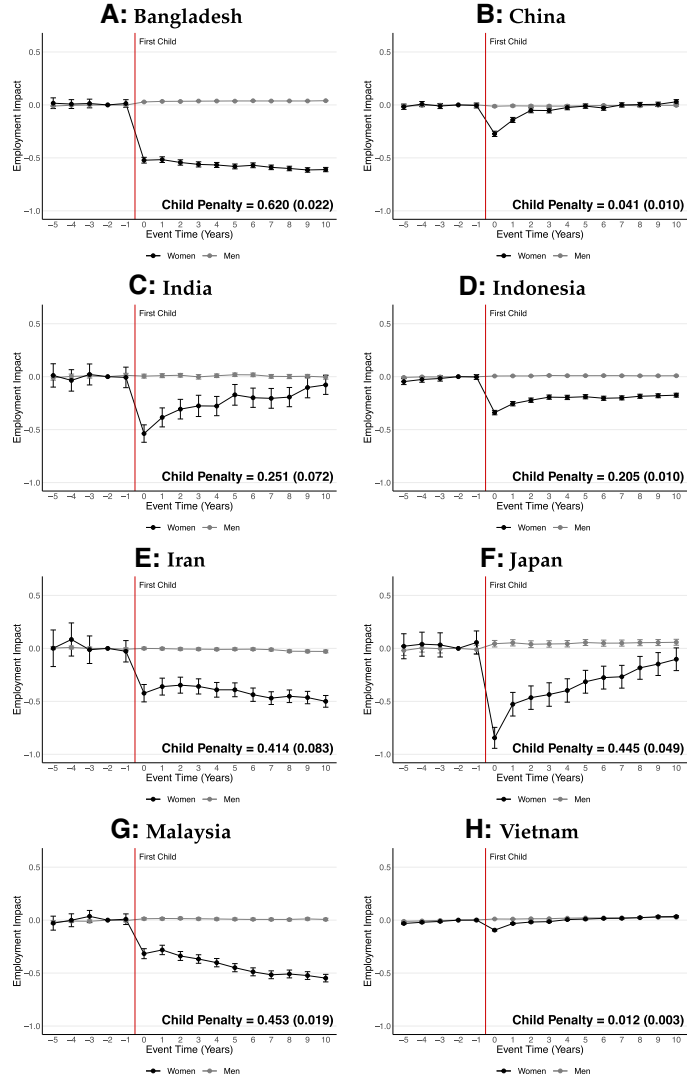


FIGURE 6
Child penalties in Asian countries

Notes: This figure presents event studies of first child birth for a selection of Asian countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of child birth at each event time t , i.e. P_t^m and P_t^w defined in equation (2). Each panel also displays the average child penalty over event times 0–10 as defined in equation (3). Age at first birth is restricted to 20–45. The data used for each country are described in Table A.1. The error bars depict 95% confidence intervals based on robust standard errors.

4.2.4. Africa. In Figure 8, we turn to African countries. Our data and pseudo-panel approach allow us to produce compelling event studies across a large number of African countries, a region where data constraints often prevent evidence of this type. Our results show that Africa is a tale of two regions. Child penalties are small or zero in the least developed countries of Central Africa, while they are very large in the more developed countries of the North and South. In the figure, child penalties are precisely estimated zeros in Mozambique and Tanzania, and they are also small in Cameroon, Ghana, Kenya, and Zimbabwe. On the other hand, child

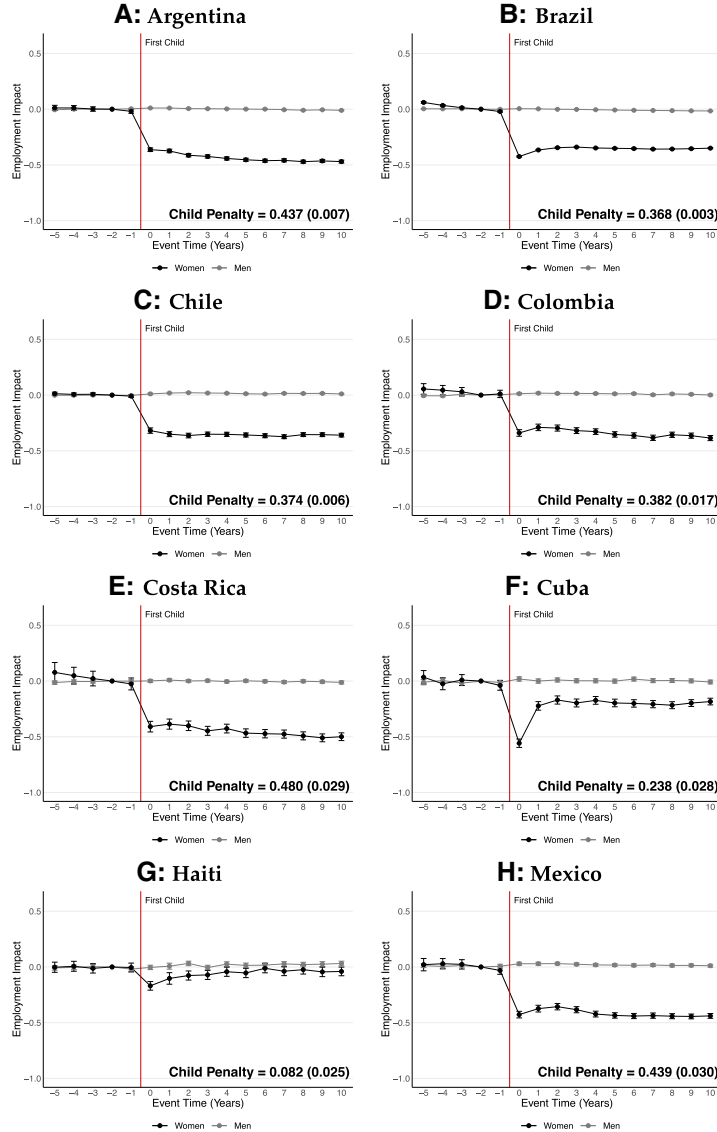


FIGURE 7
Child penalties in Latin American countries

Notes: This figure presents event studies of first child birth for a selection of Latin American countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of child birth at each event time t , *i.e.* P_t^m and P_t^w defined in equation (2). Each panel also displays the average child penalty over event times 0–10 as defined in equation (3). Age at first birth is restricted to 20–45. The data used for each country are described in Table A.1. The error bars depict 95% confidence intervals based on robust standard errors.

penalties are large in Morocco (41%) and in South Africa (28%). Appendix Figure A.5 shows that this regional divide is a general feature of the data.

