

The Child Penalty Atlas*

Henrik Kleven[†]

Camille Landais[‡]

Gabriel Leite-Mariante[§]

September 2023

Abstract

This paper builds a world atlas of child penalties in employment based on micro data 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 labor market transitions from subsistence agriculture towards salaried work in industry and services — child penalties take over as the dominant driver of gender inequality. 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 economic development. The development process is associated with a substitution from marriage penalties to child penalties, with the former gradually converging to zero.

*We thank Oriana Bandiera, 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, and Christian Höhne for their outstanding research assistance.

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

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

[§]London School of Economics. Email: g.leite-mariante@lse.ac.uk.

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 labor market outcomes (Kleven, Landais, and Sogaard 2019; Kleven, Landais, Posch, Steinhauer, and Zweimüller 2019). These findings suggest that the study of child penalties should take center stage in gender inequality research. A key qualification, however, 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 micro data from 134 countries, representing more than 6 billion people or three-quarters 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 this 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 across five different continents, showing that the pseudo-event study approach produces results that are virtually identical to those one would obtain from a true event study approach.

We focus on child penalties in employment. The event study evidence on employment effects is generally very compelling, 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 tend to have much larger sample sizes than panel data. 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 precisely 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? There is a simple translation between the level of child penalties and the fraction of gender inequality explained by child penalties. Due to the correlation between child penalties and raw gender gaps in the data, 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 until they start falling slightly at very high 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 very high levels of development, eliminating gender inequality is synonymous with eliminating child penalties.

We argue that the increasing importance of child penalties over the course of development results from the structural transformation of labor 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 US and Denmark (Goldin 2014; Goldin and Katz 2016; Kleven, Landais, and Sogaard 2019). 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 labor 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 labor market outcomes.

This 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 labor market outcomes, including [Bertrand, Goldin, and Katz \(2010\)](#), [Angelov, Johansson, and Lindahl \(2016\)](#), [Kleven, Landais, and Sogaard \(2019\)](#), [Kleven, Landais, Posch, Steinhauer, and Zweimüller \(2019\)](#), [Cortés and Pan \(2021\)](#), [Kleven, Landais, and Sogaard \(2021\)](#), [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, constructing a world atlas of child penalties and investigating their importance for gender inequality at different levels of development and under a wide range of institutions. In addition, 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.

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 that account for the combined impact of marriage and parenthood on gender gaps. Section 6 analyzes the relationship between child penalties (or family penalties) and economic development. Section 7 concludes.

2 Data

To construct the child penalty atlas, we harmonize nationally representative micro data from 134 countries. Our sample is representative of approximately 6.1 billion people or three-quarters of the global population in 2023. 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 harmo-

nized 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 seventy-seven countries, our data come from de-identified samples of national censuses. For seventy-five 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 vs 1970-2015 for Mexico. They cover countries in the Americas, Africa, Asia, and Europe, spanning a wide range of income levels.

For sixteen 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 labor 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 seventeen 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 using the same procedure as for LIS data. 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 it also contains 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 United Kingdom, we use data from the Annual Population Survey (APS), an expanded version of the national Labor Force Survey covering the period 2012-2020. It contains information on weekly employment and household composition, from which we infer the presence of children. For the United States, we use pooled data from the Current Population Survey (CPS) and the American Community Survey (ACS), covering the years from 1968 to 2020.¹

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.² For the remaining two countries, we use data from nationally representative panel surveys. For Australia, we use data from the

¹The estimates for the US come from [Kleven \(2023\)](#).

²The 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, Landais, and Søgård \(2019\)](#) and [Kleven, Landais, Posch, Steinhauer, and Zweimüller \(2019\)](#), but the estimates presented here are slightly different as the data have been updated to include more recent years.

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 UK, the Panel Study of Income Dynamics (PSID) for the US, the National Longitudinal Survey of Youth (NLSY) also for the US, 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 duration. 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 duration.

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.³ The map divides countries into deciles of the raw gender gap, with darker colors representing larger gaps. There is enormous variation in employment gaps around the world, from about 5pp in gender-equal countries to about 90pp in gender-unequal countries. Gender gaps are relatively small in high-income countries, East Asia, and Central Africa, while they are large in Latin America, North Africa, the Middle East, and South Asia. There is also significant variation within these regions. An objective of this paper is to investigate if these large differences in gender gaps can be explained by differences in the impact of family formation — marriage and children — on women relative to men. We will show that most of the global variation in gender inequality reflects variation in child penalties (and to a smaller extent marriage penalties as well).

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, Landais, and Sogaard \(2019\)](#), the following specification is run separately for men and women:

$$Y_{it}^g = \alpha^g \cdot D_{it}^{Event} + \beta^g \cdot D_{it}^{Age} + \gamma^g \cdot D_{it}^{Year} + \nu_{it}^g, \quad (1)$$

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 $\alpha_t^g \in \alpha^g$ measure the impact

³Denoting the employment rate by L^g for gender $g = w, m$, the raw gender gap is defined as $(L^m - L^w) / L^m$.

of child birth on gender g in event year t , relative to the base year.⁴ 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 labor market outcomes of women relative to men. [Kleven, Landais, and Sogaard \(2019\)](#) 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.⁵ A positive child penalty implies that parenthood increases the gender gap.

3.2 Pseudo-Event Study Approach

Description The event study approach described above is straightforward to implement given access to high-quality panel data. Unfortunately, as explained in Section 2, such data are not available outside a handful of high-income countries, limiting our understanding of how child penalties vary across geography. However, 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

⁴Throughout 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.

⁵Due 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.

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 a particular event time relative to first child birth. For individuals observed with children in the cross-section, this is easily done: the age of their oldest child determines the position in positive event time, $t \geq 0$. The main challenge pertains to individuals observed without children in the cross-section: 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 \mathbf{X}_i . This parent is matched to a childless individual j observed in year $y - n$ with age $a - n$ and the same characteristics $\mathbf{X}_j = \mathbf{X}_i$. This gives a surrogate observation for $t = -n$.⁶ 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.⁷ This procedure implies that age and calendar time are changing along the event time dimension, exactly as in the panel data approach of Kleven, Landaïs, and Sogaard (2019).⁸ 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).⁹

Apart from year, age, and gender, the variables we include in our matching algorithm are 4 categories of education (less than primary education, primary education, secondary education, and post-secondary education), 4 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 vs urban

⁶A 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.

⁷To clarify, there is a slight difference in the matching protocol depending on the definition of the employment measure available in the data source considered. The protocol described above is used when considering cases where we observe weekly employment (obtained from a question about work activities *last week*). This protocol needs to be adjusted when considering cases where we observe 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 $t = 1, \dots, T + 1$.

⁸For 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 $\mathbf{X}_j = \mathbf{X}_i$, 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.

⁹Similar 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.

area. This represents a relatively light set of matching variables, consistently defined and available across all our data sources.

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 labor market outcomes of men around $t = 0$, this reflects bias from positive selection. Figures A.2-A.3 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 evaluating the pseudo-event study approach is to validate it against 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 presents such a comparison of pseudo-event studies and actual event studies in four specific countries: Austria, Chile, Denmark, and Japan. These countries differ in context as well as in the size and dynamics of their child penalties. Moreover, while administrative panel data are available for the universe of the population in the first three countries, the panel data available in Japan come from a survey with a much smaller sample size. The figure shows that, for all four countries, and irrespective of their differences in context and data sources, the pseudo-panel and panel estimates align remarkably well. The two

methods give virtually identical estimates of the child penalty. Figure 4 confirms this conclusion in a broader set of countries. The figure 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 more than .95.

4 The Child Penalty Atlas

4.1 A View From Above: Continent-Level Child Penalties

Figure 5 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)-(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 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 employment for women. While qualitatively similar, the quantitative magnitudes vary substantially across continents. Child penalties are largest in Latin America: the average impact of parenthood on the employment rate of women relative to men equals 38%. Child penalties are more modest in Africa and Asia, average impacts of 9% and 18% respectively, and at intermediate levels in Europe, North America, and Oceania. In the next section, we investigate the variation across countries within each continent.

4.2 A Tale of Great Variation: Country-Level Child Penalties

Figure A.4 in the appendix provides event studies of first child birth for all 134 countries in our sample. Here we consider a subset of these countries, selected to highlight key features of the data. The discussion is organized by continent.

Europe: Figure 6 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 divergence between the neighboring 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 6 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, Landais, Posch, Steinhauer, and Zweimüller 2022](#)) as well as the timing and number of subsequent child births.

Asia: Figure 7 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.4), 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 neighboring countries such as China vs South Korea (4% vs 49%), Thailand vs Malaysia (5% vs 45%), and Myanmar vs Bangladesh (19% vs 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.

Latin America: Figure 8 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 percent, 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 8, 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.

Africa: In Figure 9, 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 penalties are large in Morocco (41%) and in South Africa (28%). Appendix Figure A.4 shows that this regional divide is a general feature of the data.

4.3 Urban Areas See Larger Effects: City-Level Child Penalties

Our data and empirical approach allow for estimating child penalties at the city level. Figure 10 presents 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 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 UK), 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 United States, child penalties are positively associated with urbanization, and he argues that this is an effect of differences in urban vs 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 11 presents a heatmap of child penalties. In this map, countries are divided into deciles of the child penalty, with darker colors representing larger penalties.¹⁰ 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 United States, Canada, France, and Australia. As mentioned in the beginning of the paper, the countries currently in the atlas represent more than 6 billion people or three-quarters 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.¹¹ Our ultimate goal is to build out the atlas to cover the last quarter of the global population as well, or at least most of it.

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

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

where 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 absent children for men and women vs actual employment for men — although in practice that makes little difference as men are largely unaffected by child birth.

¹⁰Countries for which we do not have data are depicted in gray.

¹¹Although Greenland is part of the Kingdom of Denmark, the Danish administrative data do not include information on Greenlanders.

Figure 12 provides a heatmap of the fraction of the gender gap in employment explained by child penalties. Countries are again divided into deciles, and darker colors represent 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 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.

5 Marriage Penalties vs 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 13 shows the distribution of time between first child and first marriage, pooling countries in the IPUMS census data.¹² 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 labor 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 differ-

¹²The figure is based on the subset of IPUMS countries for which the timing of marriage is observed.

ently, 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 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 14 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 newlyweds 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 in this sample. The upshot of Figure 14 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.5 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.

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

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.¹³ The coefficients $\alpha_t^g \in \alpha^g$ and $\eta_{t'}^g \in \eta^g$ 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.¹⁴ 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 15 presents results for a subset of countries: Brazil, China, Japan, Mauritius, Rwanda, Sweden, United Kingdom, 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 United Kingdom 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

¹³Specifically, we omit the dummies for $t, t' = -2$.

¹⁴Specifically, counterfactual employment is calculated as the average predicted outcome from (5) when omitting the contribution of both child and marriage coefficients.

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:

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

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

Figure 16 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 a 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 on gender gaps are always larger than child penalties and they increase less with development. As a result, the child-to-family penalty ratio increase 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 to child penalties. As far as we are aware, this is a novel finding in the gender literature. It reinforces the general idea — emphasized in this paper and in the recent literature — that gender inequality in labor 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 pre-condition 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 17 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.

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., ?). 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 start 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 level of income. 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.¹⁵ 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 at the highest income levels. This means that, if not for the impact of family formation on the gender division of labor, 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 18, we show bin-

¹⁵This is consistent with the historical patterns documented for Denmark (Kleven, Landais, and Sogaard 2019) and Austria (Kleven, Landais, Posch, Steinhauer, and Zweimüller 2022).

scatters of child penalties against the fraction of the labor force in agriculture (Panel A), the fraction of the labor force in industry (Panel B), the fraction of the labor 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 data. To exploit within-country variation in child penalties and development indicators, large countries are disaggregated into subnational regions.¹⁶

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 — produce 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 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 labor 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 — perhaps due to women’s comparative advantage in child bearing and infant child care, reinforced by cultural norms — 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 labor market outcomes is virtually synonymous with eliminating child penalties, consistent with recent findings from specific high-income countries such as Denmark and the United States (Kleven, Landais, and Sogaard 2019; Kleven, Landais, Posch, Steinhauer, and Zweimüller 2019; Kleven 2023).

¹⁶Specifically, 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 US (states).

7 Conclusion

This paper builds a global atlas of child penalties in employment based on harmonized and nationally representative micro data for 134 countries. The estimates are based on pseudo-event studies of first child birth using cross-sectional data. Our data include about 312 million individual \times year observations across all countries and represent three-quarters 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 vs Central Europe and Central Africa vs 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 labor 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 labor market.

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 analyzed in conjunction. Marriage penalties are sizable at early and intermediate stages of development, but converges 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 labor 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, labor market structure, and social norms together shape the evolution of child penalties and gender convergence. This is an important agenda for future research.

References

- ALTONJI, JOSEPH G., AND REBECCA M. BLANK (1999): "Race and Gender in the Labor Market," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3, chap. 48. Elsevier: Amsterdam. [4](#)
- ANDRESEN, MARTIN E., AND EMILY NIX (2022): "What Causes the Child Penalty? Evidence from Adopting and Same Sex Couples," *Journal of Labor Economics*, 40(4), 971–1004. [4](#), [6](#)
- ANGELOV, NIKOLAY, PER JOHANSSON, AND ERICA LINDAHL (2016): "Parenthood and the Gender Gap in Pay," *Journal of Labor Economics*, 34(3), 545–579. [4](#)
- BERTRAND, MARIANNE (2011): "New Perspectives on Gender," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4b, chap. 17. Elsevier: Amsterdam. [4](#)
- BERTRAND, MARIANNE, CLAUDIA GOLDIN, AND LAWRENCE F. KATZ (2010): "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2(3), 228–255. [4](#)
- BLAU, FRANCINE D., AND LAWRENCE M. KAHN (2017): "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55(3), 789–865. [4](#)
- CORTÉS, PATRICIA, AND JESSICA PAN (2021): "Children and the Remaining Gender Gaps in the Labor Market," *Journal of Economic Literature*, Forthcoming. [4](#)
- DEATON, ANGUS (1985): "Panel Data from Time Series of Cross-Sections," *Journal of Econometrics*, 30(1), 109–126. [10](#)
- GOLDIN, CLAUDIA (2014): "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 104(4), 1091–1119. [3](#)
- GOLDIN, CLAUDIA, AND LAWRENCE F. KATZ (2016): "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation," *Journal of Labor Economics*, 34(3), 705–746. [3](#)
- KLEVEN, HENRIK (2023): "The Geography of Child Penalties and Gender Norms: Evidence from the United States," NBER Working Paper No. 30176. [1](#), [4](#), [6](#), [9](#), [15](#), [21](#)

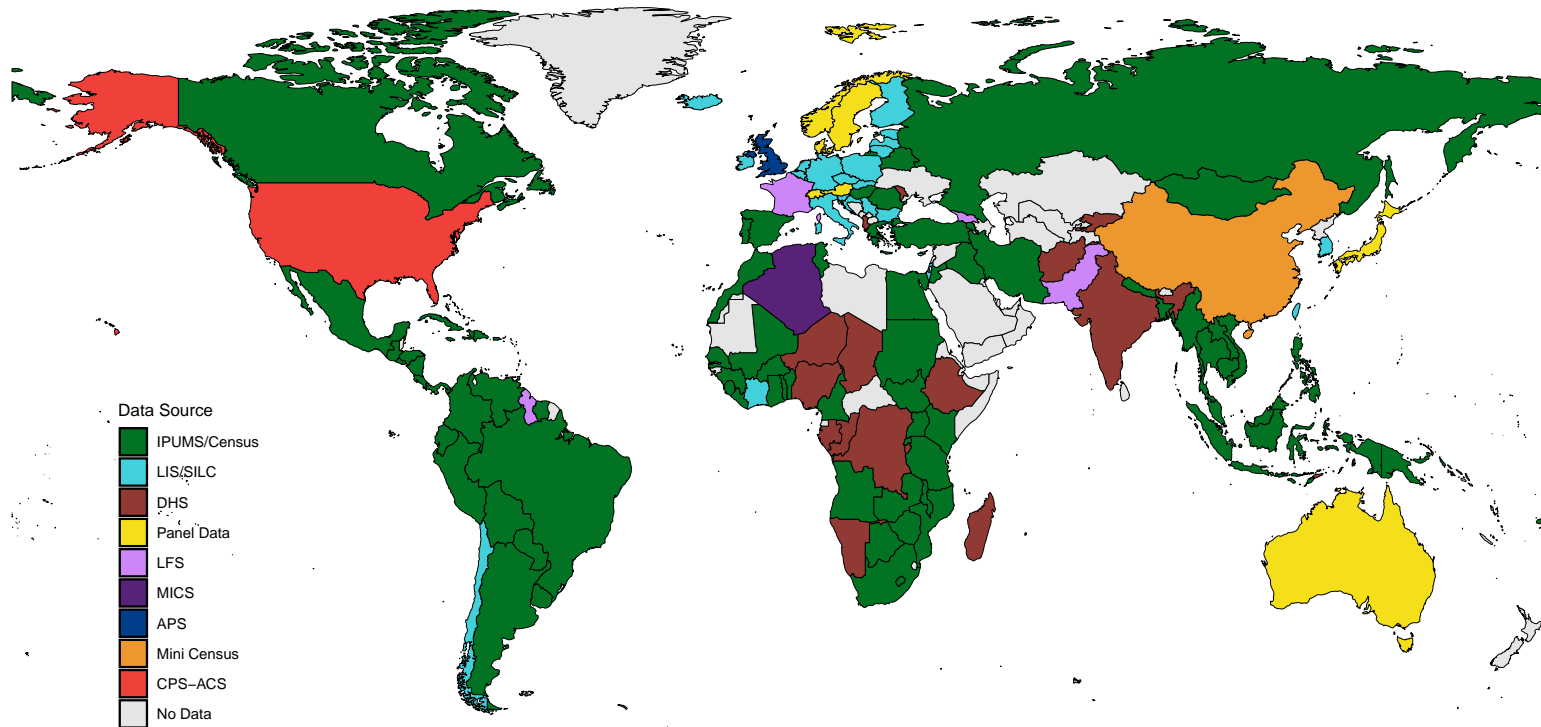
KLEVEN, HENRIK, CAMILLE LANDAIS, JOHANNA POSCH, ANDREAS STEINHAUER, AND JOSEF ZWEIMÜLLER (2019): "Child Penalties Across Countries: Evidence and Explanations," *AEA Papers and Proceedings*, 109, 122–126. [1](#), [4](#), [6](#), [21](#)

