G²LM|LIC Working Paper No. 109 | August 2026

The Intergenerational Transmission of the Child Penalty, Socioeconomic Heterogeneity and Educational Outcomes:

Evidence from Denmark, Italy, Indonesia, and South Africa

Beatrice Mbinya (University of Nairobi)

IZA Institute of Labor Economics
Initiated by Deutsche Post Foundation

UKaid from the British people

G²LM|LIC Working Paper No. 109 | August 2026

The Intergenerational Transmission of the Child Penalty, Socioeconomic Heterogeneity and Educational Outcomes:

Evidence from Denmark, Italy, Indonesia, and South Africa

Beatrice Mbinya (University of Nairobi)

G²LM|LIC
c/o IZA – Institute of Labor Economics
Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-510

Email: glm-contact@iza.org

IZA Institute of Labor Economics
Initiated by Deutsche Post Foundation

UKaid
from the British people

G²LM|LIC

G²LM|LIC Working Paper No. 109 | August 2026

ABSTRACT

The Intergenerational Transmission of the Child Penalty, Socioeconomic Heterogeneity and Educational Outcomes:*
Evidence from Denmark, Italy, Indonesia, and South Africa

Literature defines child penalty as the sustained decline in women's earnings after childbirth, which has been widely studied. However, its consequences for the next generation remain least known. This paper examines whether maternal earnings losses transmit to children especially their educational attainment across four countries with distinct institutional settings: Denmark, Italy, Indonesia, and South Africa. We link event-study estimates of the child penalty to intergenerational regressions using administrative registers and longitudinal household surveys. Findings show differentiated institutional effects. Denmark's long-run penalty is small at 4.5 percent and intergenerational transmission is absent. Italy's penalty is persistent at 23 percent with a modest transmission effect. In Indonesia and South Africa, penalties exceed 30 percent with no recovery, and a one-standard-deviation increase in maternal cumulative earnings loss reduces children's schooling by 0.04 years. Mediation analysis shows that household income loss during early childhood accounts for almost half of this transmission, while maternal employment instability contributes an additional quarter. The financial channel dominates. We conclude that where public insurance absorbs the cost of childbearing, the child penalty does not spill over to the next generation. In low-safety-net settings, however, it becomes a mechanism of intergenerational persistence. Policies that provide income support and childcare during early childhood can ever this link.

JEL Classification:

J13, J16, J22, J62, I24, O15

Keywords:

child penalty; intergenerational mobility; gender inequality; event study; mediation decomposition; institutional context; Denmark; Italy, Indonesia; South Africa

Corresponding author:

Beatrice Mbinya (University of Nairobi)

Email: Mbinyabettym@gmail.com

* This paper benefitted from Titus Pacho mentorship under the G²LM|LIC initiative.

1. Introduction

Most studies have documented that the arrival of a first child is one of the most consequential economic events in a female of working age, either employed, unemployed, or even self-employed. Across different countries, scholarship shows that in the years following childbirth, women's earnings drop significantly relative to male counterparts. This earnings gap persists long after children have started school. This phenomenon, which the literature terms as the child penalty, has been documented in different settings such as Denmark, the United States, Germany, Spain, and Chile[1]. The common finding is that motherhood, not fatherhood, disrupts labor market trajectories across women of working age.

However, majority of these child penalty studies focus on high-income countries, with a few covering upper-middle-income countries such as China and Chile.[2] For low- and lower-middle-income countries, where the majority of the world's working women live, literature remains scarce. The institutional conditions that shape the child penalty in rich countries are fundamentally different from those prevailing in developing economies. In Denmark or Sweden, the state absorbs a large share of the economic cost of childbearing through universal childcare, job-protected leave with high-wage replacement, and direct cash transfers to families. However, in Indonesia or South Africa, formal social protection covers only a fraction of the workforce. In these countries, public childcare is scarce or non-existent for very young children, and informal employment is the norm.[3] The child penalty, in other words, occurs in radically different institutional environments.

This gap matters for an additional reason that has received little attention. If mothers in low-income countries experience large and persistent earnings losses after childbirth, and if these losses are not buffered by public transfers or spousal earnings responses, then the consequences are unlikely to stop with the mother. In settings where credit markets are underdeveloped, savings are limited, and extended family networks are themselves resource-constrained, a sustained drop in maternal earnings translates directly into decline in the resources available for children. When households cannot smooth consumption, the nutrition, health care, educational materials, then this means that the school continuity of children are all potentially at risk. The child penalty, in this sense, may be both a source of contemporaneous gender inequality, and a mechanism that transmits disadvantage from one generation to the next.