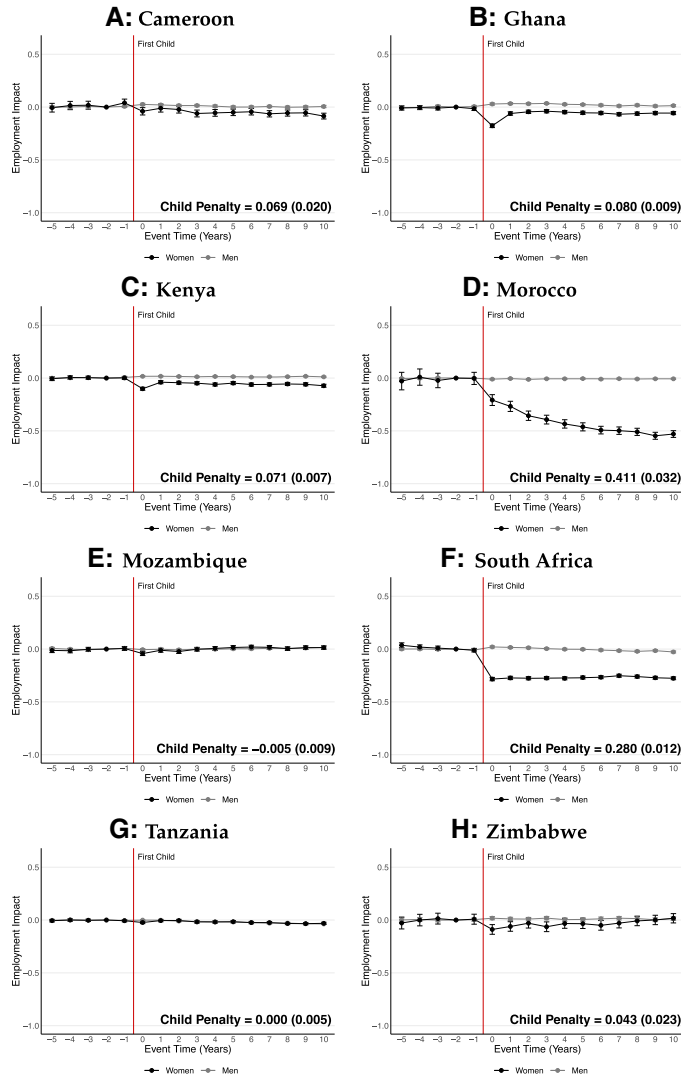


FIGURE 8
Child penalties in African countries

Notes: This figure presents event studies of first child birth for a selection of African countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of child birth at each event time t , i.e. P_t^m and P_t^w defined in equation (2). Each panel also displays the average child penalty over event times 0–10 as defined in equation (3). Age at first birth is restricted to 20–45. The data used for each country are described in Table A.1. The error bars depict 95% confidence intervals based on robust standard errors.

4.3. Urban areas see larger effects: city-level child penalties

Our data and empirical approach also allow us to estimate child penalties at the city level. Figure A.6 in the Appendix provides event studies of first child birth in selected cities, comparing them to event studies for countries as a whole. It is striking to see how sharp and precise the event studies are, even at the city level. Cities tend to have larger child penalties than countries overall. In the figure, Beijing has a child penalty of 12% (compared to 4% in China as a

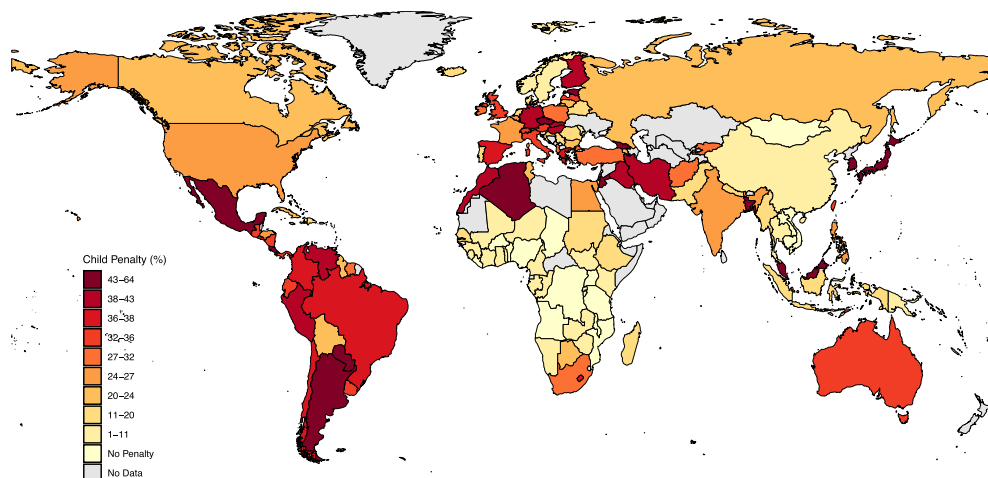


FIGURE 9
Heatmap of child penalties

Notes: This figure presents a heatmap of child penalties in employment, defined in equation (3). The map divides countries into deciles of the child penalty, with darker colours representing larger penalties. Countries with a penalty less than or equal to zero fall in the “No Penalty” category. Countries for which we do not have data fall in the “No Data” category. The event studies of child birth used to estimate child penalties for all 134 countries in our dataset are provided in [Appendix Figure A.5](#).

whole), Ho Chi Minh City has a penalty of 25% (compared to 1% in Vietnam), Istanbul has a child penalty of 61% (compared to 29% in Turkey), London has a child penalty of 43% (compared to 34% in the U.K.), and Nairobi has a child penalty of 22% (compared to 7% in Kenya). The presence of excess penalties in large cities is a feature of most countries in the world, not only those shown in the figure. More broadly, urban areas see larger child penalties than rural areas. [Kleven \(2023\)](#) shows that, in the U.S., child penalties are positively associated with urbanization, arguing that this is an effect of differences in urban versus rural job structure (such as the flexibility and family friendliness of jobs). We come back to the role of urbanization when discussing economic development and structural transformation in [Section 6](#).

4.4. Putting it all together: world maps

Compiling the estimates from all 134 countries, [Figure 9](#) presents a heatmap of child penalties. In this map, countries are divided into deciles of the child penalty, with darker colours representing larger penalties.¹² The map visualizes the main insights discussed above such as the large penalties in Latin America, the small penalties in Central Africa, and the great heterogeneity in Europe and Asia. The map also visualizes the magnitude of child penalties in a number of major countries not included in the previous figures such as the U.S., Canada, France, and Australia. As mentioned in the beginning of the paper, the countries currently in the atlas represent more than 7 billion people, or 95.5% of the global population. The missing countries in the map are located mostly in Central Asia, the Middle East, and Northern Africa, along with New Zealand and Greenland.¹³ Our ultimate goal is to build out the atlas to cover the remaining 4.5% of the global population or at least most of it.

12. Countries for which we do not have data are depicted in grey.

13. Although Greenland is part of the Kingdom of Denmark, the Danish administrative data do not include information on Greenlanders.

How much of gender inequality can be explained by child penalties? There is a straightforward translation between the level of child penalties and the fraction of gender inequality explained. We have

$$\text{Fraction Explained} = \frac{\text{Child Penalty} \times \text{Fraction Mothers}}{\text{Raw Gender Gap}}. \quad (4)$$

In this expression, the child penalty and gender gap are both unscaled, *i.e.* measured in absolute rather than percentage terms. Using scaled penalties and gender gaps would not be strictly correct as they have different denominators—counterfactual employment rates absent children (child penalty) versus actual employment rate for men (gender gap)—although in practice this makes little difference as these denominators tend to be similar. We consider the difference between scaled and unscaled child penalties in the next section.

Figure 10 provides a heatmap of the fraction of the gender gap in employment explained by child penalties. Countries are again divided into deciles, with darker colours representing larger fractions explained. The map is quite striking. The strong within-continent heterogeneity in child penalties largely disappears when considering the fraction of gender gaps explained by child penalties. This reflects that, within continents, variation in child penalties is strongly positively associated with variation in gender gaps such that the fraction explained in equation (4) is fairly uniform. The relative importance of child penalties is strongest in North America, Europe, and Australia, followed by Latin America. Although Latin America tops the world in terms of the magnitude of child penalties, the relative importance of child penalties is weaker there than in North America and Europe. The reason is that raw gender gaps in employment are proportionally larger in Latin America. In Africa and Asia, child penalties explain a much smaller share of gender gaps. Even in the large-penalty countries of the Middle East (such as Iran) and South Asia (such as Bangladesh), child penalties explain a relatively modest share of overall gender inequality. As with Latin America, the reason is that raw gender inequality is proportionally higher in those places. Conversely, in a number of low-penalty countries—take Scandinavia as an example—child penalties explain virtually all of the observed inequality.