——— (2022): "Do Family Policies Reduce Gender Inequality? Evidence from 60 Years of Policy Experimentation," *American Economic Journal: Economic Policy*, Forthcoming. [13](#), [20](#)

KLEVEN, HENRIK, CAMILLE LANDAIS, AND JAKOB E. SØGAARD (2019): "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 11(4), 181–209. [1](#), [3](#), [4](#), [6](#), [8](#), [9](#), [10](#), [20](#), [21](#)

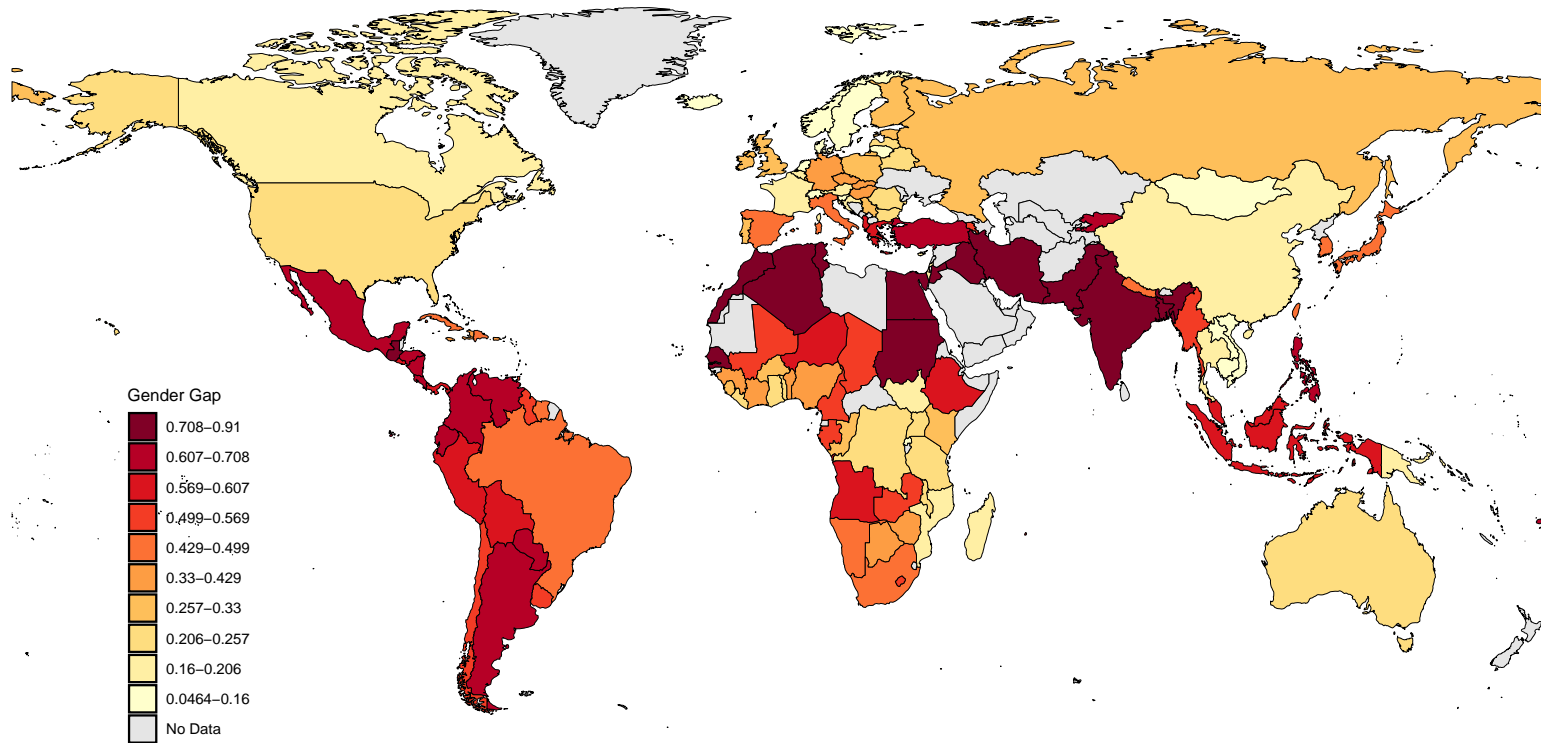
——— (2021): "Does Biology Drive Child Penalties? Evidence from Biological and Adoptive Families," *American Economic Review: Insights*, 3(2), 183–198. [4](#)

FIGURE 1: MAP OF DATA SOURCES



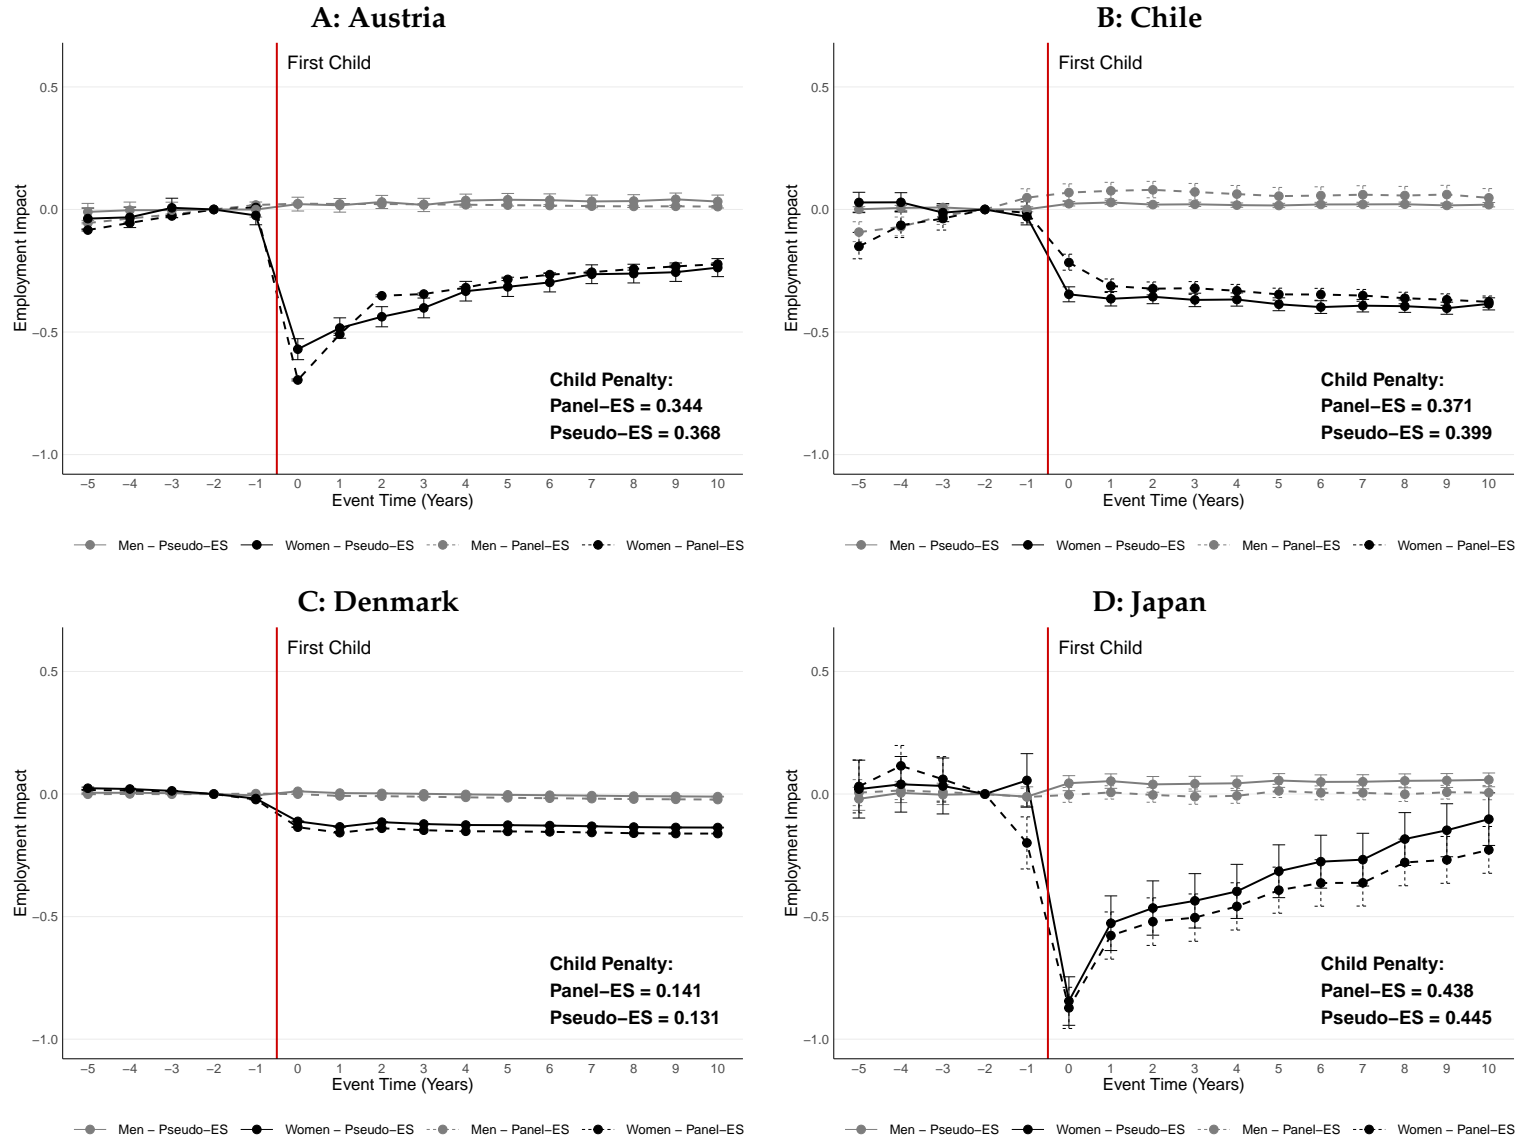
Notes: This map summarizes the data sources used for the 134 countries included in our child penalty atlas. Each color represents a specific data source. Data for the countries labeled as “Panel Data” refer to either administrative panels or panel surveys. Data for all other countries come from either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

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 colors representing larger gaps. The estimation sample includes individuals both with and without children. We restrict attention to individuals between the ages of 15 and 50. For individuals observed with children, we further restrict attention to those who had their first child between the ages of 20 and 45.

FIGURE 3: VALIDATION OF PSEUDO-EVENT STUDY APPROACH
EVENT STUDIES OF FIRST CHILD BIRTH



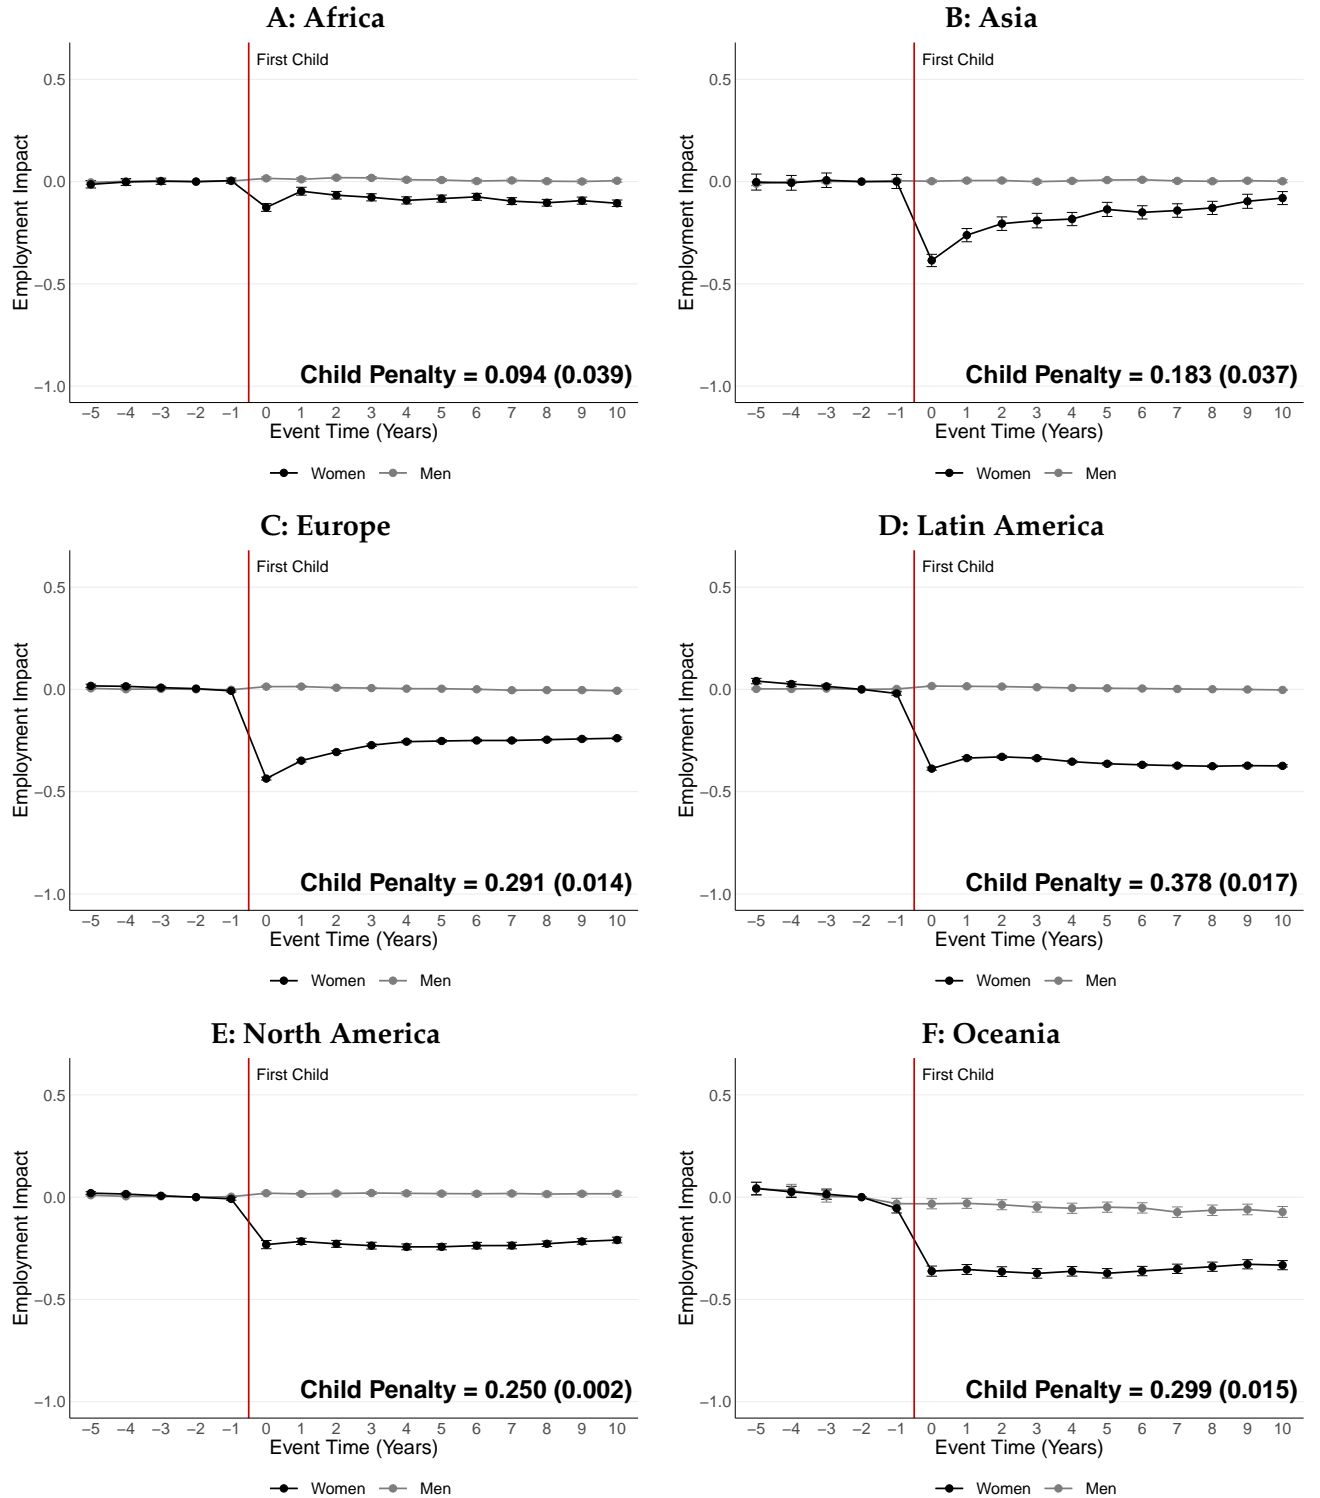
Notes: This figure validates the pseudo-event study approach (solid lines) against an actual event study approach using panel data (dashed lines) for a selection of countries. Each panel shows, for both approaches, event studies of impact of first child birth on the employment rate of men (gray series) 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 average child penalties over event times 0-10 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.