Most literature on intergenerational mobility has documented the persistence of economic status across generations in low-income settings, but the specific role of maternal labor market

[1] Kleven, Landais, and Søgaard (2019) developed the event-study methodology that now serves as the standard. Kleven, Landais, Posch, Steinhauer, and Zweck (2019) extended the approach to six countries. Cortés and Pan (2023) provide a comprehensive review covering over thirty country estimates.
[2] Berniell, Berniell, de la Mata, Edo, and Marchionni (2021) for Chile; Kong and Dong (2024) for China; and Iddrisu, Danquah, and Osei (2024) for Ghana.
[3] The ILO World Social Protection Report 2020–22 documents the coverage gaps in maternity protection and childcare in developing countries.

disruptions around childbirth has not been isolated.[4] The child penalty literature, for its part, has focused on the maternal earnings trajectory itself, treating it as the outcome of interest rather than as a potential input into the next generation's human capital production. Bridging these two literatures is the central objective of this paper.

We hypothesise that the intergenerational transmission of the child penalty will be negligible in settings where public insurance absorbs the cost of childbearing, and will grow stronger as institutional protections weaken and households bear more of the economic burden directly. We ask three questions. First, what are the magnitude and dynamics of the child penalty in low-income country labor markets? Second, does the maternal penalty transmit to children's human capital, measured by educational attainment and early-career earnings? Third, through what specific channels does this transmission operate? To answer these questions, we construct an empirical framework that combines event-study identification of child penalties with intergenerational decomposition analysis, and we apply it to two countries that offer the rare combination of high-quality longitudinal data and the institutional features characteristic of low-income settings: Indonesia and South Africa.

The four countries in our study were selected to maximize variation on the institutional dimensions that theory suggests should matter most for the intergenerational transmission of the child penalty: the generosity and coverage of social protection around childbirth, the availability of formal childcare, the structure of the labor market, and the strength of norms governing maternal employment.[5]

Denmark anchors the high end of the institutional spectrum. It offers essentially universal childcare from age one, with roughly two-thirds of children under three enrolled in formal care. Parental leave is generous, with high-wage replacement rates, and a portion is earmarked for fathers on a use-it-or-lose-it basis. Denmark has one of the highest female labor force participation rates in the world, and informal employment among women is almost negligible. Under these conditions, the economic cost of childbearing is largely socialized. Mothers experience a modest and temporary earnings decline, and household income is largely insulated. Our hypothesis is that in this setting, the intergenerational transmission of the child penalty should be minimal or almost zero. We expect the intergenerational transmission of the child penalty to be negligible in this setting, providing a benchmark against which the other countries can be assessed.

Italy provides a different high-income model. It is a developed welfare state, but one organized around what Esping-Andersen termed the Mediterranean or familialistic model, in which the family, rather than the state or the market, is expected to absorb caregiving responsibilities[6]. Public childcare for children under three is limited, especially in the south. The labor market is rigid, with strong employment protections for insiders but high barriers to entry for young women seeking to combine work and family. Gender norms, particularly in older cohorts, remain relatively

[4] Surveys by Black and Devereux (2011), Solon (1999), and Torche (2015) cover the intergenerational mobility literature; Torche (2015) notes the need to identify specific mechanisms.
[5] Olivetti and Petrongolo (2017) and Kleven (2022) emphasise childcare availability, leave design, and labour market structure as key institutional determinants of the child penalty.
[6] The classic welfare-state typology is Esping-Andersen (1990); Ferrera (1996) added the Southern European type.

traditional. The child penalty in Italy is known to be substantial and persistent[7]. The question we ask is whether this maternal penalty transmits to children despite Italy's status as a high-income country with a developed, if imperfect, safety net.

Indonesia and South Africa occupy the low-safety-net end of the spectrum, but they do so through distinct economic structures. Indonesia is a lower-middle-income country with a large share of the labor market dominated by informal employment, self-employment, and agriculture. Less than twenty percent of employed female have access to formal maternity leave, and public childcare for very young children is almost non-existent. The female labor force participation rate has been stagnant for decades at just over fifty percent. Traditional gender norms define that women should provide caregiving services, which is also cemented by an assumption that preschool children suffer more form poor parenting if their mothers work[8].