4.5. *Scaled versus unscaled child penalties*

Through most of the analysis, we focus on scaled child penalties as defined in equations (2)–(3). They capture the impact of child birth on employment in per cent of counterfactual employment absent children. It is worth asking if the results would be different if we consider unscaled child penalties, *i.e.* the impact of child birth on the employment rate in percentage points. Figure 11 plots scaled child penalties against unscaled child penalties for all 134 countries in our sample. Three takeaways are worth highlighting. First, the child penalty pairs lie above the 45-degree line as scaled penalties, by construction, are larger than unscaled penalties.¹⁴ Second, there is a strong positive association between scaled and unscaled child penalties. This implies that most of the cross-country variation in scaled penalties documented above is driven by variation in unscaled penalties rather than by variation in counterfactual employment rates. Third, there is some dispersion in scaled child penalties conditional on unscaled child penalties, with a few countries in the Middle East (*e.g.* Jordan) and South Asia (*e.g.* Bangladesh) standing out. In these countries, the large magnitude of scaled penalties stem from a combination of large unscaled penalties and

14. To be clear, this requires the unscaled penalty to be non-negative. In the rare cases where the unscaled penalty is negative, the child penalty pairs lie below the 45-degree line.

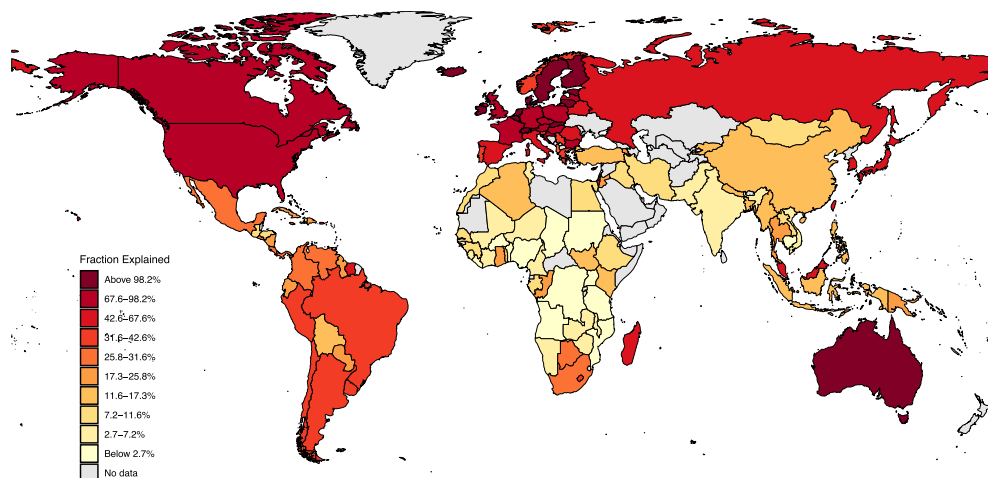


FIGURE 10

Heatmap of fraction of gender gap explained by child penalties

Notes: This figure presents a heatmap of the fraction of the gender gap in employment explained by child penalties. The fraction explained is calculated using equation (4). The map divides countries into deciles of the fraction explained, with darker colours representing larger fractions explained.

exceptionally low female employment rates even absent children. In the next section, we dig deeper into the drivers of low female employment rates that preexist the arrival of children.

5. MARRIAGE PENALTIES VERSUS CHILD PENALTIES

Up to this point, we have focused on the impact of parenthood on gender gaps. However, the arrival of children is linked to the institution of marriage, especially in low-income countries where child birth almost never precedes marriage and often follows soon after marriage. To illustrate this point, [Figure A.7 in the Appendix](#) shows the distribution of time between first child and first marriage, pooling countries in the IPUMS census data.¹⁵ These are all lower- and middle-income countries. The figure shows that very few families have children before marriage and that a majority of families have children within a few years of marriage: 63% have their first child within 5 years of marriage, and 42% have their first child within 2 years of marriage.

The link between marriage and parenthood raises two important questions. First, does the event of marriage impose labour market penalties on women over and above those imposed by the event of parenthood? Given marriage is associated with *anticipated* parenthood, we might expect some of the effect of parenthood to materialize already at the time of marriage. To put it differently, the event study methodology used to estimate child penalties only captures effects realized after parenthood, not anticipation effects realized before parenthood. Studying marriage penalties represents a way to explore such anticipation effects. Second, given the closeness of marriage and parenthood in many families, does marriage confound the estimated impacts of parenthood presented above? It is important to note that our pseudo-event study approach is not directly confounded by marriage effects, because we match on marital status as described in Section 3.2. But an indirect threat to identification remains: by matching parents (virtually all of them married) to non-parents with the same marital status, the surrogate observations at

15. The figure is based on the subset of IPUMS countries for which the timing of marriage is observed.




FIGURE 11

Scaled versus unscaled child penalties

Notes: This figure plots scaled child penalties against unscaled child penalties across all 134 countries in our sample. Unscaled child penalties, on the x -axis, represent average effects on the employment rate of women relative to men *in absolute terms* (pp). Scaled child penalties, on the y -axis, represent average effects on the employment rate of women relative to men *in per cent of the counterfactual employment rate absent children*, as defined in equations (2)–(3).

negative event times may be selected in a way that creates bias. In this section, we investigate these questions in detail.

We begin by presenting case studies of two large middle-income countries, Brazil and China. Figure 12 shows event studies of first marriage based on a pseudo-panel approach similar to the one used for the event studies of first child. The only modification relates to the matching algorithm: while child penalties were estimated by matching parents at $t = 0$ to non-parents with the same marital status (typically married), the marriage penalties are estimated by matching newly-weds at $t = 0$ to unmarried people with the same parental status (typically childless). All other aspects of the matching algorithm and specification are the same. The left panels show marriage penalties in the full sample of people for whom we observe the timing of first marriage. These graphs suggest large and persistent effects of marriage on gender gaps. In fact, the marriage penalties for Brazil and China are larger than the child penalties for these two countries. However, as shown by the dashed lines, families in Brazil and China start having children soon after marriage. Therefore, we expect that the marriage penalties in these graphs partly capture the effects of having children. This is confirmed in the right panels of the figure. Here we exclude families who have their first child within two years of marriage, thus considering a subsample where marriage and parenthood are clearly separated in time. This significantly reduces the marriage penalty between event times 0 and 2, but sizable, short-run penalties remain even