FIGURE 4: VALIDATION OF PSEUDO-EVENT STUDY APPROACH
SCATTERPLOT



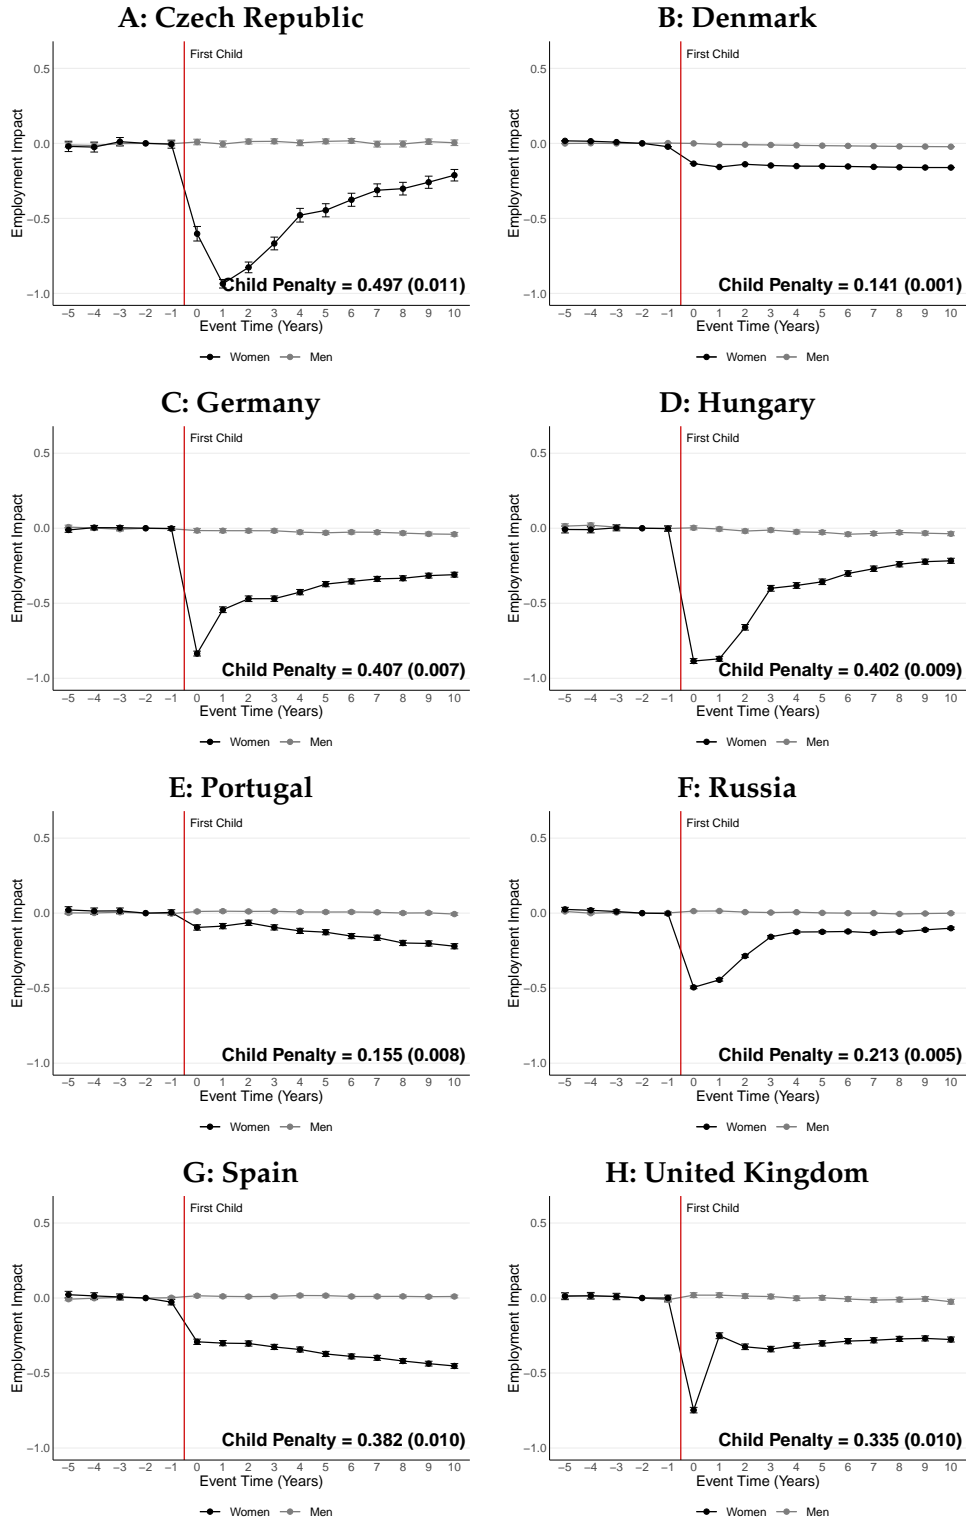
Notes: This figure validates the pseudo-event study approach against an actual event study approach for all 11 countries where representative panel data are available. The figure plots child penalties obtained from the pseudo-event study approach (y-axis) against child penalties obtained from the actual event study approach (x-axis). These are average child penalties over event times 0-10 as defined in equation (3). All child penalty pairs lie close to the 45-degree line, thus validating our empirical approach in this sample of countries.

FIGURE 5: CHILD PENALTIES BY CONTINENT



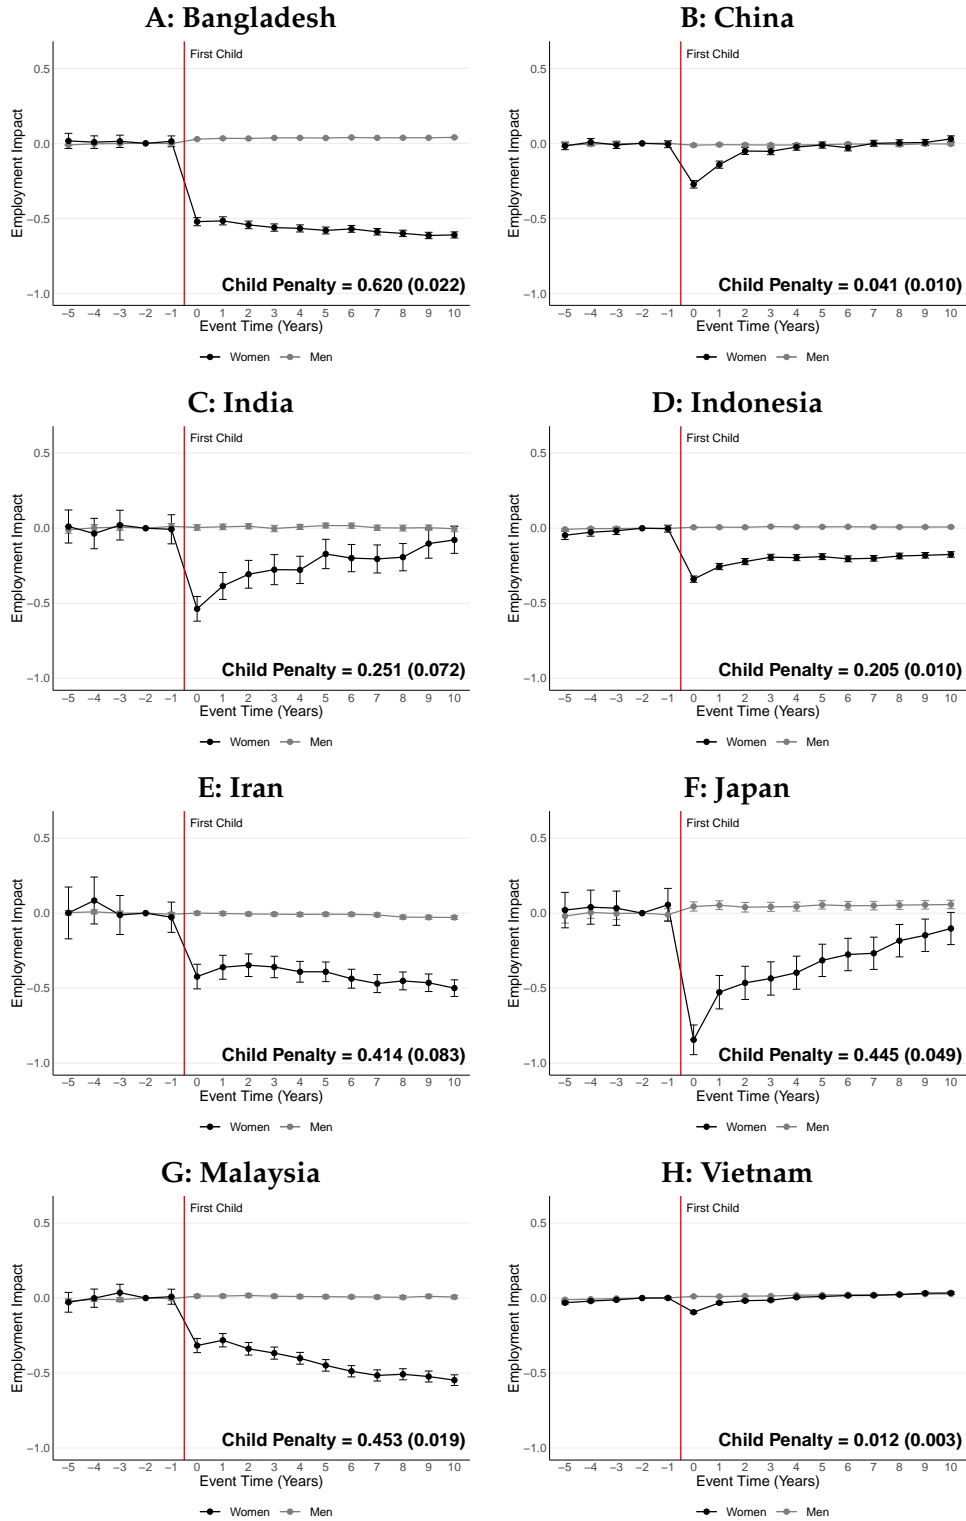
Notes: This figure presents event studies of first child birth by continent. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact of child birth at each 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.

FIGURE 6: CHILD PENALTIES IN EUROPEAN COUNTRIES



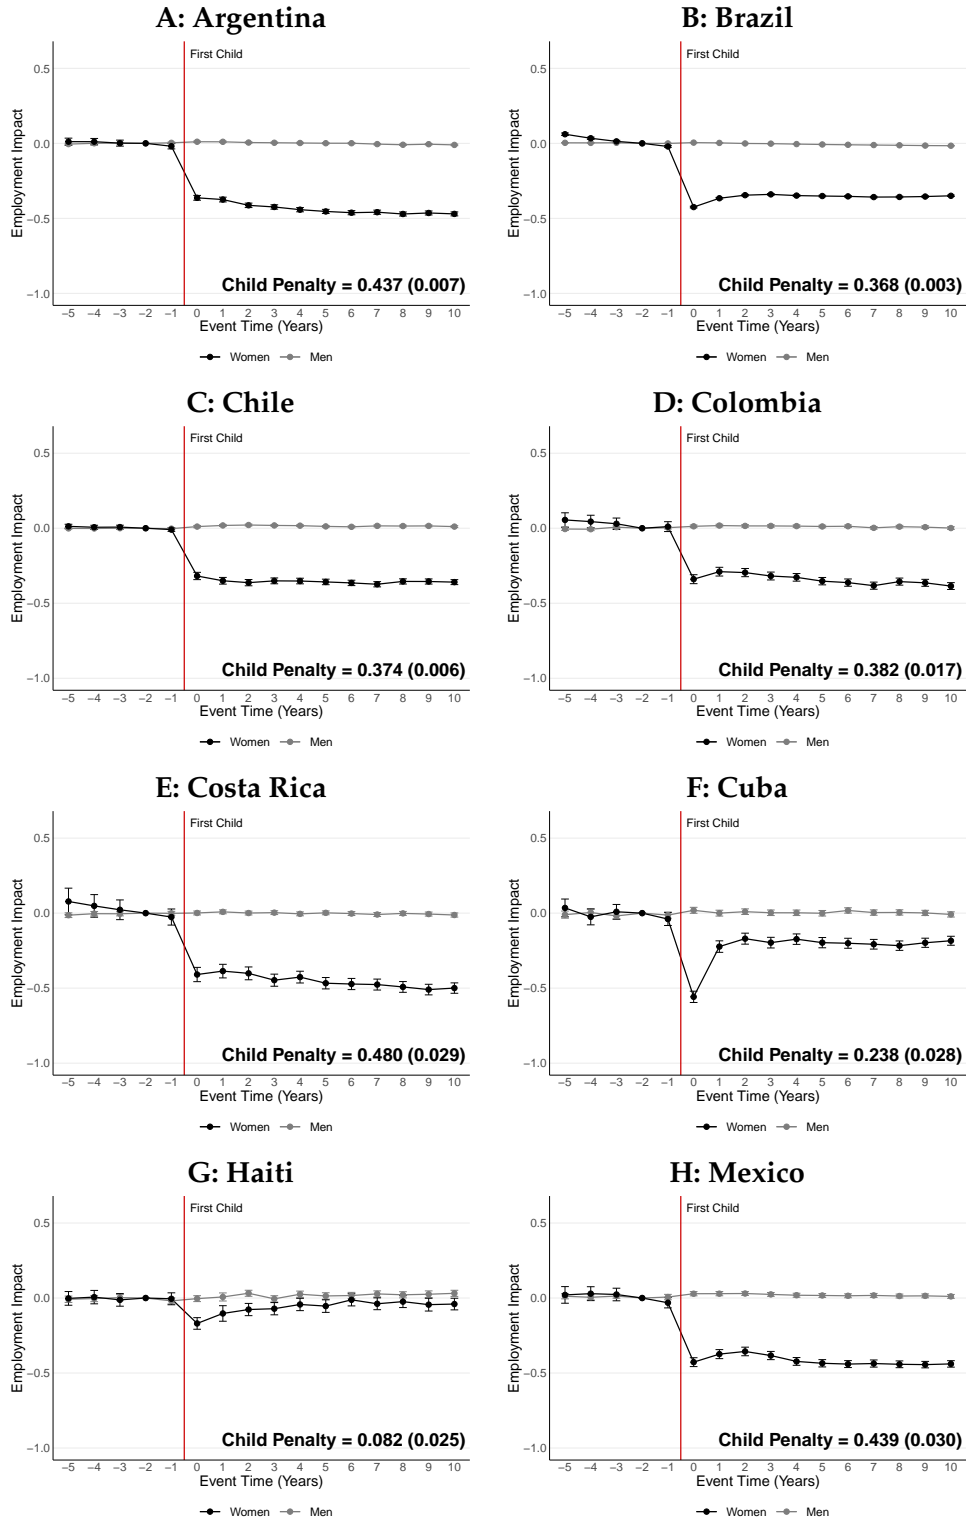
Notes: This figure presents event studies of first child birth for a selection of European countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact 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 7: CHILD PENALTIES IN ASIAN COUNTRIES