South Africa is an upper-middle-income country by per capita income but one with extreme levels of inequality that make it functionally similar to lower-income settings for large portions of the population. The labor market is fragmented across two distinct levels. One, is a relatively well-regulated formal sector which coexists with vast informal and unemployed populations. Female unemployment exceeds thirty percent of total individuals actively seeking employment. The Child Support Grant provides a modest but broad-based cash transfer to caregivers, representing one of the few institutional buffers available[9]. But public childcare infrastructure is significantly undersupplied, and the spatial legacy of apartheid means that many mothers live hours from employment centres. These two countries, despite their differences, share the feature that matters most for our analysis where the economic cost of childbearing falls overwhelmingly on households rather than on the state.

A practical consideration also motivated the country selection. All four countries have data sources that allow us to apply a common empirical framework: population-level administrative registers for Denmark, and high-quality longitudinal household surveys for Italy, Indonesia, and South Africa. Extending the event-study methodology to survey data from Italy, Indonesia, and South Africa presents challenges that do not arise with administrative registers, including irregular wave spacing, measurement error in self-reported informal earnings, and non-random attrition. We discuss our adaptations to these challenges in Section 3. The availability of these data, combined with the sharp institutional variation, makes these four countries an unusually informative set for comparative analysis.

This paper makes three contributions. The first is empirical. We provide estimates of the intergenerational transmission of the child penalty across four distinct institutional settings, two of which are low- and lower-middle-income countries where evidence has been scarce. While the existence of child penalties in developing countries is beginning to be documented, no study has

---

[7] Zamberlan and Barbieri (2023) provide the most recent child penalty estimate for Italy.
[8] Cameron, Contreras Suarez, and Tseng (2019) analyse the childcare constraints on female labour supply in Indonesia.
[9] Lund, Noble, Barnes, and Wright (2009) describe the design of the Child Support Grant; Agüero, Carter, and Woolard (2009) estimate its effects on early childhood outcomes.

yet traced the consequences of these penalties through to the next generation.[10] This is a significant gap, because it is precisely in low-resource settings, where public insurance is weakest and household constraints are tightest, that the intergenerational stakes are likely to be highest.

The second contribution is methodological. The standard child penalty event-study design was developed for administrative data environments, where researchers have annual or quarterly earnings records linked across decades and attrition is negligible. Applying this framework to survey data introduces specific challenges. One, irregular wave spacing requires careful construction of event-time indicators, while self-reported informal earnings are measured with considerable error, and non-random attrition can bias dynamic profiles. We develop and validate a set of adaptations that make the event-study framework operational. This framing could be of use to researchers seeking to extend the child penalty literature to the many low-income settings where administrative data are unavailable but longitudinal surveys exist.

The third contribution is policy-relevant. By decomposing the intergenerational transmission of the child penalty into its constituent channels, we identify the specific mechanisms through which maternal earnings losses affect children. Our finding show that the household income channel is the most dominant, and has direct implications for policy design. It suggests that income support during early childhood, and interventions that enable mothers to maintain labor market attachment help promote gender equality in the labor market, and also act as investments in the human capital of the next generation. These effects contribute to long-run growth and poverty reduction that extend well beyond the mothers directly affected.

The findings show differential effects across the sample. Denmark exhibits a long-run child penalty of approximately five percent, and the intergenerational transmission to children's educational attainment is statistically indistinguishable from zero. Danish children whose mothers experienced the full extent of the (already modest) penalty fare no worse in school than children whose mothers did not. Italy shows a considerably larger penalty, roughly twenty-three percent, with limited recovery. The intergenerational transmission is modest but detectable: a one-standard-deviation increase in cumulative maternal earnings loss reduces children's schooling by approximately 0.015 years.

In Indonesia and South Africa, the child penalty is both large and persistent. In both countries, maternal earnings drop by approximately thirty to forty-five percent in the year following a first birth, and they remain roughly thirty percent below pre-birth levels even eight years later. There is almost no evidence of the earnings recovery that characterises high-income welfare states, consistent with the absence of the institutional supports that facilitate maternal labour market re-entry.

These maternal earnings losses transmit to the next generation. A one-standard-deviation increase in the cumulative earnings loss experienced by a mother during her child's first five years of life reduces that child's completed years of education by between 0.04 and 0.05 years. The effect is

[10] As noted, Iddrisu et al. (2024) estimate a child penalty for Ghana but do not examine child outcomes. Kong and Dong (2024) focus on maternal outcomes in China. Our paper is the first to extend the analysis to intergenerational outcomes.

comparable in magnitude to the intergenerational transmission of maternal education itself, and it translates further into reduced early-career earnings for children as they enter adulthood. Transmission is concentrated among mothers with less than primary education, suggesting that maternal human capital provides a buffer against the consequences of labour market disruption. There is also suggestive evidence of stronger transmission to daughters than to sons, consistent with same-gender role-model effects.