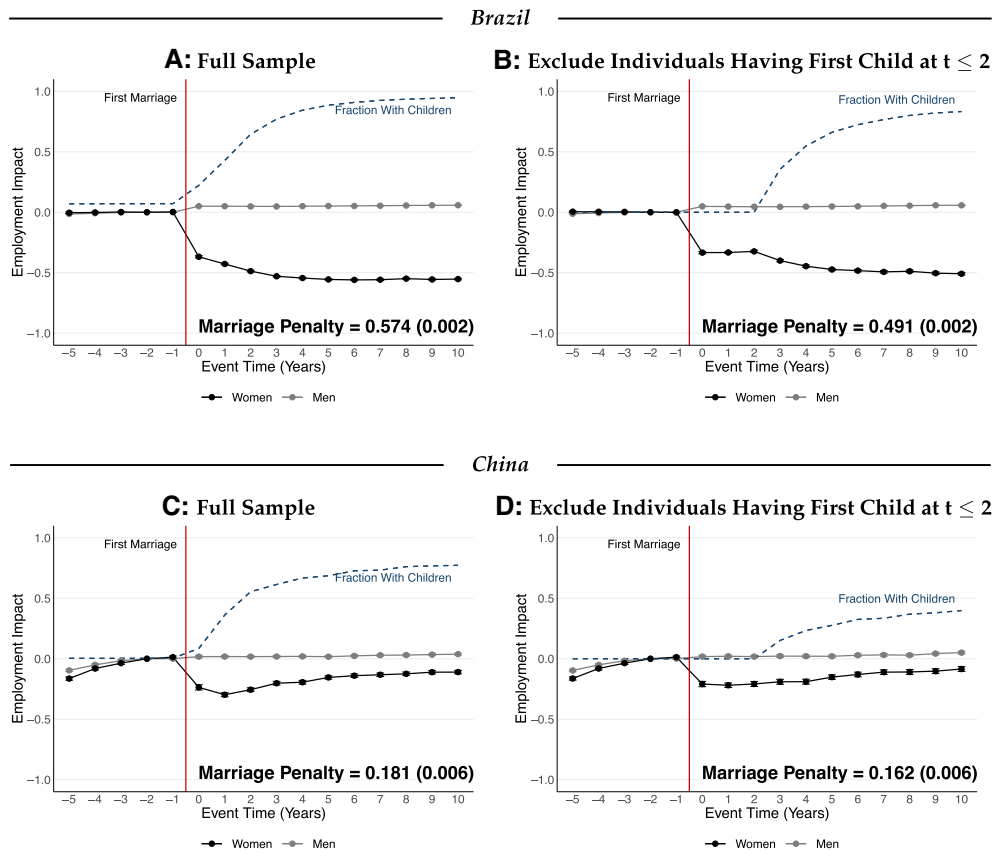


FIGURE 12

Marriage penalties. A case study of Brazil and China

Notes: This figure presents event studies of first marriage in Brazil and China. Each panel shows an event study of the impact of marriage, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of marriage at each event time t , *i.e.* P_t^m and P_t^w defined in equation (2). Each panel also displays the average marriage penalty over event times 0–10 as defined in equation (3). The fraction of individuals with children at each marriage event time is shown by the dashed blue line. Two samples are considered: the full sample of families for whom marriage timing is observed (left panels), and the subsample of families who have their first child at least three years after marriage (right panels). Age at first marriage is restricted to 20–45. The error bars (hardly visible) depict 95% confidence intervals based on robust standard errors.

in this sample. The upshot of Figure 12 is that there are clear marriage penalties in Brazil and China, but a naive approach to estimating them will be confounded by child penalties.

Appendix Figure A.8 presents a similar analysis, but for child penalties rather than marriage penalties. Still focusing on Brazil and China, the figure shows event studies of first child in two samples: the full sample of individuals for whom marriage timing is observed (left panels) and a subsample of individuals who got married at least three years before parenthood. We find that the estimated child penalties are less sensitive to the timing of marriage than the estimated marriage penalties were to the timing of parenthood. Excluding individuals who got married just before (or after) having their first child reduces the child penalty somewhat in Brazil, while it hardly affects the child penalty in China.

The preceding discussion suggests that, to fully address the temporal proximity of marriage and parenthood, marriage and child penalties should be estimated jointly. Fortunately, even if

many couples get married and have children in relatively close succession, there is enough variation in the time between first marriage and first child to conduct joint event studies of the two. We therefore extend our baseline specification (1) to include event time dummies for both first child and first marriage, *i.e.*

$$Y_{itt'}^g = \alpha^g \cdot D_{it}^{\text{Child}} + \eta^g \cdot D_{it'}^{\text{Marriage}} + \beta^g \cdot D_{it}^{\text{Age}} + \gamma^g \cdot D_{it}^{\text{Year}} + v_{itt'}^g. \quad (5)$$

Here, the employment outcome of individual i at child event time t and marriage event time t' , $Y_{itt'}^g$, is regressed on event time dummies for each event, age dummies, and year dummies. The vectors of event time dummies for first child birth and first marriage omit a base year before the event.¹⁶ The coefficients $\alpha_i^g \in \alpha^g$ and $\eta_i^g \in \eta^g$ measure the employment effect of child birth and marriage, respectively, relative to the base year. As before, we scale the estimated employment effects by a counterfactual employment level, but here we consider a counterfactual without the effects of both children and marriage.¹⁷ Finally, we define child and marriage penalties as we have done previously—the average effect of each event on women relative to men over a ten-year horizon—as well as a *family penalty* that capture the total effect of marriage and parenthood.

Figure 13 presents results for a subset of countries: Brazil, China, Japan, Mauritius, Rwanda, Sweden, U.K., and Zambia. These countries have been chosen to span a wide range of geography and income levels. Several insights are worth emphasizing. First, high-income countries feature virtually no marriage penalty. The graphs for Sweden and the U.K. look similar to the baseline child penalty graphs presented above and the overall family penalty is about the same as the child penalty alone. Second, most low- and middle-income countries feature sizable marriage penalties in addition to child penalties. The family penalties for these countries are therefore considerably larger than the child penalties discussed above. Third, compared to the baseline child penalty specification, the horse-race specification in (5) reassigns part of the child penalty to a marriage penalty in some countries. In Brazil and Japan, for instance, the child penalties are smaller than those presented earlier, while the overall family penalties from both children and marriage are larger. Finally, in some African countries—Rwanda is an example—we see neither child nor marriage penalties.

The joint event study specification (5) requires information on the precise timing of marriage, which is not available for all countries in our data. To investigate the relative importance of marriage and parenthood across all countries, we consider a less granular specification:

$$Y_{it}^g = \alpha^g \cdot \text{Child}_i + \eta^g \cdot \text{EverMarried}_i + \beta^g \cdot D_{it}^{\text{Age}} + \gamma^g \cdot D_{it}^{\text{Year}} + v_{it}^g, \quad (6)$$

where the event time dummies for first marriage are replaced with a single dummy for being married ever, and the event time dummies for first child birth are replaced with a single dummy for having a child. Otherwise, the specification is the same as before.

Figure 14 presents results from this simpler specification. It shows binscatters of child penalties, marriage penalties, family penalties, and child-to-family penalty ratios by GDP per capita. Each panel includes a quadratic fit, depicted by the red line. Child penalties increase strongly and almost linearly over the course of economic development. Marriage penalties are large at low and intermediate levels of development, but fall towards zero at higher levels of development. Family penalties that capture the combined impact of marriage and parenthood are always larger than child penalties and they increase less with development. As a result, the child-to-family penalty

16. Specifically, we omit the dummies for $t, t' = -2$.

17. Specifically, counterfactual employment is calculated as the average predicted outcome from (5) when omitting the contribution of both child and marriage coefficients.

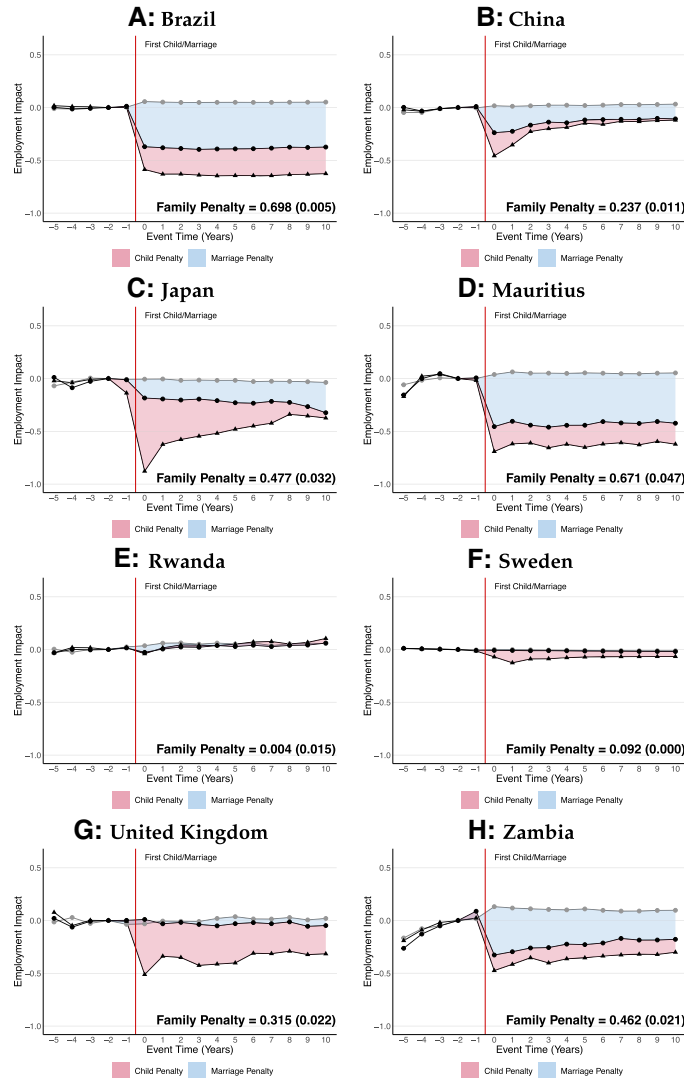


FIGURE 13

Family penalties by country. Joint event studies of first marriage and first child

Notes: This figure presents stacked event studies of first marriage and first child in selected countries. These are based on the joint event study specification (5). Each panel shows an event study of the impact of first marriage (blue area) and first child (pink area) on the employment rate of men (grey series) and women (black series). The series depict the percentage impact of marriage and children, respectively, at each event time t . Each panel also displays an average family penalty—the sum of marriage and child penalties—over event times 0–10. Age at first marriage and age at first birth are both restricted to 20–45.

ratio increases strongly with development, from about 25% to almost 100%. In other words, the development process is associated with a substitution from marriage penalties to child penalties.