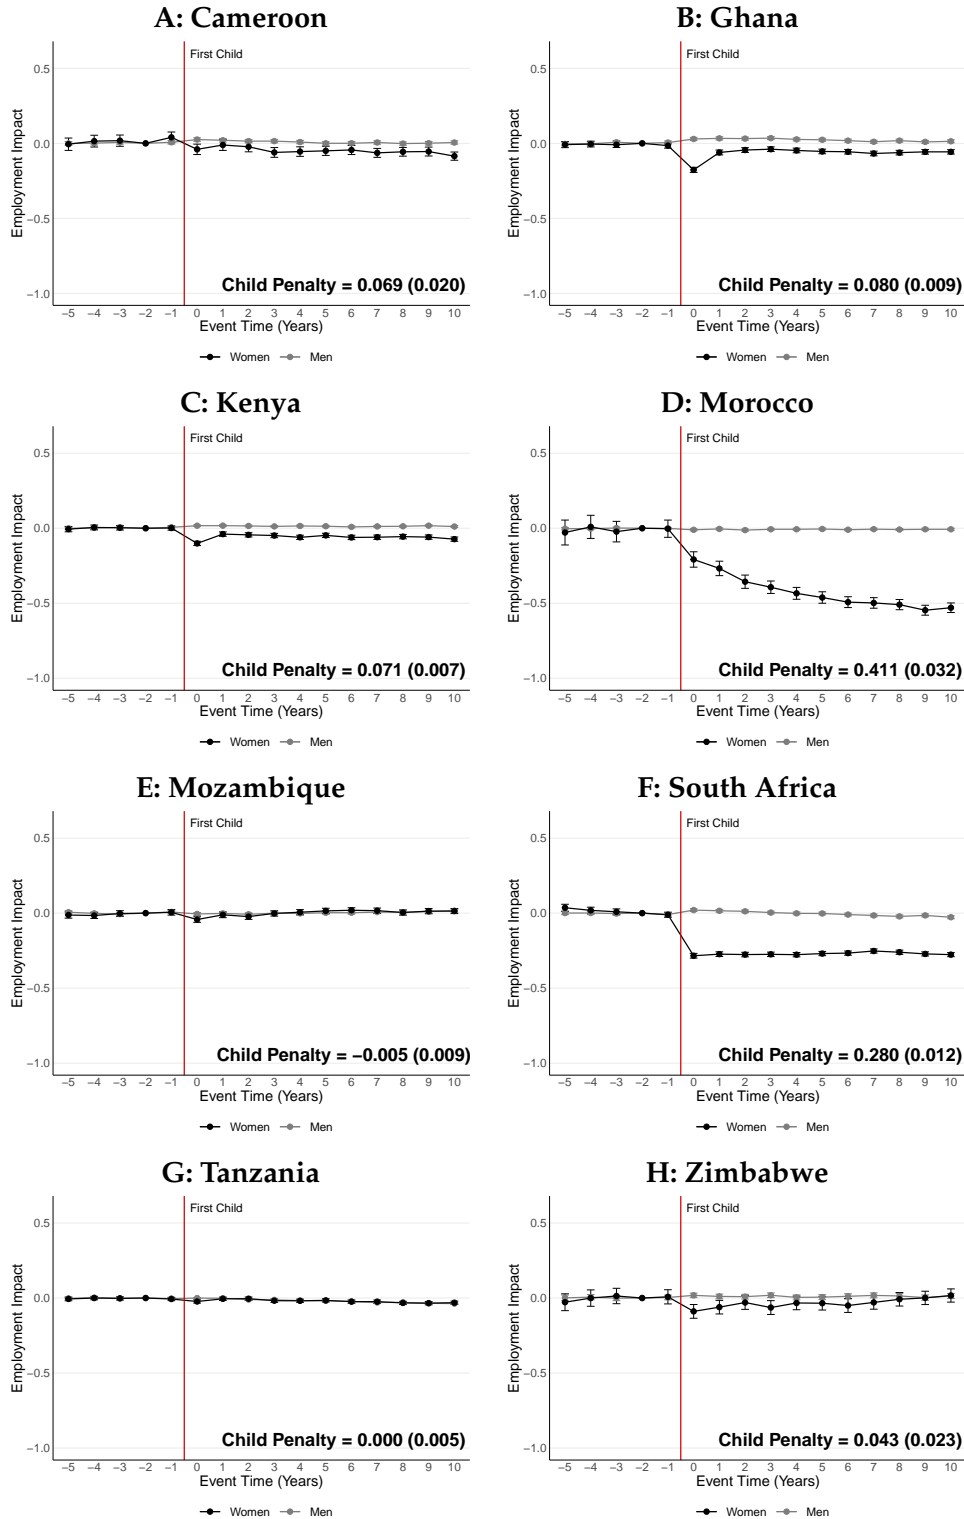
Notes: This figure presents event studies of first child birth for a selection of Asian countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact 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 8: CHILD PENALTIES IN LATIN AMERICAN COUNTRIES



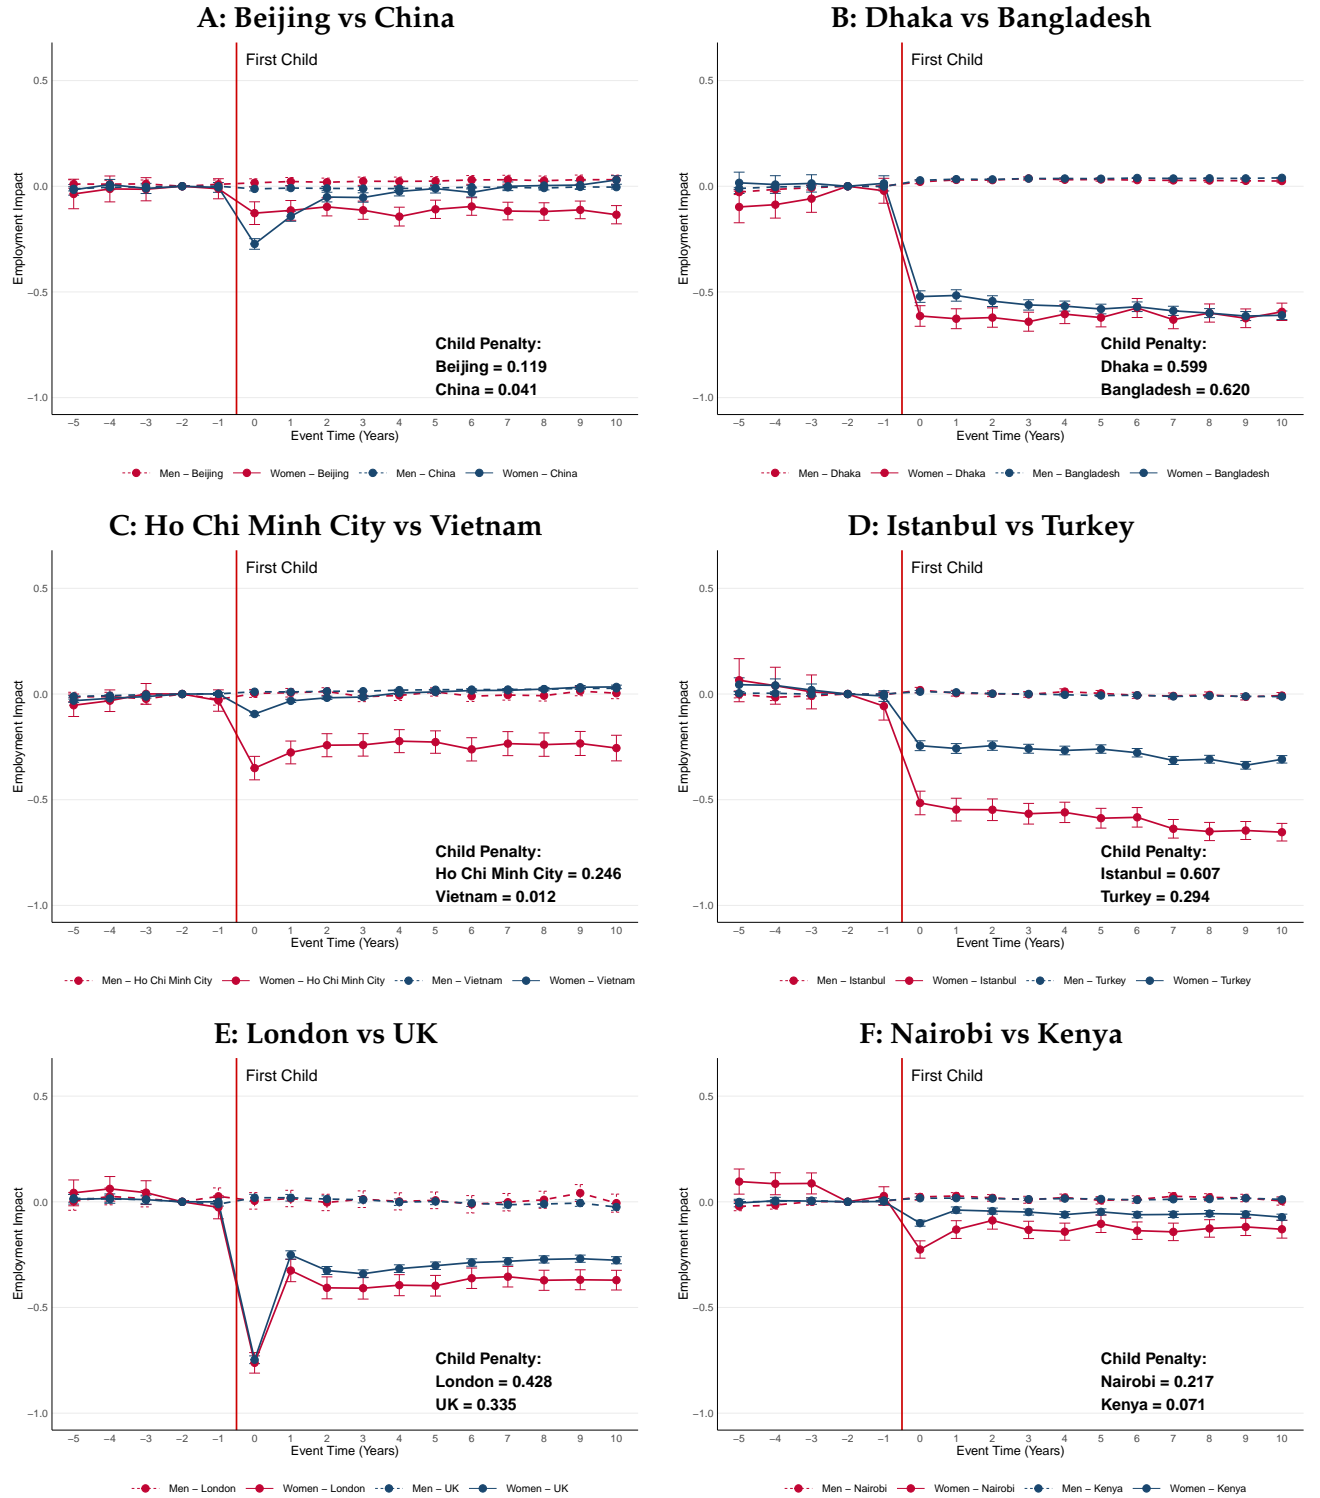
Notes: This figure presents event studies of first child birth for a selection of Latin American countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage 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 9: CHILD PENALTIES IN AFRICAN COUNTRIES



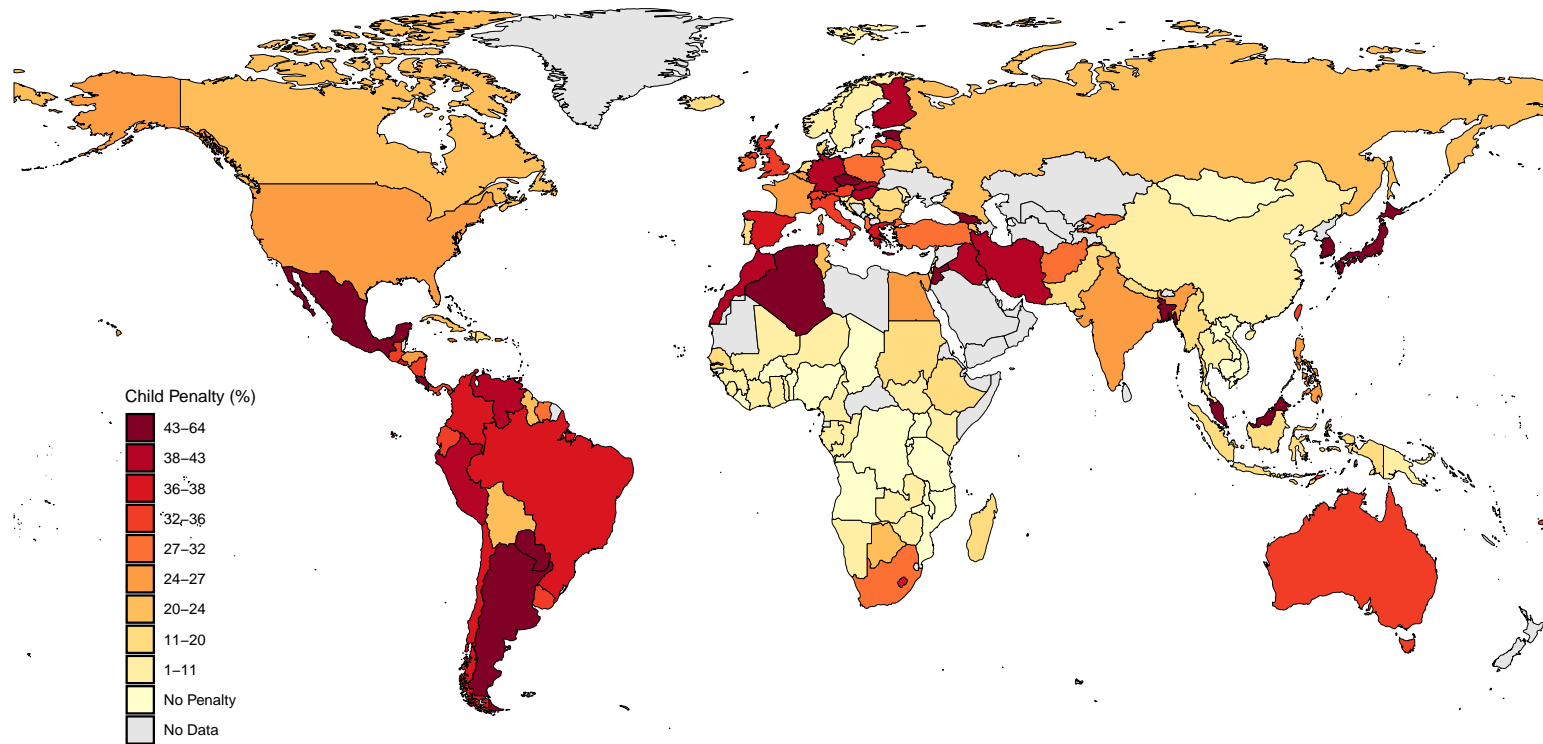
Notes: This figure presents event studies of first child birth for a selection of African countries. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact 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 10: CHILD PENALTIES IN CITIES VS COUNTRIES



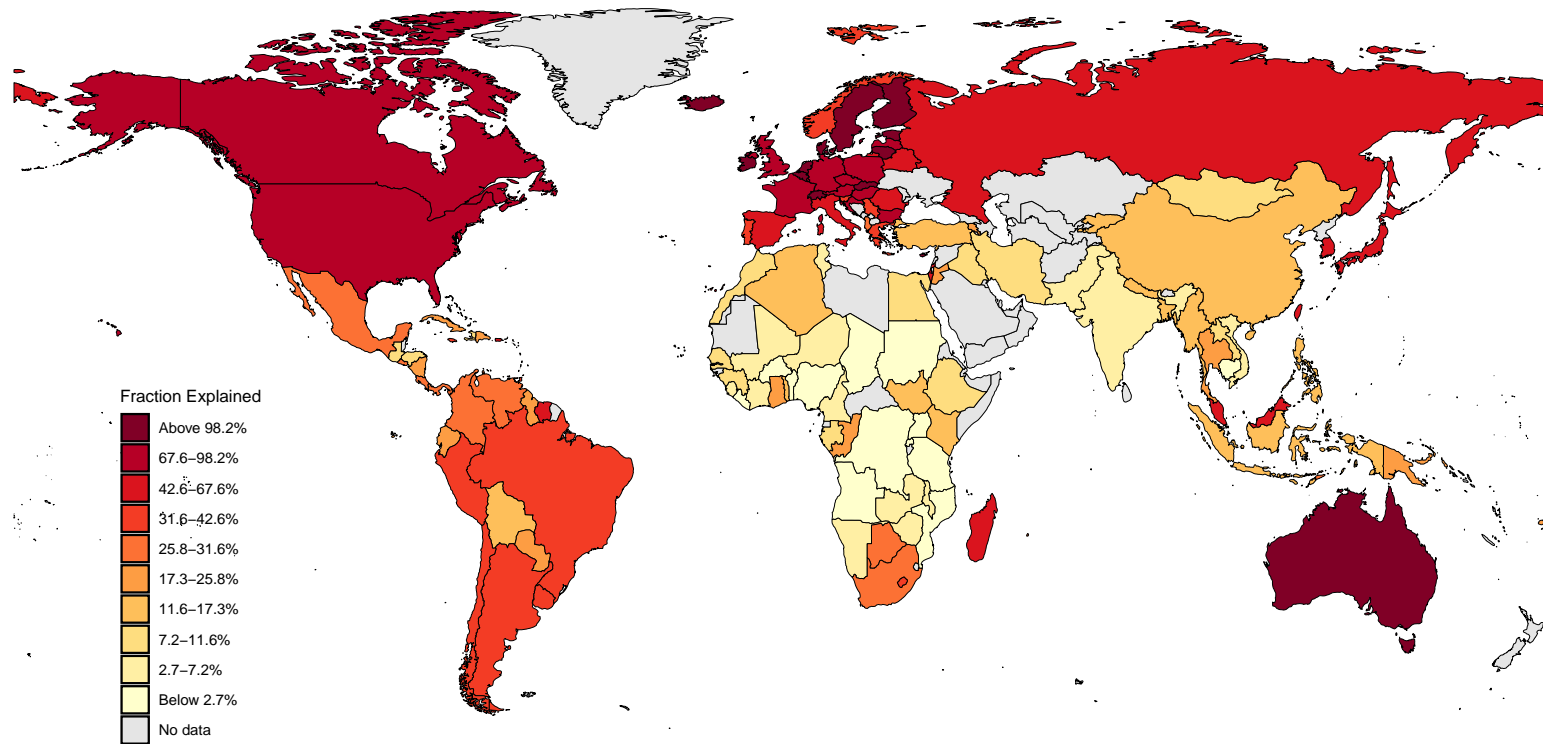
Notes: This figure compares event studies of first child birth in large cities (red lines) and countries as a whole (blue lines). Each panel shows event studies of the impact of first child on the employment rate of men (dashed lines) and women (solid lines) in cities and countries, respectively. The 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 error bars depict 95% confidence intervals based on robust standard errors. Our pseudo-event study approach yields precise estimates even at the city-level.

FIGURE 11: 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 colors 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.4.