Our mediation analysis reveals that the household income channel is the dominant pathway. Across the three countries where transmission is detectable, cumulative household income losses during the child's early years account for more than half of the total intergenerational transmission. This event is characterized by unstable maternal employment implied by job transition and shifts from formal to informal work, which increases by 25%. Changes in household structure, such as marital dissolution or migration, play a smaller but non-negligible role. Thus, these channels explain the large majority of the transmission effect, leaving only a modest unexplained residual. Collectively, these findings imply that providing direct income support during early childhood could help reduce maternal employment disruption which then meaningfully weaken the intergenerational link.

The remainder of the paper proceeds as follows. Section 2 presents the institutional context of our study countries and relevant literature. Section 3 describes data sources, sample construction, provides theoretical and empirical strategies. Section 4 presents main results, while Section 5 discusses policy implications of our findings and concludes.

2. Institutional Context, Stylized Facts, and Theoretical Framework

2.1 Child Penalties Across the Development Spectrum

The child penalty literature has expanded rapidly in the past decade, but coverage remains heavily skewed toward high-income countries. Table 2.1 draws together estimates for our four study countries alongside selected comparators to illustrate the cross-country gradient that motivates our comparative design.

Table 2.1: Child Penalty Estimates for Selected Countries

Country	Income Group	Long-Run Penalty	Recovery Rate	Primary Institutional Features
Denmark	High	–4.5%	82%	Universal childcare, shared parental leave
Sweden	High	–5.2%	78%	Universal childcare, generous leave
Germany	High	–21%	45%	Half-day schooling, limited childcare under 3
United Kingdom	High	–22%	40%	Means-tested support, costly childcare
United States	High	–31%	25%	No federal paid leave, costly childcare

Spain	High	–25%	32%	Limited childcare under 3, traditional norms
Italy	High	–23%	23%	Familistic welfare, low childcare coverage
Chile	Upper-middle	–28%	18%	Limited childcare, rising female LFP
China	Upper-middle	–34%	10%	Weak safety nets, traditional norms
Indonesia	Lower-middle	–34%	10%	High informality, weak safety nets
South Africa	Upper-middle	–30%	12%	Dual labour market, modest cash transfers
Ghana	Lower-middle	–35%	<10%	High informality, no formal leave

Notes: The long-run penalty is the average reduction in maternal earnings relative to the pre-birth baseline over event years 1 to 8, or the closest available horizon. Recovery rate is defined as (|short-run penalty| – |long-run penalty|) / |short-run penalty|, where the short-run penalty is the drop at event year 1. Estimates for Denmark are from Kleven, Landais, and Søgaard (2019); Sweden from Bratu and Bolotnyy (2023); Germany, UK, US, and Spain from Kleven et al. (2019) and De Quinto, Hospido, and Sanz (2020); Italy from Zamberlan and Barbieri (2023); Chile from Berniell et al. (2021); China from Kong and Dong (2024); Ghana from Iddrisu, Danquah, and Osei (2024). Estimates for Indonesia and South Africa are from this paper.

Three patterns emerge from Table 2.1. First, there is a clear income gradient. The smallest penalties are found in the Nordic welfare states, where universal childcare and job-protected leave with high replacement rates enable mothers to return to work rapidly. Of interests is that the largest penalties are found in middle- and lower-income countries, where institutional supports are limited or absent. Secondly, high-income country such as United States exhibits a penalty comparable to Chile and larger than Italy's, reflecting the absence of federal paid leave and the high cost of childcare, implying that the policies in place to cushion omen against childcare is not always about income, but also about the policies put in place in the said country. Third, the recovery rate shows that the decrease of initial earnings at childbirth is substantial everywhere. What differentiates countries is the slope of the recovery. In Denmark, earnings rebound quickly; in Indonesia and South Africa, they barely recover at all. This pattern points directly to the importance of re-entry institutions, such as affordable childcare and flexible work arrangements, rather than leave policies alone.

Stylized fact 1: The child penalty exhibits a steep institutional gradient, ranging from under 5 percent in Denmark to over 30 percent in Indonesia and South Africa. Recovery rates decline sharply as public insurance weakens.