To conclude, the analysis in this section refines our understanding of child penalties and the impact of family formation on gender gaps. We have shown that child and marriage penalties can be separately identified based on a horse-race event study specification that uses variation in the time between first marriage and first child. In a number of countries, particularly those at low and intermediate income levels, women face significant marriage penalties in addition

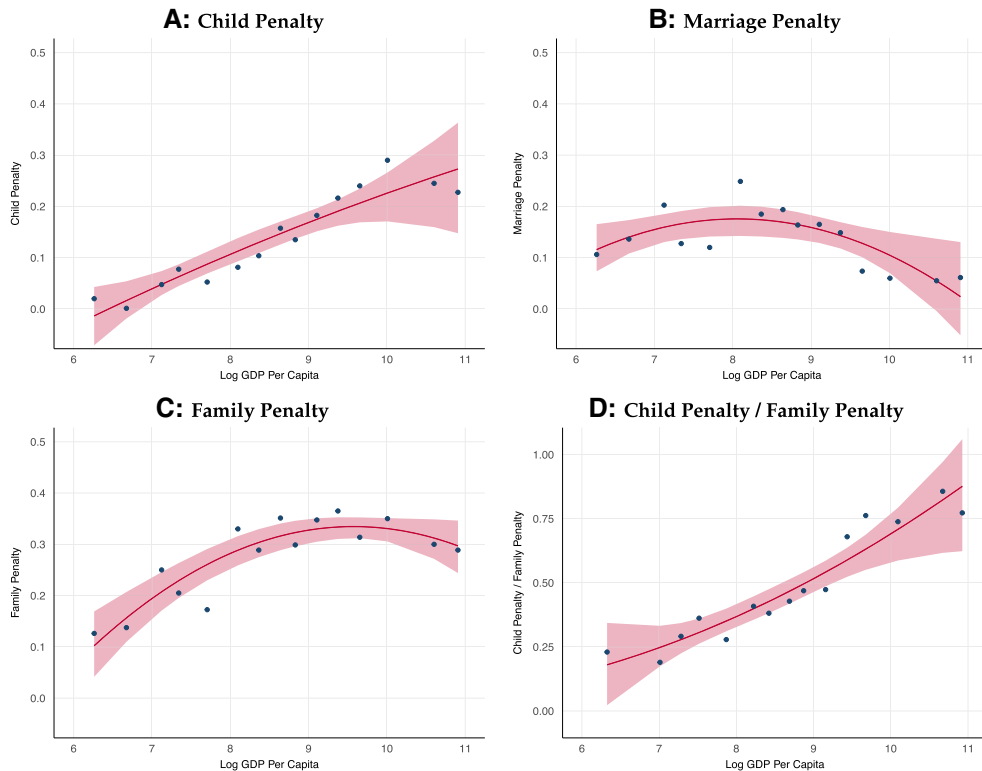


FIGURE 14

Substitution between marriage and child penalties over the course of development

Notes: This figure shows binscatters of child penalties (Panel A), marriage penalties (Panel B), family penalties (Panel C), and child-to-family penalty ratios (Panel D) by GDP per capita. These are unscaled penalties in the employment rate obtained from specification (6). The family penalty is defined as the sum of child and marriage penalties. Countries have been grouped into 15 equal-sized bins of GDP per capita, and the y-axis values show averages across all countries in each bin. Each panel also includes a quadratic fit, depicted by the red lines. The GDP data come from the World Bank's WDI for the year 2019.

to child penalties. As far as we are aware, this is a novel finding in the gender literature. It reinforces the general idea that gender inequality in labour market outcomes is driven largely by the unequal impact of family formation on women and men. In fact, a natural interpretation of marriage penalties is that they represent *anticipated* child penalties considering that marriage is essentially a precondition for procreation in many cultures.

6. CHILD PENALTIES, MARRIAGE PENALTIES, AND ECONOMIC DEVELOPMENT

In this section, we investigate the relationship between child penalties, marriage penalties, and economic development in greater detail. Figure 15 decomposes the gender gap in employment at each level of GDP per capita into a child-related gender gap, a marriage-related gender gap, and a residual gender gap. To create the decomposition, we use local linear regressions (LOESS) to smooth the GDP profile of each gender gap—the total gap, the child-related gap, and the marriage-related gap, respectively. The residual gap equals the difference between the smoothed total gap and the sum of the smoothed child- and marriage-related gaps. The residual captures all factors that cannot be directly attributed to family formation such as education differences and discrimination.




FIGURE 15
Decomposition of gender gap by levels of development

Notes: This figure decomposes the gender gap in employment at each level of GDP per capita into a child-related gender gap (dark blue), a marriage-related gender gap (light blue), and a residual gender gap (grey). The gaps are shown in absolute terms, *i.e.* the percentage point difference in the employment rates of men and women. The estimates of child and marriage penalties used for the decomposition are obtained from specification (6). The decomposition is implemented using LOESS to smooth the GDP profile of each gender gap—the total gap, the child-related gap, and the marriage-related gap—and calculating the residual gap as the difference between the smoothed total gap and the sum of the smoothed child- and marriage-related gaps. The GDP data come from the World Bank’s WDI for the year 2019.

Three key patterns emerge from the figure. The first is the inverted U-shaped relationship between the total gender gap and development, a pattern that is now well-documented (see *e.g.* Ngai *et al.* 2022). At early stages of development, female employment is decreasing in GDP per capita and the gender gap is therefore increasing. The gender gap reaches a peak around midway in the development process and then starts to decline. The second pattern is the secular rise in child-related gender inequality. Child penalties are virtually non-existent at early stages of development, but increase steadily as income per capita grows. They reach a plateau at late stages of development before decreasing slightly at the very highest income levels. This implies that almost all of the decline in gender inequality during the transition from middle- to high-income country is due to factors other than parenthood.¹⁸ The third pattern is the evolution of marriage-related gender inequality. Unlike child penalties, marriage penalties are sizable even at low levels of development. They stay sizable at intermediate levels of development and then decline at advanced levels. The total impact of family formation—marriage and children combined—explains about 50% of the gender gap at low and intermediate income levels and about 100% at high-income levels. In fact, the residual gender gap turns marginally negative

18. This is consistent with the historical patterns documented for Denmark (Kleven *et al.* 2019b) and Austria (Kleven *et al.* 2024).