FIGURE 12: HEATMAP OF FRACTION OF GENDER GAP EXPLAINED BY CHILD PENALTIES



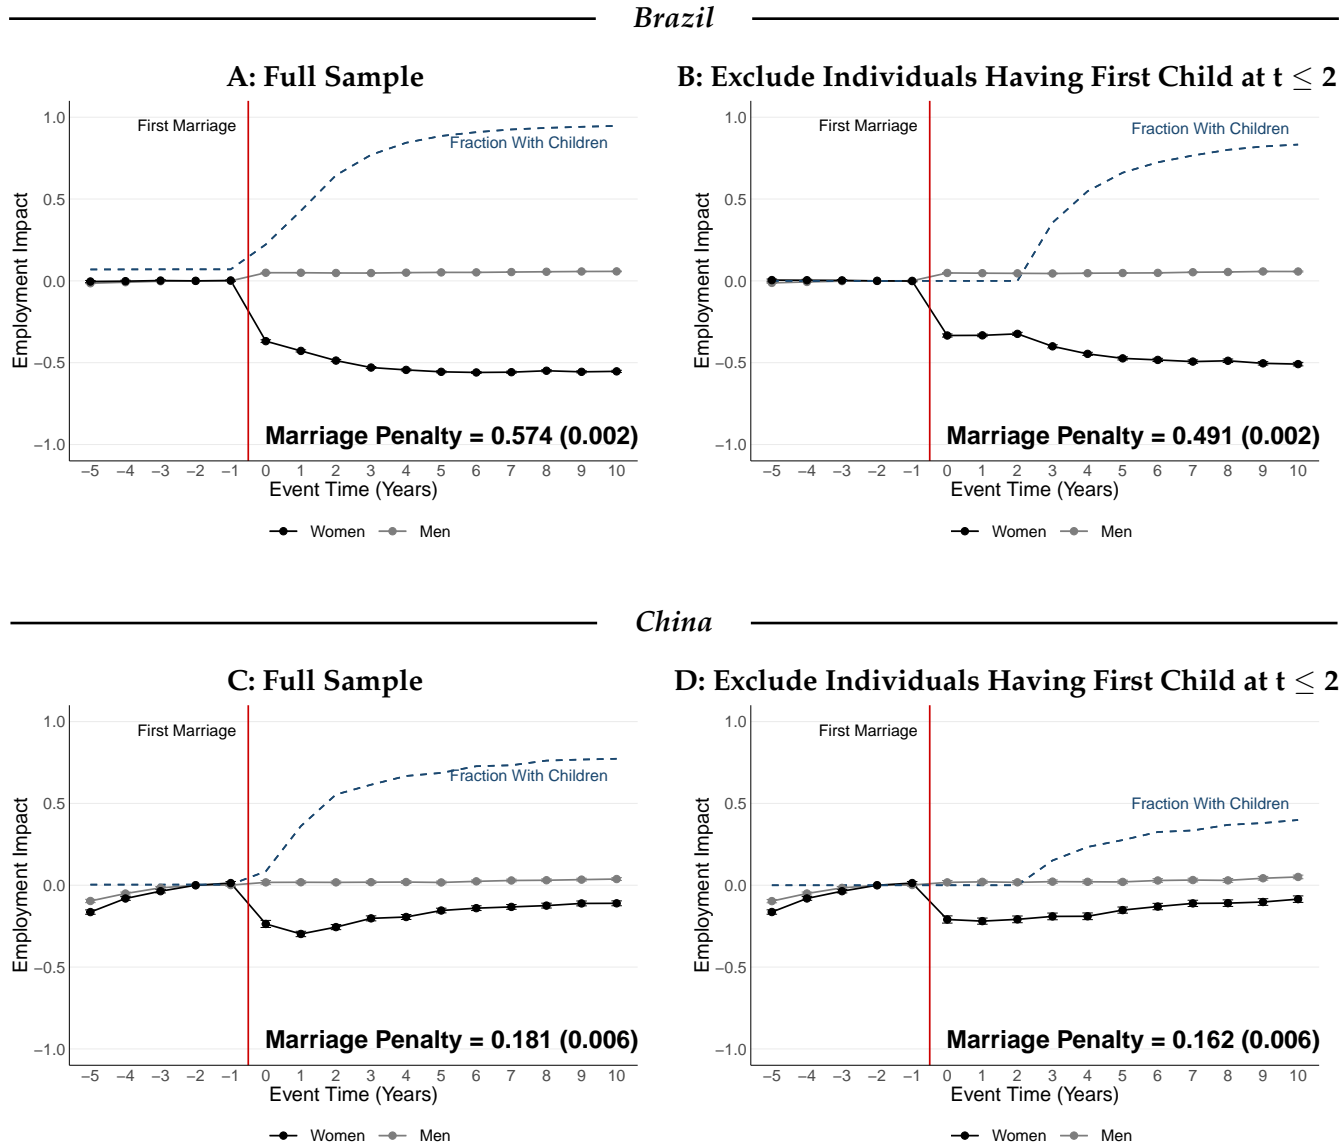
Notes: This figure presents a heatmap of the fraction of the gender gap in employment explained by child penalties. The fraction explained is calculated using equation (4). The map divides countries into deciles of the fraction explained, with darker colors representing larger fractions explained.

FIGURE 13: TIME BETWEEN FIRST CHILD AND FIRST MARRIAGE



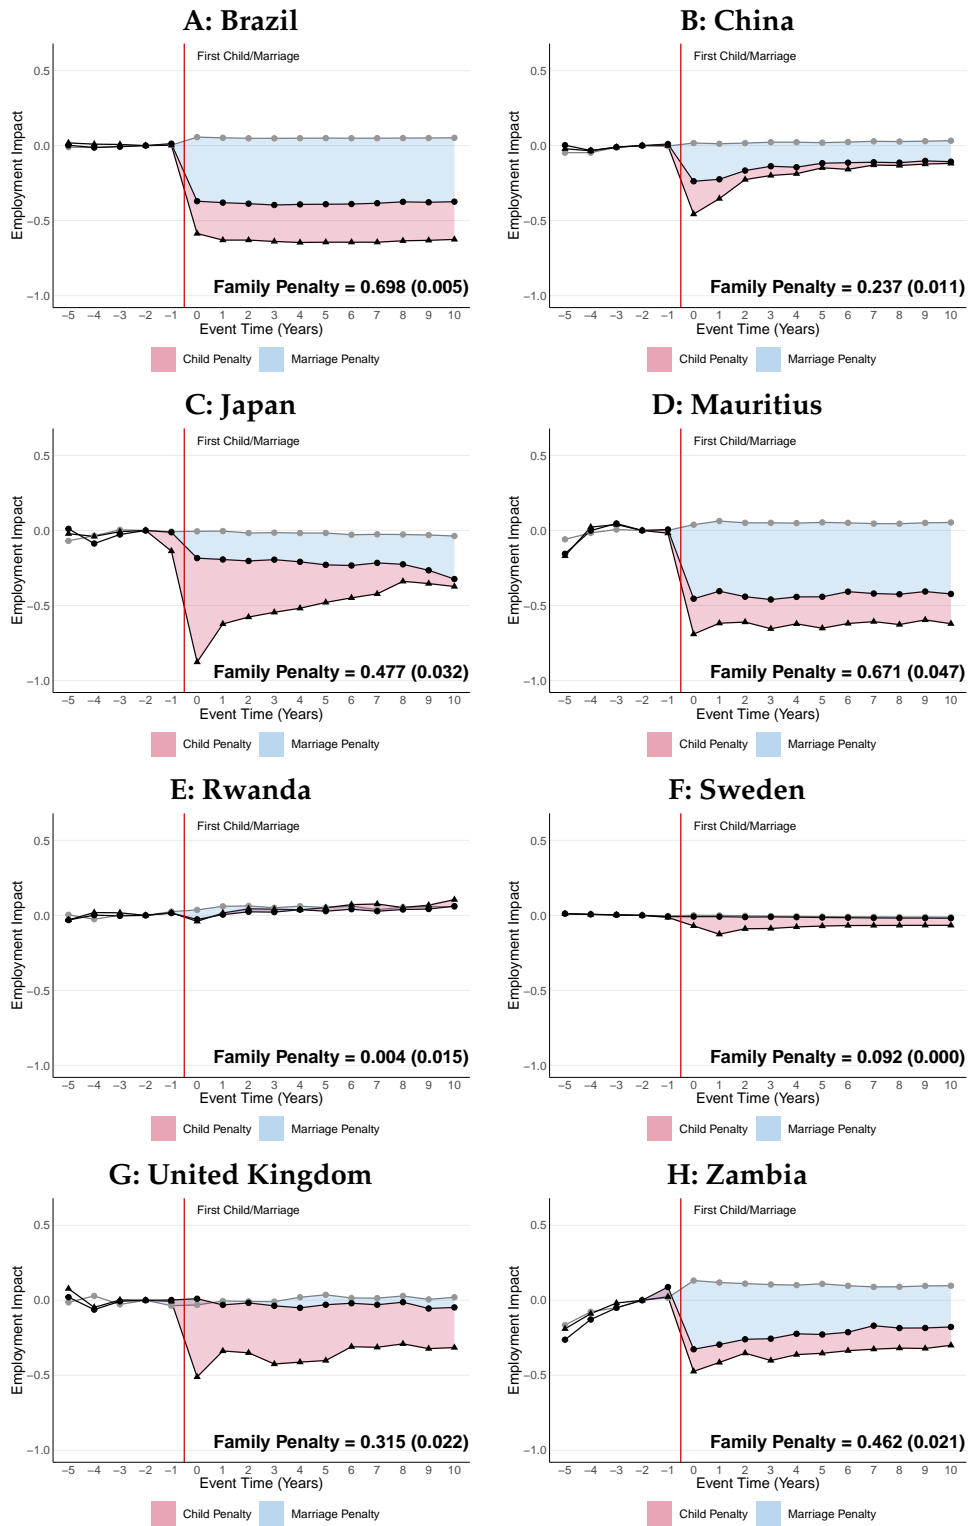
Notes: This figure shows the distribution of time between first child and first marriage, pooling countries in the IPUMS census data for which we observe the timing of both events. These are all low- and middle-income countries. The sample is restricted to individuals who are married or have ever been married. Negative numbers indicate that the first child was born before marriage, while positive numbers indicate that the first child was born after marriage. Married individuals who do not (yet) have children are assigned to the “>10” category, together with individuals whose first child was born more than 10 years after marriage. The figure also displays the total fraction of individuals who had their first child before marriage (only 6%) and the total fraction of individuals who had their first child within two years of marriage (42%).

FIGURE 14: MARRIAGE PENALTIES
A CASE STUDY OF BRAZIL AND CHINA



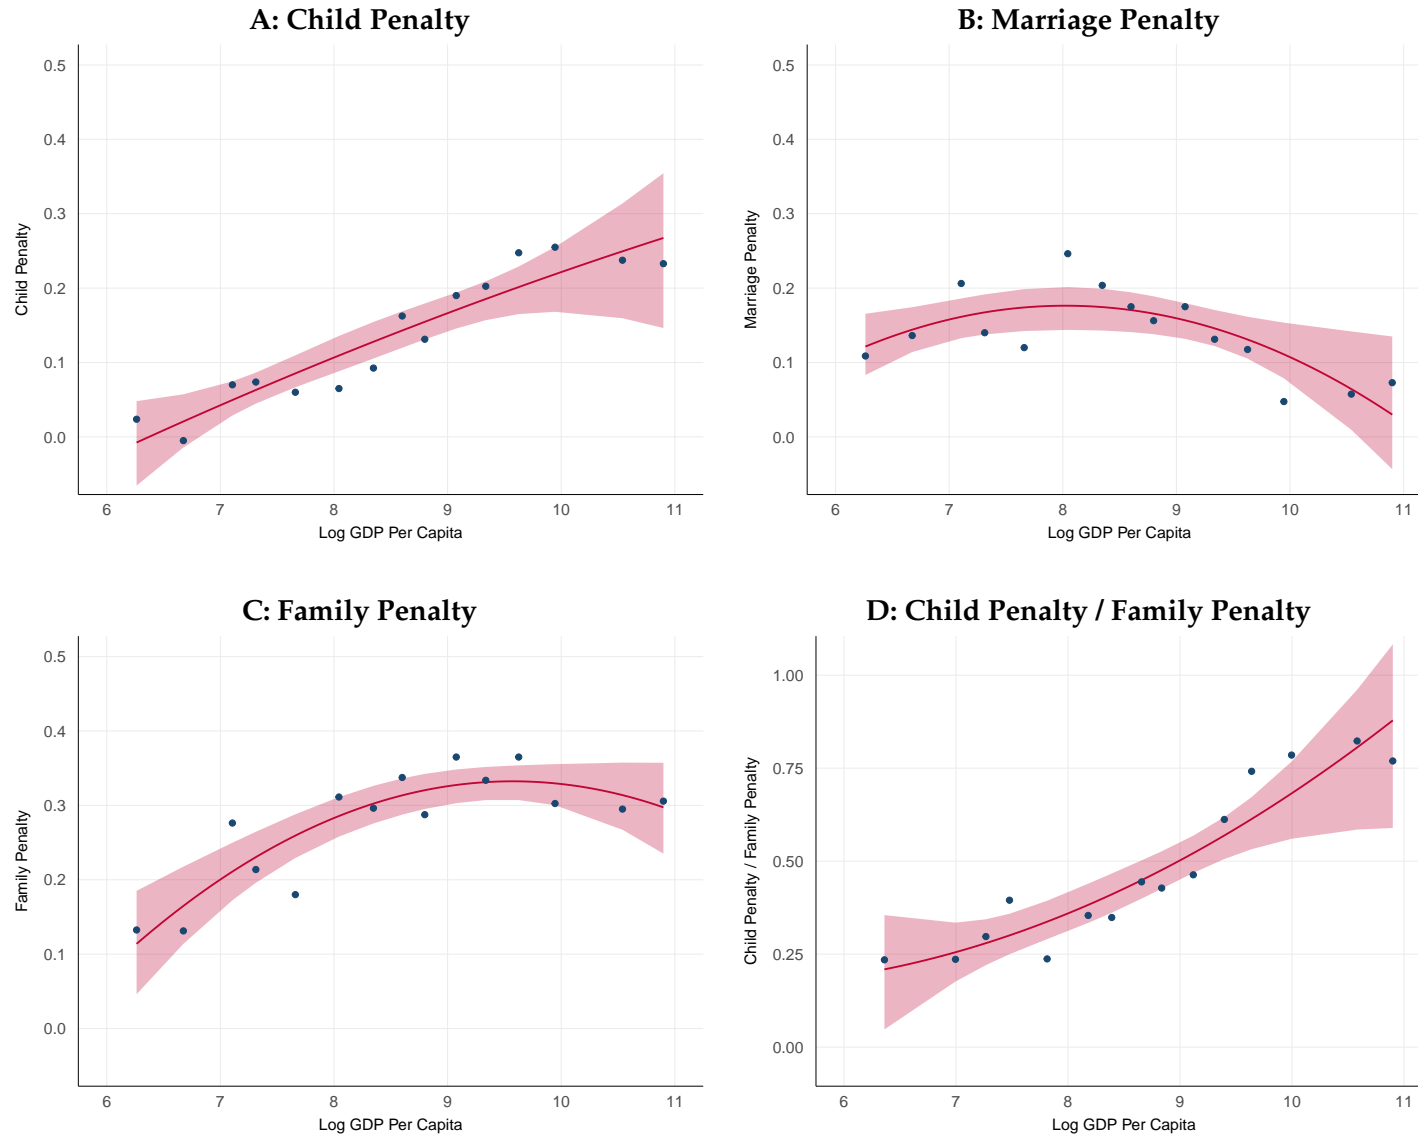
Notes: This figure presents event studies of first marriage in Brazil and China. Each panel shows an event study of the impact of marriage, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact of marriage at each 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.

FIGURE 15: 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 (gray 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.

FIGURE 16: SUBSTITUTION BETWEEN MARRIAGE AND CHILD PENALTIES OVER THE COURSE OF DEVELOPMENT



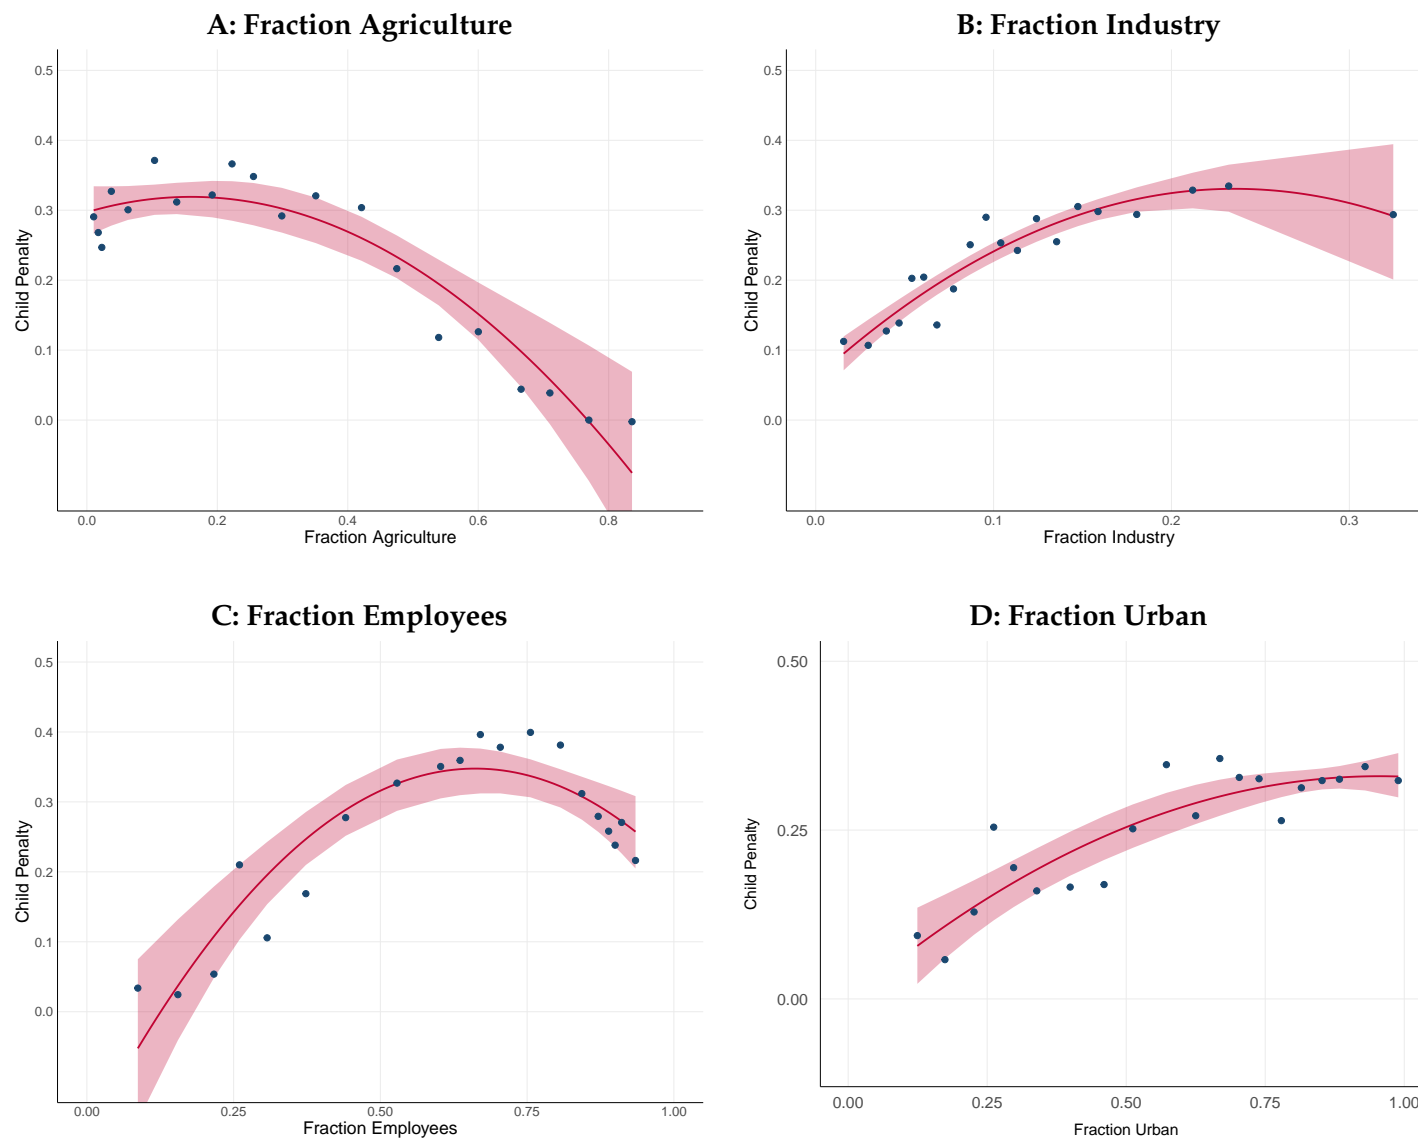
Notes: This figure shows binscatters of child penalties (Panel A), marriage penalties (Panel B), family penalties (Panel C), and child-to-family penalty ratios (Panel D) by GDP per capita. These are unscaled penalties in the employment rate obtained from specification (6). The family penalty is defined 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 World Development Indicators for the year 2019.

FIGURE 17: 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 (gray). 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 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 — 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 World Development Indicators for the year 2019.

FIGURE 18: CHILD PENALTIES AND STRUCTURAL TRANSFORMATION



Notes: This figure shows binscatters of child penalties against the fraction of the labor force in agriculture (Panel A), the fraction of the labor force in industry (Panel B), the fraction of the labor 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.

Online Appendix

A Supplementary Figures and Tables

TABLE A.1: DATA SUMMARY

Country	Source	Years	No. Waves	Employment Measure	Maternity Leave		No. Obs.
					Excluded	Max Days	
Afghanistan	DHS	2015	1	Weekly	Yes	60	28,709
Albania	DHS	2008-2017	2	Weekly	No	330	15,392
Algeria	MICS	2012-2019	4	Weekly	No	91	91,816
Angola	Census	2014	1	Weekly	Yes	56	541,452
Argentina	IPUMS	1970-2001	4	Weekly	Yes	70	3,391,576
Armenia	IPUMS	2001-2011	2	Weekly	No	1,095	189,992
Australia	Panel Data	2001-2019	19	Weekly	No	143	387,695
Austria	Panel Data	1981-2017	38	Annual	No	730	14,104,688
Bangladesh	IPUMS	1991-2011	3	Weekly	No	84	7,250,559
Belarus	IPUMS	1999-2009	2	Weekly	Yes	1,095	660,960
Belgium	LIS	1985-2017	21	Annual	No	344	88,473
Benin	IPUMS	1979-2013	4	Weekly	No	56	536,792
Bolivia	IPUMS	1976-2001	3	Weekly	Yes	30	486,681
Botswana	IPUMS	1991-2011	3	Weekly	No	84	77,961
Brazil	IPUMS	1991-2010	3	Weekly	Yes	183	18,447,111
Bulgaria	SILC	2007-2020	14	Weekly	No	730	58,282
Burkina Faso	IPUMS	1996-2006	2	Weekly	No	70	577,352
Burundi	DHS	2010-2016	2	Weekly	Yes	84	22,320
Cambodia	IPUMS	1998-2008	2	Weekly	No	90	730,506
Cameroon	IPUMS	1976-2005	3	Weekly	No	70	662,598
Canada	IPUMS	2011	1	Weekly	Yes	590	292,388
Chad	DHS	1996-2014	3	Weekly	Yes	70	17,343
Chile	LIS	1990-2017	13	Weekly	No	167	839,202
China	Mini Census	2005	1	Weekly	Yes	152	811,326
Colombia	IPUMS	1973-2005	3	Weekly	Yes	91	2,130,657
Congo, Dem. Rep.	DHS	2007-2013	2	Weekly	Yes	56	23,862
Congo, Rep.	DHS	2005-2011	2	Weekly	No	105	14,376
Costa Rica	IPUMS	1973-2011	4	Weekly	Yes	91	368,480
Cote d'Ivoire	LIS	2002-2015	3	Weekly	No	98	38,799
Croatia	SILC	2010-2020	11	Weekly	No	669	56,232
Cuba	IPUMS	2002-2012	2	Weekly	Yes	365	637,368
Cyprus	SILC	2009-2020	11	Weekly	No	113	45,089
Czech Republic	LIS	1992-2016	7	Weekly	No	1,095	51,163
Denmark	Panel Data	1980-2017	38	Annual	Yes	551	20,729,380
Dominican Rep.	IPUMS	1981-2010	3	Weekly	No	84	594,864

TABLE A.1: DATA SUMMARY, CONTINUED

Country	Source	Years	No. Waves	Employment Measure	Maternity Leave		No. Obs.
					Excluded	Max Days	
Ecuador	IPUMS	1982-2010	4	Weekly	Yes	70	1,161,328
Egypt	IPUMS	1996-2006	2	Weekly	No	91	4,022,460
El Salvador	IPUMS	1992-2007	2	Weekly	Yes	84	267,580
Estonia	SILC	2009-2020	12	Weekly	No	517	48,276
Ethiopia	DHS	2000-2016	4	Weekly	Yes	60	49,732
Fiji	IPUMS	1976-2014	5	Weekly	No	84	97,640
Finland	SILC	2009-2020	12	Weekly	No	338	86,340
France	LFS	1990-2020	31	Weekly	Yes	1,095	1,113,179
Gabon	DHS	2000-2012	2	Weekly	No	56	9,908
Gambia	LFS	2010-2015	2	Weekly	Yes	42	22,336
Georgia	LFS	2020-2021	2	Weekly	Yes	183	41,424
Germany	LIS	1989-2005	35	Weekly	No	909	322,910
Ghana	IPUMS	2000-2010	2	Weekly	Yes	84	827,642
Greece	IPUMS	1981-2001	3	Weekly	No	246	1,065,021
Guatemala	IPUMS	1964-2002	5	Weekly	Yes	84	660,925
Guinea	IPUMS	1983-2014	3	Weekly	No	56	387,468
Guyana	LFS	2017-2021	5	Weekly	Yes	49	28,535
Haiti	IPUMS	1971-2003	2	Weekly	Yes	42	285,874
Honduras	IPUMS	1974-2001	3	Weekly	Yes	42	302,181
Hungary	IPUMS	1990-2011	3	Weekly	No	1,095	482,283
Iceland	SILC	2004-2018	15	Weekly	No	365	45,120
India	DHS	2005-2015	2	Weekly	Yes	42	238,228
Indonesia	IPUMS	1971-2010	4	Weekly	Yes	46	10,847,064
Iran	IPUMS	2006	1	Weekly	Yes	122	480,443
Iraq	IPUMS	1997	1	Weekly	No	32	383,254
Ireland	LIS	1994-2018	20	Weekly	No	252	74,060
Israel	LIS	1986-2013	15	Weekly	No	106	111,300
Italy	LIS	1986-2020	14	Weekly	No	478	117,180
Jamaica	IPUMS	1982-2001	3	Weekly	Yes	56	127,494
Japan	Panel Data	2004-2020	17	Weekly	Yes	426	58,667
Jordan	IPUMS	2004	1	Weekly	Yes	70	149,651
Kenya	IPUMS	1989-2009	3	Weekly	Yes	91	1,397,772
Kyrgyz Republic	DHS	1997-2012	2	Weekly	No	126	9,576
Laos	IPUMS	2005	1	Weekly	No	90	138,689
Latvia	SILC	2009-2020	12	Weekly	No	605	38,556

TABLE A.1: DATA SUMMARY, CONTINUED


Country	Source	Years	No. Waves	Employment Measure	Maternity Leave		No. Obs.
					Excluded	Max Days	
Lesotho	IPUMS	1996-2006	2	Weekly	No	84	78,768
Liberia	IPUMS	2008	1	Weekly	No	98	63,017
Lithuania	LIS	2009-2018	10	Annual	No	730	28,470
Luxembourg	LIS	1985-2013	9	Weekly	No	450	28,467
Madagascar	DHS	1992-2008	4	Weekly	Yes	98	27,120
Malawi	IPUMS	1987-2008	3	Weekly	Yes	56	794,223
Malaysia	IPUMS	1991-2000	2	Weekly	Yes	60	230,466
Maldives	DHS	2009-2016	2	Weekly	No	60	15,456
Mali	IPUMS	1987-2009	3	Weekly	No	56	682,800
Mauritius	IPUMS	1990-2011	3	Weekly	Yes	84	123,057
Mexico	IPUMS	1970-2015	6	Weekly	Yes	70	11,766,780
Moldova	DHS	2005	1	Weekly	No	1,095	4,949
Mongolia	IPUMS	2000	1	Weekly	No	120	80,348
Morocco	IPUMS	1982-2004	3	Weekly	No	49	1,420,896
Mozambique	IPUMS	1997-2007	2	Weekly	Yes	60	800,136
Myanmar	IPUMS	2014	1	Weekly	No	42	1,466,036
Namibia	DHS	1992-2013	4	Weekly	Yes	84	19,460
Nepal	IPUMS	2001-2011	2	Weekly	No	52	1,298,668
Netherlands	LIS	1990-2018	11	Annual	No	85	86,317
Nicaragua	IPUMS	1995-2005	2	Weekly	Yes	56	202,638
Niger	DHS	1992-2012	4	Weekly	Yes	56	27,392
Nigeria	DHS	1990-2018	6	Weekly	Yes	84	102,960
Norway	Panel Data	1993-2017	25	Annual	No	395	8,176,350
Pakistan	LFS	2010-2021	11	Weekly	Yes	42	464,002
Panama	IPUMS	1960-2010	6	Weekly	Yes	56	311,922
Papua New Guinea	IPUMS	1980-2000	3	Weekly	Yes	0	292,932
Paraguay	IPUMS	1962-2002	5	Weekly	Yes	42	400,195
Peru	IPUMS	1993-2007	2	Weekly	Yes	90	1,324,900
Philippines	IPUMS	1990	1	Weekly	No	60	1,811,807
Poland	LIS	1992-2020	20	Weekly	No	365	682,580
Portugal	IPUMS	1981-2011	4	Weekly	No	61	740,056
Puerto Rico	IPUMS	1990-2010	4	Weekly	Yes	56	131,064
Romania	IPUMS	1992-2011	3	Weekly	No	730	2,076,342
Russia	IPUMS	2002-2010	2	Weekly	No	548	4,308,868
Rwanda	IPUMS	2002-2012	2	Weekly	No	84	598,204

TABLE A.1: DATA SUMMARY, CONTINUED

Country	Source	Years	No. Waves	Employment Measure	Maternity Leave		No. Obs.
					Excluded	Max Days	
Senegal	IPUMS	1988-2002	2	Weekly	No	98	234,000
Serbia	SILC	2013-2020	8	Weekly	No	341	51,688
Sierra Leone	IPUMS	2004	1	Weekly	No	84	81,262
Slovakia	LIS	1992-2018	10	Weekly	No	1,095	64,970
Slovenia	LIS	1997-2012	6	Weekly	No	341	24,306
South Africa	IPUMS	1996-2011	3	Weekly	No	0	2,268,054
South Korea	LIS	2006-2016	6	Weekly	No	456	74,316
South Sudan	IPUMS	2008	1	Weekly	Yes	90	88,697
Spain	IPUMS	1991-2001	2	Weekly	No	113	1,508,936
Sudan	IPUMS	2008	1	Weekly	Yes	42	1,027,125
Suriname	IPUMS	2012	1	Weekly	No	112	12,141
Sweden	Panel Data	1997-2017	21	Annual	Yes	484	30,606,597
Switzerland	Panel Data	1981-2020	40	Annual	No	100	32,503,240
Taiwan	LIS	1981-2016	11	Annual	No	NA	209,495
Tanzania	IPUMS	1988-2012	3	Weekly	No	84	2,263,515
Thailand	IPUMS	1990-2000	2	Weekly	No	90	337,138
Timor-Leste	DHS	2009-2016	2	Weekly	No	84	19,276
Togo	IPUMS	2010	1	Weekly	No	98	122,987
Trinidad & Tobago	IPUMS	1970-2011	4	Weekly	Yes	91	101,348
Tunisia	Census	2004	1	Weekly	Yes	183	485,077
Turkey	IPUMS	1985-2000	3	Weekly	Yes	91	2,405,829
Uganda	IPUMS	1991-2014	3	Weekly	No	60	1,681,773
United Kingdom	APS	2012-2020	9	Weekly	Yes	274	612,621
United States	CPS/ACS	1968-2020	53	Weekly	Yes	0	78,872,671
Uruguay	IPUMS	1963-2011	5	Weekly	Yes	192	444,810
Venezuela	IPUMS	1971-2001	4	Weekly	Yes	126	1,626,772
Vietnam	IPUMS	1989-2009	3	Weekly	Yes	122	6,910,548
Zambia	IPUMS	1990-2010	3	Weekly	Yes	84	630,762
Zimbabwe	IPUMS	2012	1	Weekly	No	77	166,126