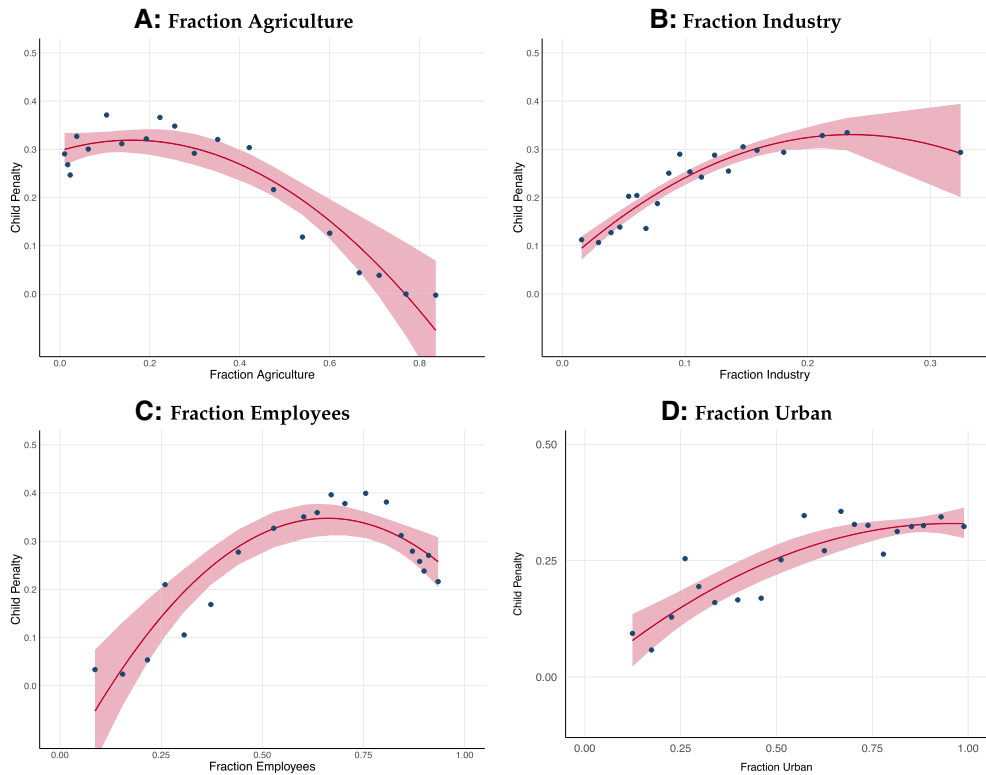


FIGURE 16

Child penalties and structural transformation

Notes: This figure shows bincatters of child penalties against the fraction of the labour force in agriculture (Panel A), the fraction of the labour force in industry (Panel B), the fraction of the labour force in salaried employment (Panel C), and the fraction of the population living in urban areas (Panel D). Each panel includes a quadratic fit, depicted by a red line. The analysis includes all 134 countries in our dataset, and it exploits within-country variation by disaggregating large countries into subnational regions. Countries are disaggregated when (i) the population is greater than 100 million people and (ii) our data include information on region of residence. In each panel, different locations (countries or regions) have been grouped into 20 ventiles of the given development indicator, and the y-axis shows the average child penalty across all locations in each ventile. All development indicators are obtained from the same data and sample as the child penalty estimates.

at the highest income levels. This means that, if not for the impact of family formation on the gender division of labour, women would be ahead of men in advanced economies.

What explains the rise of child penalties over the course of economic development? To answer this question, we investigate how child penalties correlate with different dimensions of structural transformation: industrial structure, salaried work, and urbanization. In Figure 16, we show bincatters of child penalties against the fraction of the labour force in agriculture (Panel A), the fraction of the labour force in industry (Panel B), the fraction of the labour force in salaried employment (Panel C), and the fraction of the population living in urban areas (Panel D). Each panel includes a quadratic fit, depicted by the red line. The analysis includes all 134 countries in our data. To exploit within-country variation in child penalties and development indicators, large countries are disaggregated into subnational regions.¹⁹

19. Specifically, we disaggregate countries where (i) the population is greater than 100 million people and (ii) our data include information on region of residence. These are Indonesia (provinces), Mexico (states), Vietnam (regions), China (provinces), Brazil (states), Egypt (governorates), and the U.S. (states).

The relationships are very strong. We see that child penalties are declining in the importance of agriculture, increasing in the importance of industry, increasing in the importance of salaried work and formalization, and finally increasing in urbanization. Repeating the exercise for family penalties—the sum of child and marriage penalties—produces qualitatively similar relationships, but somewhat less strong. This reflects that, over the course of economic development, the rise in child penalties is partially offset by a fall in marriage penalties as we have documented above.

Taken together, our evidence provides a clearer picture of the nexus between child penalties, marriage penalties, gender inequality, and economic development. At early stages of development, the economy is close to subsistence levels and most of the population is self-employed, works in agriculture, and lives in rural areas. The gender gap is significant, but it is not driven by the arrival of children. It is instead driven by factors that predate the arrival of children, including marriage, education, and norms. As economies develop, the structure of the labour market changes: there is a shift from agriculture to industry and services along with a rise in salaried employment and urbanization. This creates a separation between the home and the workplace, which leads to a specialization between child care and market work. This specialization is highly gendered, thus increasing child penalties and creating the institution of the stay-at-home mom. At late stages of development, women close both the education gap and the marriage-related gap, but child penalties are sticky and prevent full gender convergence. At this point, eliminating gender inequality in labour market outcomes becomes almost synonymous with eliminating child penalties, consistent with recent findings from specific high-income countries such as Denmark and the U.S. (Kleven *et al.* 2019a, 2019b; Kleven 2023).

7. CHILD PENALTIES OVER THE LONG RUN OF HISTORY

Thus far, our analysis of the link between economic development, child penalties, and gender inequality has been cross-sectional, comparing countries at different development levels today. In this section, we take advantage of the historical data described in Section 2.5 to demonstrate the robustness of our results to looking at countries longitudinally, over the course of their own development. A compelling feature of the pseudo-event study framework is that it allows for quasi-experimental analyses of child penalties over the very long run (back to the 1700s for some countries).

7.1. *United States 1880–2020*

We begin the investigation with the U.S., where cross-sectional data on family structure, gender, and employment are continuously available for the full population (or large representative samples of the population) between 1880 and 2020. This allows us to estimate child penalties over 140 years. During this period, the U.S. went from having a GDP per capita similar to that of lower-income countries today (around the level of Namibia or Mongolia) to being one of the richest and most developed countries in the world. In Figure 17, we present event studies of first child birth from decadal census data, grouping the years 1880–1910 (Panel A), 1920–40 (Panel B), and 1950–70 (Panel C). These panels show unscaled impacts on the employment rates of men and women. Panel D of the figure shows the time series of both unscaled and scaled child penalties between 1880 and 2020, adding estimates for recent decades from Kleven (2023). The decade-level event studies that underlie the estimates in this time-series graph are provided in Figure A.9 of the Appendix.

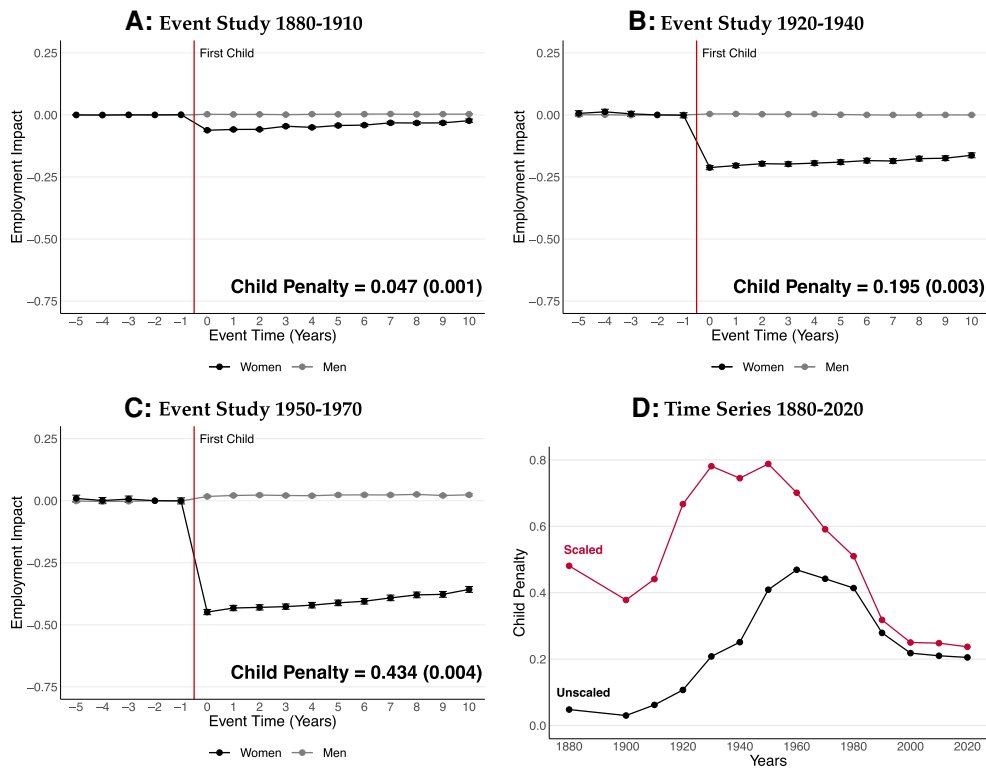


FIGURE 17
 Child penalties in the U.S., 1880–2020

Notes: This figure illustrates the evolution of child penalties in the U.S. between 1880 and 2020. Panels A–C present event studies of the impact of first child birth, depicted by the vertical line, on the employment rate of men (grey series) and women (black series) in different historical time periods. Each series depicts the unscaled impact (pp) of child birth at each event time t , i.e. \hat{a}_t^m and \hat{a}_t^w in equation (1). Each panel also displays the average unscaled child penalty over event times 0–10. Panel D presents the full time series of decade-level child penalties between 1880 and 2020. This panel shows both unscaled penalties (pp) and scaled penalties (%). Age at first birth is restricted to 20–45. The error bars depict 95% confidence intervals based on robust standard errors.

U.S. child penalties have changed enormously over the course of history. Four insights are worth highlighting. First, unscaled child penalties were remarkably small in the late 19th century. The arrival of children was associated with a decline in the female employment rate of less than 5pp. The scaled child penalty, on the other hand, was much larger (about 50%) because the baseline employment rate of women without children was very small at the time. In other words, in the 1800s, gender inequality mostly predated the arrival of children. This is consistent with our analysis of lower-income countries today, where gender inequality is mostly driven by marriage penalties and other factors that predate family formation. Second, child penalties increased steeply over the first half of the 20th century. The unscaled penalty increased by a factor of more than 10 between 1900 and 1960. Around 1960, the drop in the female employment rate at first child birth was 47pp—similar to the staggering numbers for many middle-income countries in Latin America and Asia today. Third, following their mid-century peak, child penalties began to decline steadily up to the 1990s. The unscaled child penalty was roughly halved over this period, while the scaled child penalty fell by a factor of almost four. The difference in the decline of scaled and unscaled penalties reflects that baseline employment increased strongly over this period, resulting from a massive labour market entry of young women prior to having children.

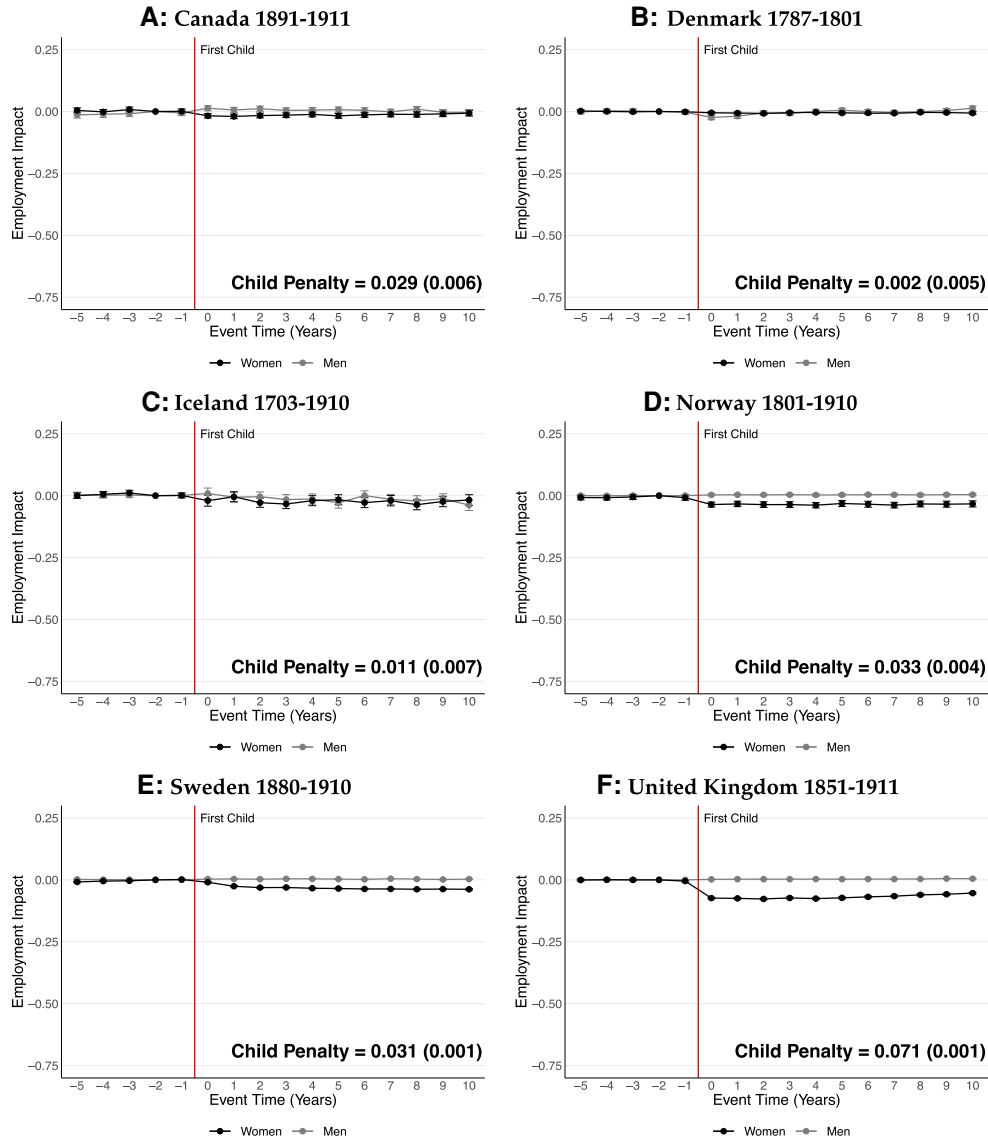


FIGURE 18
Historical child penalties in other countries

Notes: This figure presents event studies of first child birth using historical data from different high-income countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (grey series) and women (black series). The series depict the unscaled impact (pp) of child birth at each event time t , *i.e.* $\hat{\alpha}_t^m$ and $\hat{\alpha}_t^w$ in equation (1). Each panel also displays the average unscaled child penalty over event times 0–10. Age at first birth is restricted to 20–45. The error bars depict 95% confidence intervals based on robust standard errors.

This increase in the baseline pre-child employment rate of women (relative to men) is also the reason why child penalties grow in importance over the course of development: child penalties fall in the later stages of development, but the overall gender gap falls even more, and therefore child penalties gradually take over as the main driver of gender inequality. Finally, we note that both scaled and unscaled child penalties have plateaued in the U.S. over the last thirty years, explaining the slowdown in gender convergence during this period (see also [Kleven 2023](#)).

The evidence in Figure 17 demonstrates the similarities between the historical evolution of child penalties in the U.S. and the cross-sectional variation across development levels today. This strengthens our structural transformation interpretation of the global variation in child penalties.

7.2. *Other countries*

Using historical data for six other developed countries, we confirm that small child penalties are a robust feature of countries at early development stages. Figure 18 presents historical event studies of first child birth in Canada (1891–1911), Denmark (1787–1801), Iceland (1703–1910), Norway (1801–1910), Sweden (1880–1910), and the U.K. (1851–1911). The data come from historical censuses (see Section 2.5) and we have grouped all available census years for each country. The event studies look remarkably compelling and statistically precise considering the time periods being analysed, validating the usefulness of the pseudo-event study methodology. In all countries, child penalties are either zero (such as in the almost fully agricultural economies of Iceland and Denmark in the 1700s) or very small. The largest child penalty (about 7pp) is found in the U.K. in the second half of 19th century. Comparing these results to the contemporaneous atlas presented above, we see that the historical penalties are much smaller than the current penalties in all of these countries. This implies that child penalties must have grown over the course of development, similar to what we observe in the U.S.²⁰

8. CONCLUSION

This paper builds a global atlas of child penalties in employment based on harmonized and nationally representative microdata for 134 countries. The estimates are based on pseudo-event studies of first child birth using cross-sectional data. Our data include about 281 million individual \times year observations across all countries and represent more than 95% of the world's population.