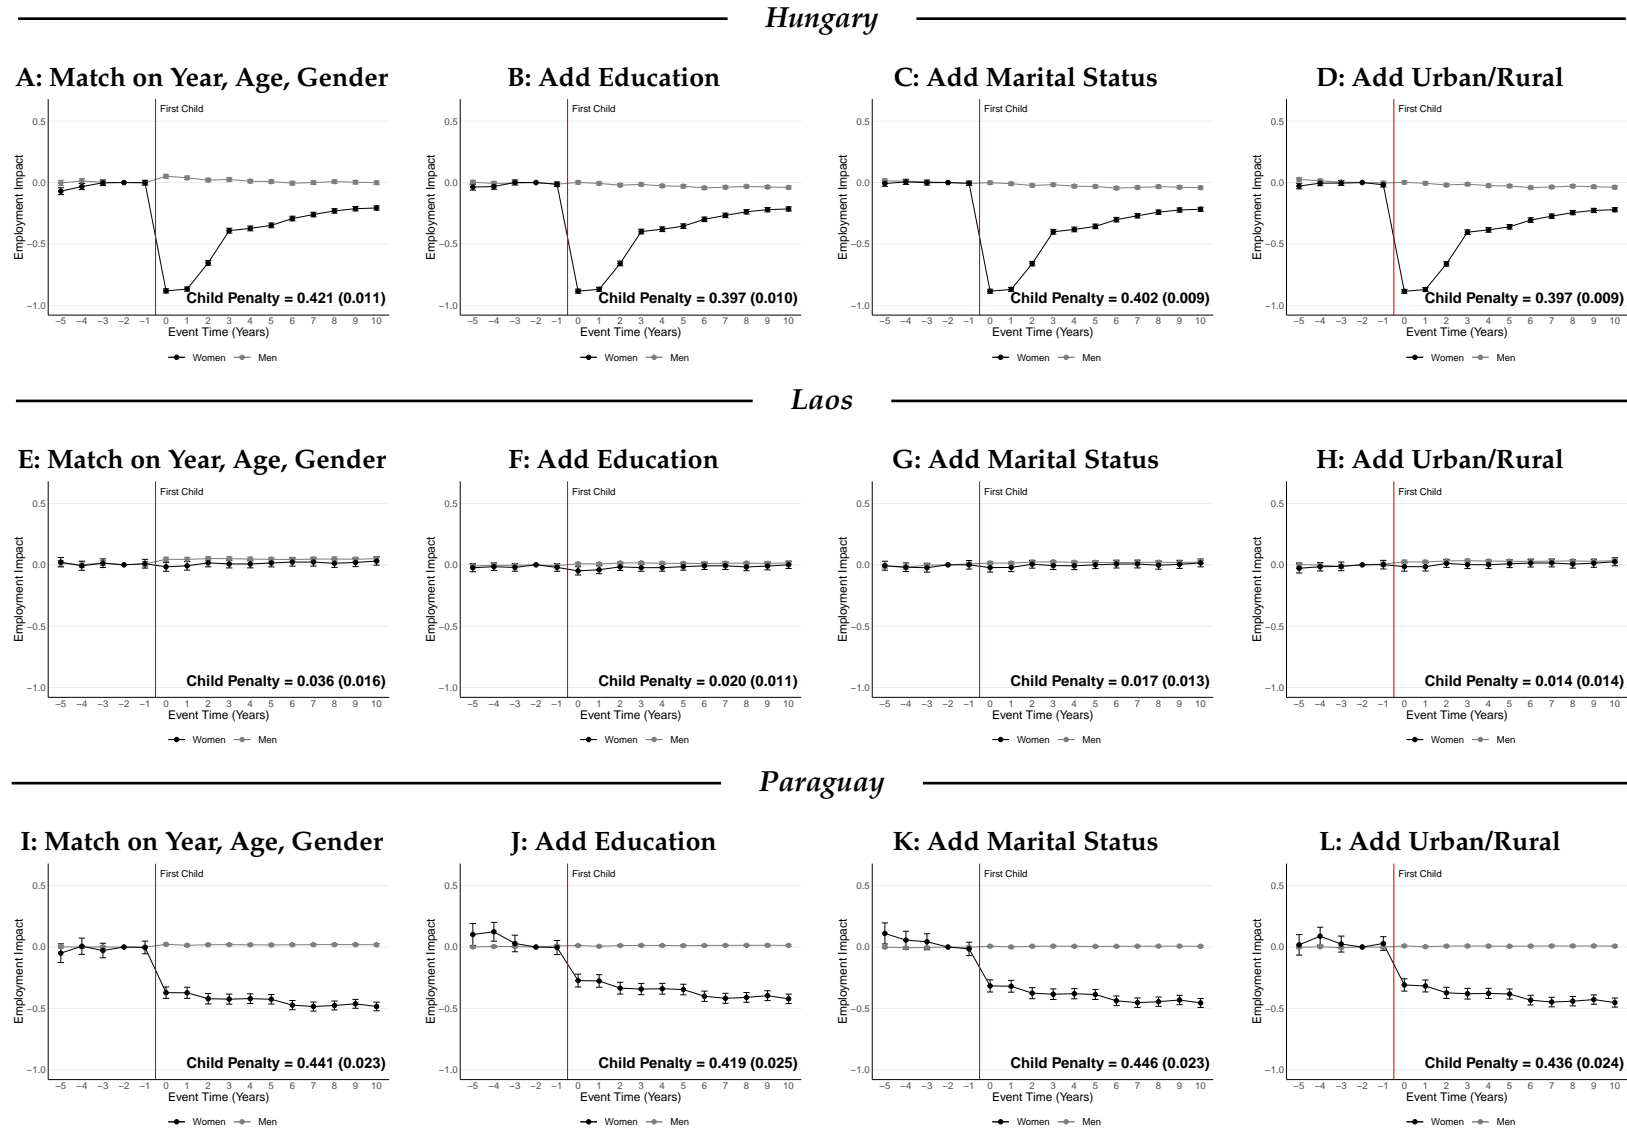
Notes: This table provides details on the data used for each of the 134 countries included in our analysis. The table lists the data source, the coverage period, the number of waves, the employment outcome (weekly vs annual), and number of individual×year observations for each country. The table also shows whether the employment outcome excludes or includes paid maternity leave and the maximum duration of paid maternity leave in each country.

FIGURE A.1: CHILD PENALTIES WHEN EXCLUDING VS INCLUDING PAID LEAVE IN EMPLOYMENT



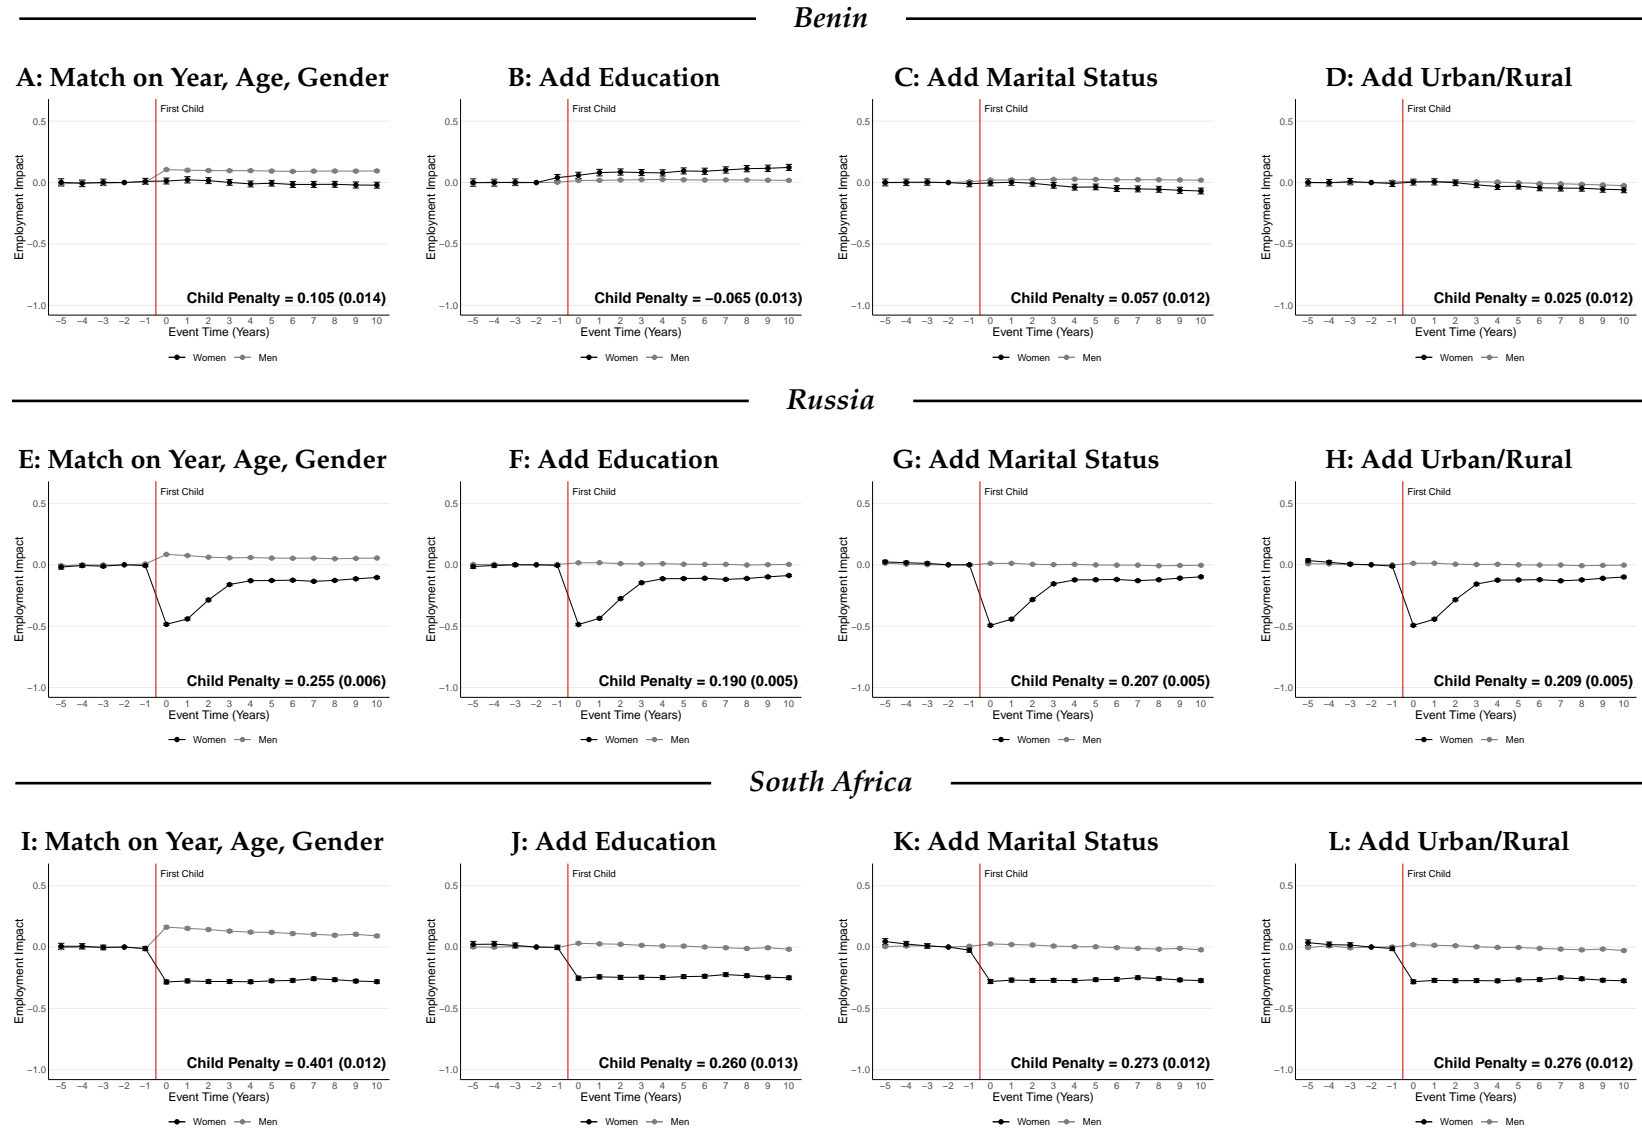
Notes: This figure compares child penalty estimates when excluding and including paid leave (maternal, parental, holiday, and sick leave) in the measure of employment. Only countries for which the data allow us to make this distinction are included in the figure. The y-axis shows child penalties when counting individuals on paid leave as employed, and the x-axis shows child penalties when counting individuals on paid leave as non-employed. Countries in light blue (dark blue) have a maximum duration of paid maternity/parental leave of less (more) than 180 days. The child penalty pairs lie very close to the 45-degree line, except for a few countries (less than five) with very long paid-leave duration. This suggests that the inability to exclude paid leave from the employment measure in some countries has a trivial effect on our results.

FIGURE A.2: PSEUDO-EVENT STUDIES UNDER DIFFERENT MATCHING SPECIFICATIONS
MATCHING SPECIFICATION MATTERS LESS



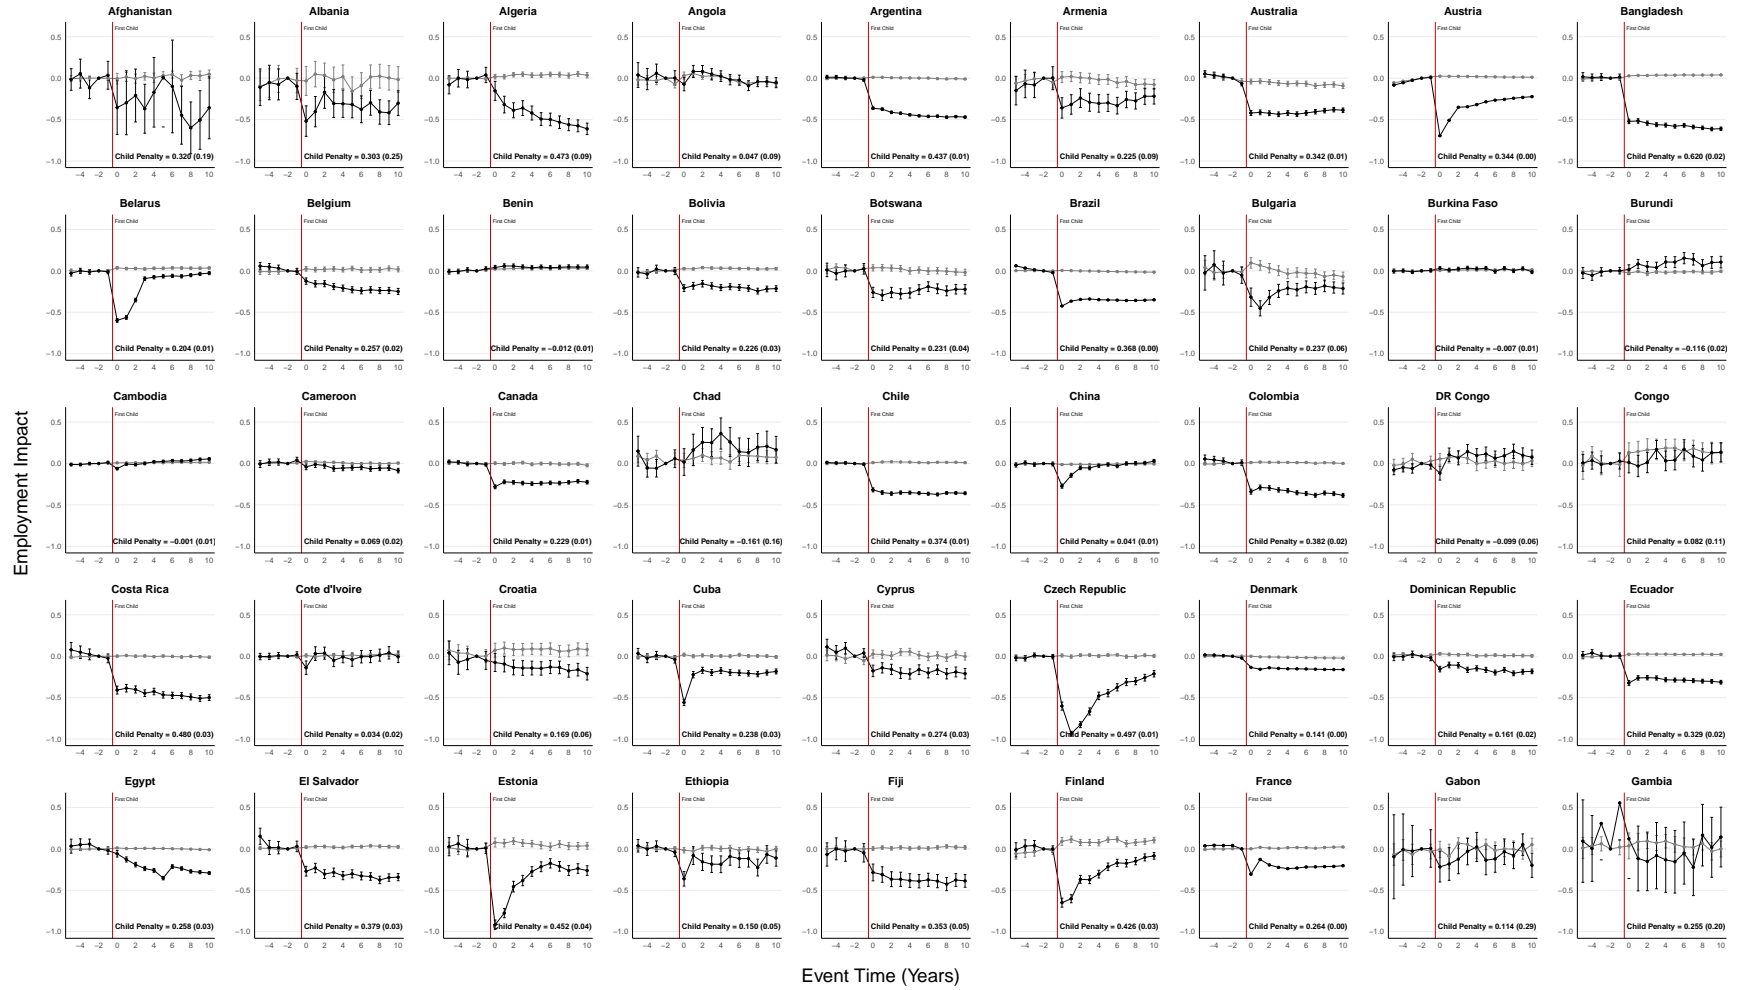
Notes: This figure presents pseudo-event studies of first child birth under different matching specifications. The figure provides examples of countries where the matching specification hardly matters for the results. Each row shows a given country, and each column shows a given matching specification. When moving 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). The results are extremely similar across the different specifications.