Most countries display clear and sizable child penalties: men and women follow parallel trends before parenthood, but diverge sharply and persistently after parenthood. While the qualitative patterns tend to be similar across countries, there is enormous variation in the magnitude of the effects in different regions of the world and often between countries in close proximity. At the continent level, Latin America has the largest child penalties in the world, while Africa and Asia have the smallest child penalties. These averages mask strong heterogeneity across countries within continents. Asia is a case in point: child penalties are small in China and much of Southeast Asia, but extremely large in the Middle East, South Asia, and Japan. Similarly stark differences are found in other parts of the world such as Scandinavia versus Central Europe and Central Africa versus Northern/Southern Africa. We show that the fraction of the gender gap explained by child penalties varies systematically with economic development and proxies for structural transformation. At early stages of development, child penalties represent a minuscule fraction of gender inequality. But as economies develop—incomes rise and the labour market transitions from subsistence agriculture towards salaried work in industry and services—child penalties gradually take over as the dominant driver of gender inequality in the labour market. The relationship between child penalties and development is validated using historical data from current high-income countries, back to the 1700s for some countries.

20. While current high-income countries had small child penalties historically, they had large marriage penalties—again consistent with what we find for low-income countries today. Figure A.10 in the Appendix compares child penalties and marriage penalties estimated from the historical data. The figure shows that, historically, marriage penalties were an order of magnitude larger than child penalties in all countries included in our study.

The paper also presents novel evidence on marriage penalties. Parenthood is closely linked to the institution of marriage—especially in lower-income settings where marriage is virtually a precondition for procreation—implying that the two should be analysed in conjunction. Marriage penalties are sizable at early and intermediate stages of development, but converge towards zero at later stages. The combined effects of marriage and children on gender gaps—what we label family penalties—are larger than child penalties and increase somewhat less over the course of development. They explain about half of gender inequality in low-income countries and all of gender inequality in high-income countries. In general, accounting for the effect of marriage reinforces the view that family formation is critical to understanding variation in gender inequality across countries and with development.

To conclude, the child penalty atlas expands our understanding of global gender inequality in the labour market. By documenting child penalties across settings with fundamentally different political, economic, and cultural institutions, the atlas allows for richer analyses of mechanisms than within-country analyses of a single institutional setting. This is necessary for understanding how factors like government policy, labour market structure, and social norms together shape the evolution of child penalties and gender convergence. This is an important agenda for future research.

Acknowledgments. We thank Oriana Bandiera, Thomas Chaney, Xavier Jaravel, Ilyana Kuziemko, Thomas Piketty, Daniel Reck, Nina Roussille, Johannes Spinnewijn, Owen Zidar, and numerous seminar/workshop participants for comments and discussions. We are grateful to Martin Eckhoff Andresen, Emily Nix, Andreas Steinhauer, and Jakob Egholt Sogaard for their help with some of the administrative data sources used in the paper. We are grateful to Mary Reader, Eva Demsky, Sveva Manfredi, and Christian Höhne for their outstanding research assistance.

Funding

C.L. acknowledges financial support from the European Research Council consolidator grant #101001464, from Pivotal Ventures, and from the Bill and Melinda Gates Foundation.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data and code underlying this research is available on Zenodo at <https://doi.org/10.5281/zenodo.13742225>

REFERENCES

- AARONSON, D., DEHEJIA, R., JORDAN, A. *et al.* (2021), “The Effect of Fertility on Mothers’ Labor Supply over the Last Two Centuries”, *The Economic Journal*, **131**, 1–32.
- AGUIERO, J. M. and MARKS, M. S. (2008), “Motherhood and Female Labor Force Participation: Evidence from Infertility Shocks”, *American Economic Review: Papers and Proceedings*, **98**, 500–504.
- ALTONJI, J. G. and BLANK, R. M. (1999), “Race and Gender in the Labor Market”, in Ashenfelter, O. and Card, D. (eds) *Handbook of Labor Economics* (Vol. 3, Chap. 48) (Amsterdam: Elsevier).
- ANDRESEN, M. E. and NIX, E. (2022), “What Causes the Child Penalty? Evidence from Adopting and Same Sex Couples”, *Journal of Labor Economics*, **40**, 971–1004.
- ANGELOV, N., JOHANSSON, P. and LINDAHL, E. (2016), “Parenthood and the Gender Gap in Pay”, *Journal of Labor Economics*, **34**, 545–579.
- ANGRIST, J. D. and EVANS, W. N. (1998), “Children and Their Parents’ Labor Supply: Evidence from Exogenous Variation in Family Size”, *American Economic Review*, **88**, 450–477.
- BERTRAND, M. (2011), “New Perspectives on Gender”, in Ashenfelter, O. and Card, D. (eds) *Handbook of Labor Economics* (Vol. 4b, Chap. 17) (Amsterdam: Elsevier).
- BERTRAND, M., GOLDIN, C. and KATZ, L. F. (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. D. and KAHN, L. M. (2017), “The Gender Wage Gap: Extent, Trends, and Explanations”, *Journal of Economic Literature*, **55**, 789–865.
- BRONARS, S. G. and GROGGER, J. (1994), “The Economic Consequences of Unwed Motherhood: Using Twin Births as a Natural Experiment”, *American Economic Review*, **84**, 1141–1156.

- CORTÉS, P. and PAN, J. (2023), “Children and the Remaining Gender Gaps in the Labor Market”, *Journal of Economic Literature*, **61**, 1359–1409.
- DEATON, A. (1985), “Panel Data from Time Series of Cross-Sections”, *Journal of Econometrics*, **30**, 109–126.
- GALLEN, Y., JOENSEN, J. S., JOHANSEN, E. R., *et al.* (2023), “The Labor Market Returns to Delaying Pregnancy” (Working Paper).
- GOLDIN, C. (2014), “A Grand Gender Convergence: Its Last Chapter”, *American Economic Review*, **104**, 1091–1119.
- GOLDIN, C. and KATZ, L. F. (2016), “A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation”, *Journal of Labor Economics*, **34**, 705–746.
- KLEVEN, H. (2023), “The Geography of Child Penalties and Gender Norms: A Pseudo-Event Study Approach” (Working Paper No. 30176, NBER).
- KLEVEN, H., LANDAIS, C., POSCH, J., *et al.* (2019a), “Child Penalties Across Countries: Evidence and Explanations”, *AEA Papers and Proceedings*, **109**, 122–126.
- , *et al.* (2024), “Do Family Policies Reduce Gender Inequality? Evidence from 60 Years of Policy Experimentation”, *American Economic Journal: Economic Policy*, **16**, 110–149.
- KLEVEN, H., LANDAIS, C. and SØGAARD, J. E. (2019b), “Children and Gender Inequality: Evidence from Denmark”, *American Economic Journal: Applied Economics*, **11**, 181–209.
- (2021), “Does Biology Drive Child Penalties? Evidence from Biological and Adoptive Families”, *American Economic Review: Insights*, **3**, 183–198.
- LUNDBORG, P., PLUG, E. and RASMUSSEN, A. W. (2017), “Can Women Have Children and a Career? IV Evidence from IVF Treatments”, *American Economic Review*, **107**, 1611–1637.
- NGAI, L. R., OLIVETTI, C. and PETRONGOLO, B. (2022), “Structural Transformation Over 150 Years of Men’s and Women’s Work” (Working Paper).
- ROSENZWEIG, M. R. and WOLPIN, K. I. (1980), “Life-Cycle Labor Supply and Fertility: Causal Inferences from Household Models”, *Journal of Political Economy*, **88**, 328–348.