FIGURE A.3: PSEUDO-EVENT STUDIES UNDER DIFFERENT MATCHING SPECIFICATIONS
MATCHING SPECIFICATION MATTERS MORE



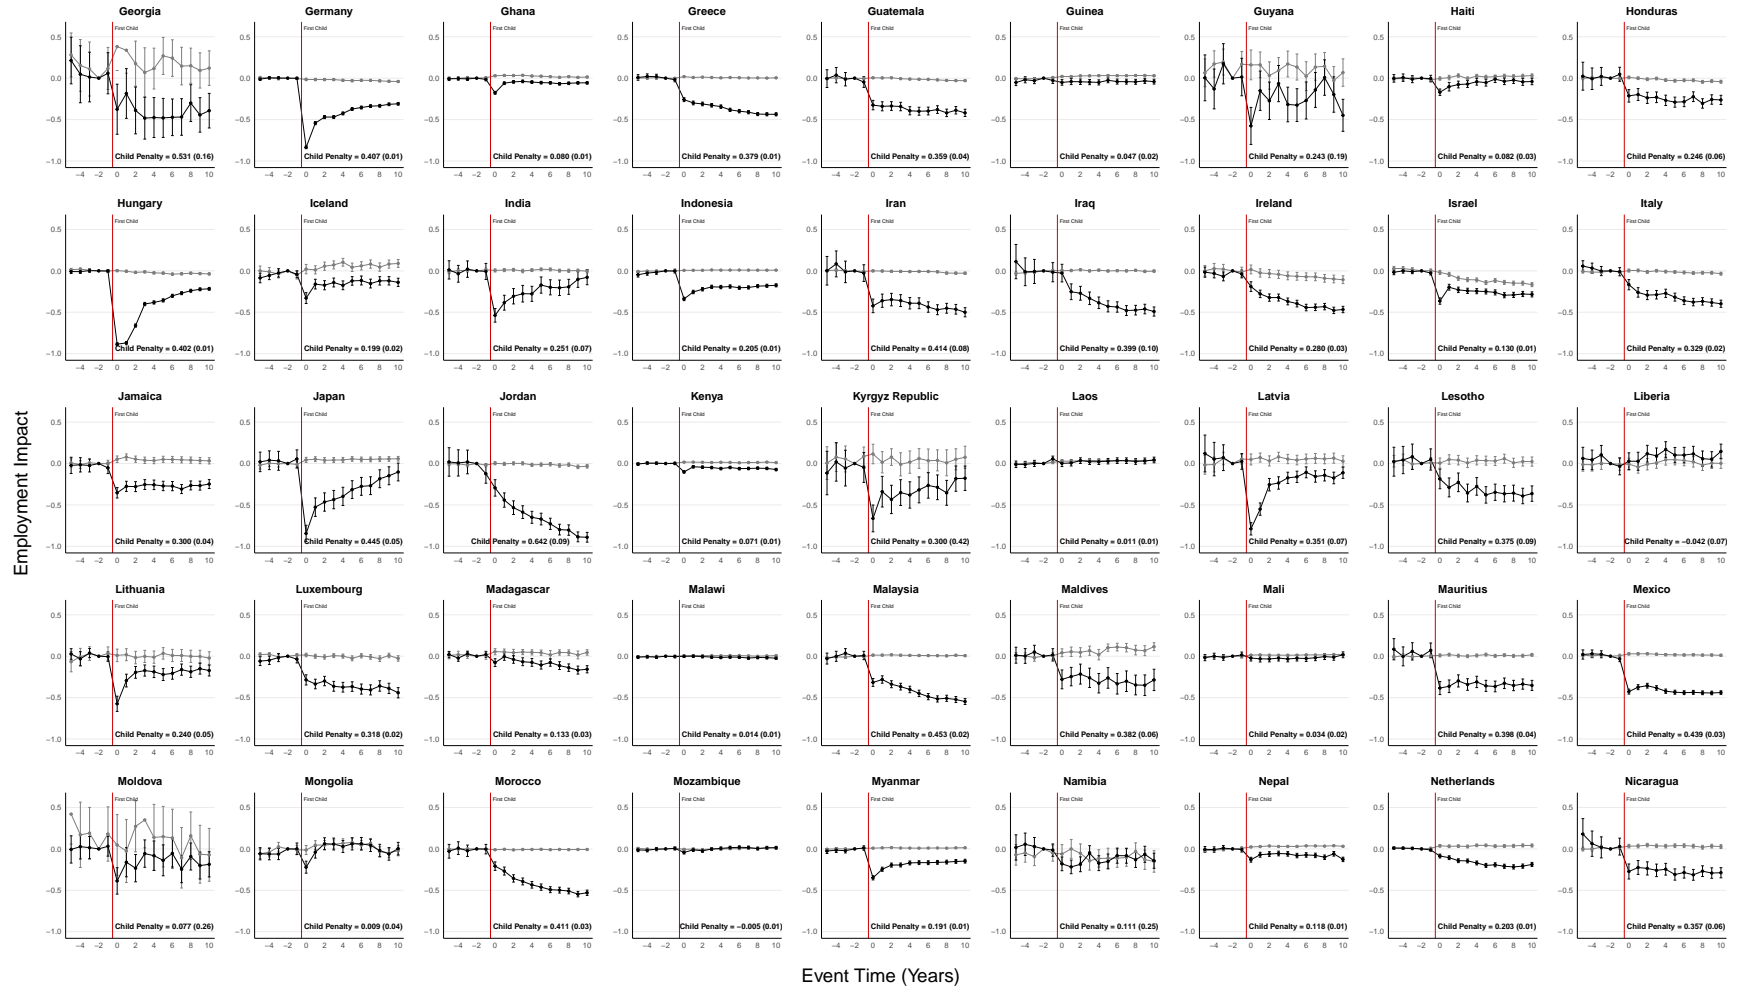
Notes: This figure presents pseudo-event studies of first child birth under different matching specifications. The figure provides examples of countries where the matching specification matters for the results (in contrast to the examples in the previous figure where the matching specification did not matter). Each row shows a given country, and each column shows a given matching specification. When moving 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). The more parsimonious specifications are associated with a positive jump in male employment at event time 0 — an indication of positive selection — while the richer specifications eliminate this jump.

FIGURE A.4: EVENT STUDIES OF FIRST CHILD BIRTH BY COUNTRY



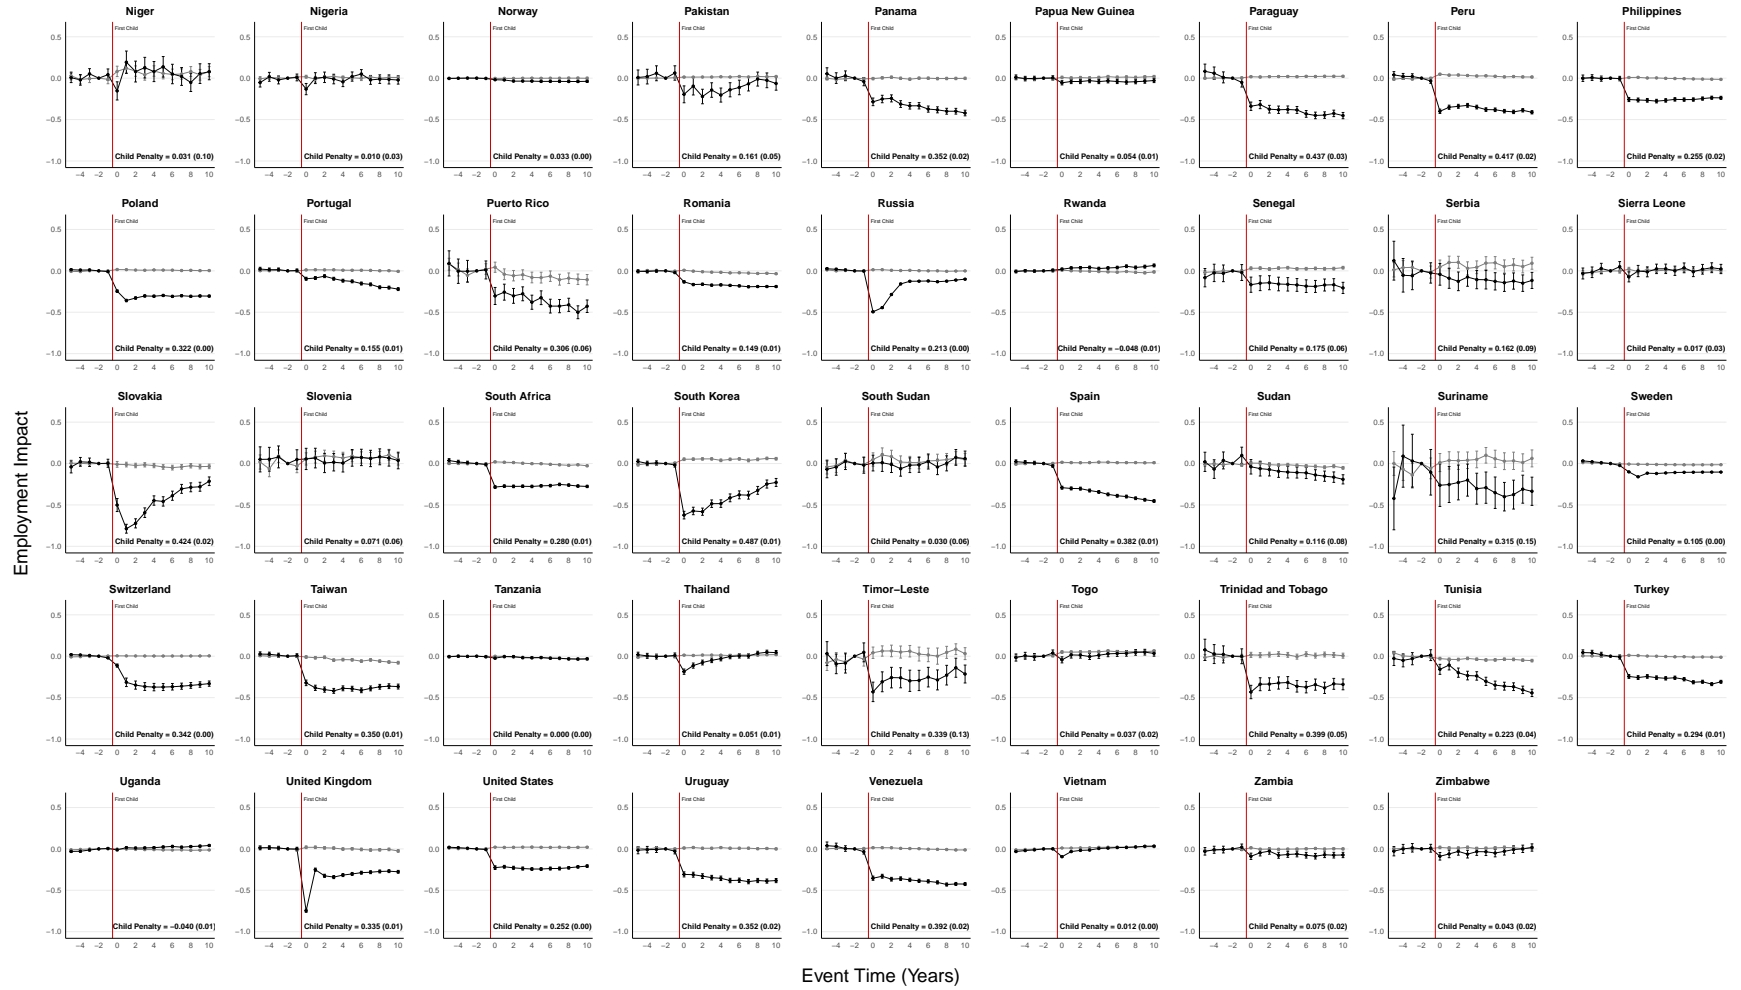
Notes: This figure presents event studies of first child birth for all 134 countries in our dataset. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact 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 A.4: EVENT STUDIES OF FIRST CHILD BIRTH BY COUNTRY, CONTINUED



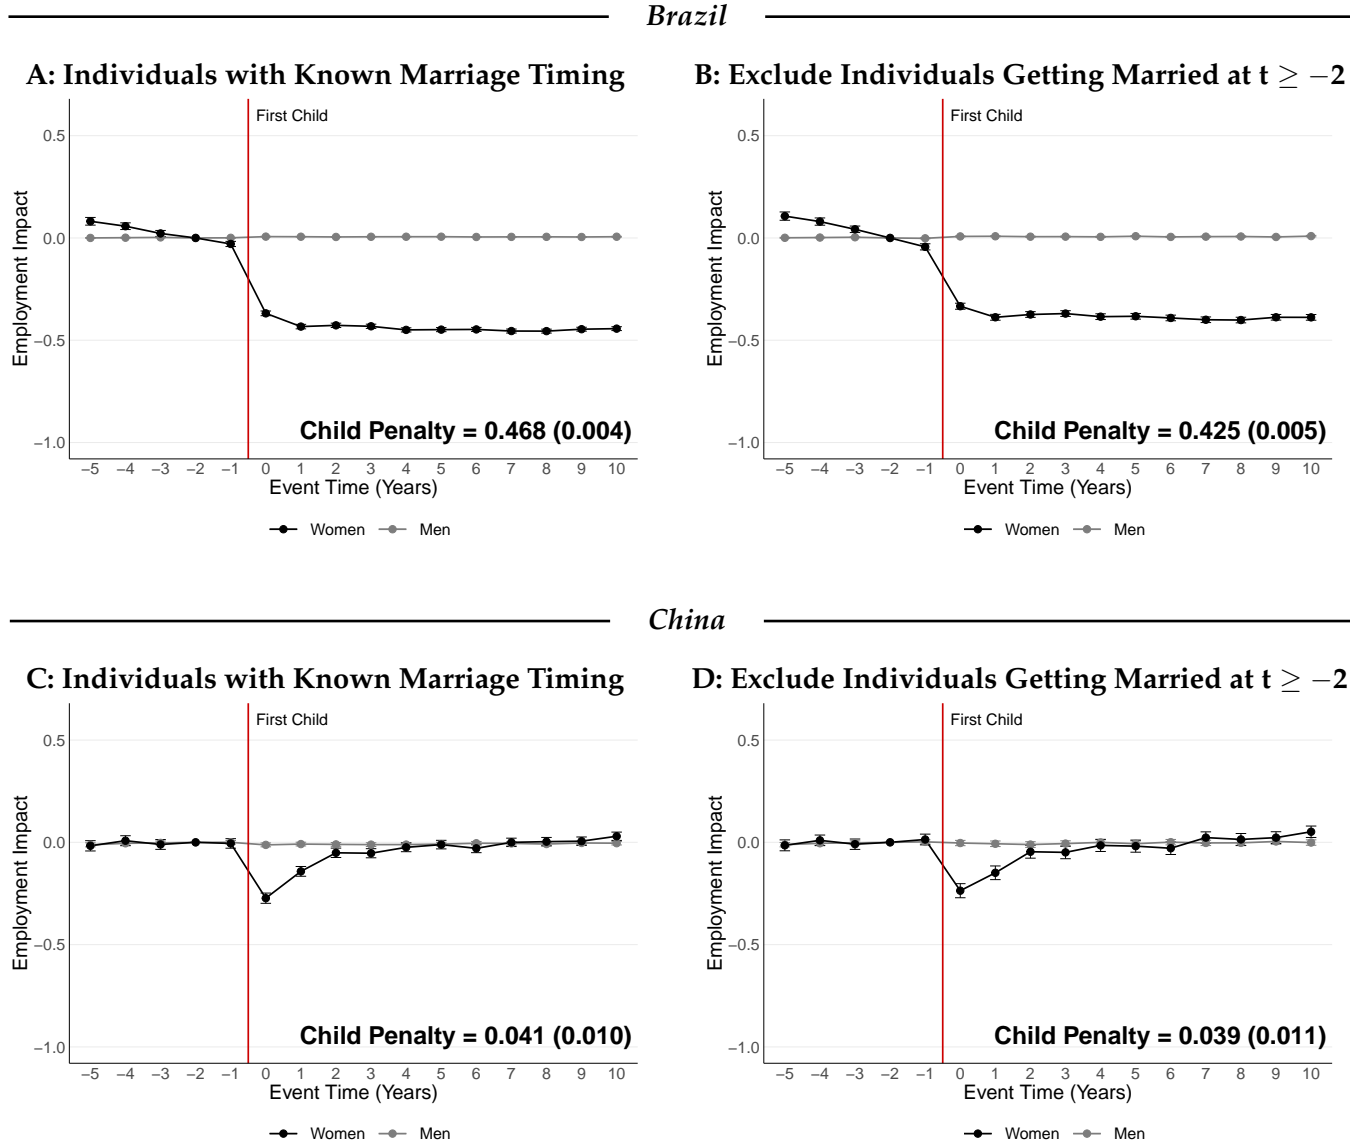
Notes: This figure presents event studies of first child birth for all 134 countries in our dataset. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray series) and women (black series). The series depict the percentage impact 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 A.4: EVENT STUDIES OF FIRST CHILD BIRTH BY COUNTRY, CONTINUED



Notes: This figure presents event studies of first child birth for all 134 countries in our dataset. Each panel shows an event study of the impact of first child, depicted by the vertical line, on the employment rate of men (gray 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 A.5: ARE CHILD PENALTIES ROBUST TO MARRIAGE EFFECTS?
CHILD PENALTIES BY MARRIAGE TIMING IN BRAZIL AND CHINA



Notes: This figure presents event studies of first child by marriage timing in Brazil and China. The left panels consider the full sample of individuals for whom marriage timing is observed, while the right panels consider the subsample of individuals who got married at least three years before their first child. Each panel shows an event study of the impact of first child on the employment rate of men (gray 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 error bars depict 95% confidence intervals based on robust standard errors.