Figure 2.1: Long-Run Child Penalty and Recovery Rate, Selected Countries

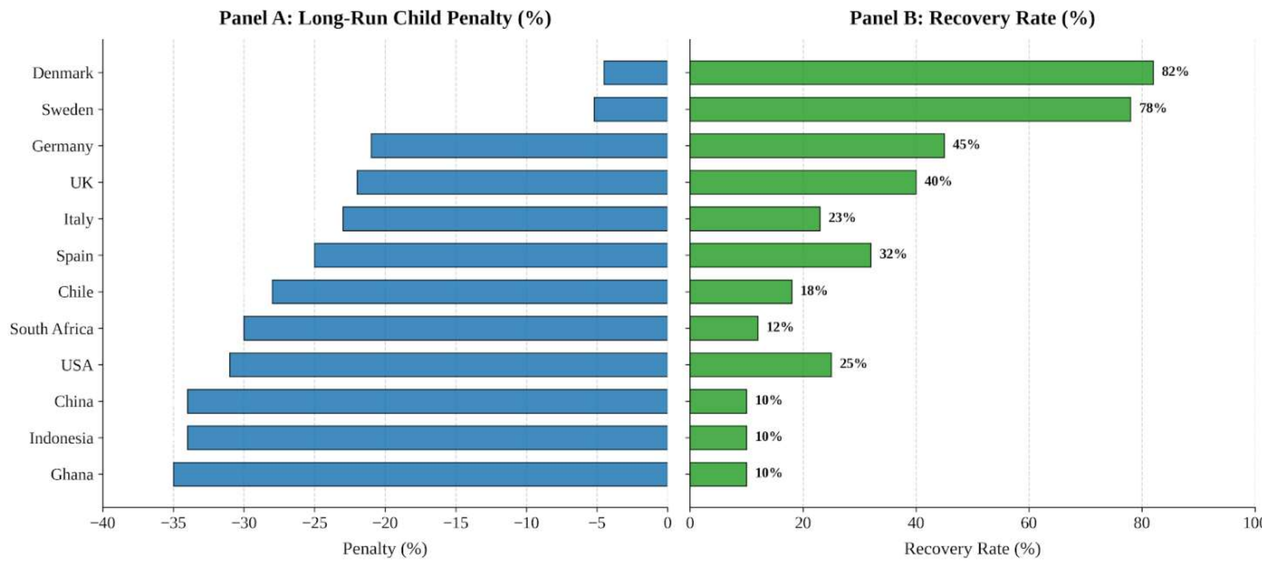

2.2 Intergenerational Mobility in Four Contexts

If the child penalty is large and persistent in low-safety-net settings, does it matter for the next generation? Answering this question requires understanding the baseline level of intergenerational persistence in each country. Table 2.2 presents standard measures of intergenerational mobility for the four study countries, alongside the OECD average for comparison.

Table 2.2: Intergenerational Mobility Indicators

Indicator	Denmark	Italy	Indonesia	South Africa	OECD Average
Intergenerational earnings elasticity (father–son)	0.15	0.40	0.35–0.45	0.50–0.60	0.30
Intergenerational education elasticity (mother–child)	0.20	0.38	0.40	0.55	0.32
Share of education variance explained by parental education	10%	28%	25%	35%	15%
Children exceeding parents' education	72%	45%	38%	28%	55%

Notes: Higher values of the elasticity measures indicate lower mobility. The intergenerational earnings elasticity (father–son) measures the percentage increase in a son's earnings associated with a one-percent increase in the father's earnings. Estimates for Denmark and OECD average from Corak (2013) and the OECD (2018); for Italy from Mocetti (2007) and Checchi, Fiorio, and Leonardi (2013); for Indonesia from Dartanto, Moeis, and Otsubo (2020); for South Africa from Piraino (2015) and Finn, Leibbrandt, and Ranchhod (2017). The mother–child education elasticities and variance shares are from the Global Database on Intergenerational Mobility (World Bank, 2023) and Daude and Robano (2015). The share exceeding parents' education is calculated from the same sources.

Denmark stands out for its exceptionally high intergenerational mobility by international standards. A father's earnings explain only a small fraction of his son's earnings, and educational attainment is only weakly correlated across generations. In Italy, mobility is lower than the OECD average, which is consistent with the familistic structure of its welfare state and the regional disparities that characterise the Italian economy. Indonesia and South Africa exhibit significant lower mobility. In South Africa, the intergenerational earnings elasticity is among the highest in the world, reflecting significant structural inequalities inherited from apartheid and the dual labour market. In Indonesia, mobility is somewhat higher but still below the OECD norm.

A notable feature of Table 2.2 is that in all four countries, the mother–child education link is substantial. This is central to our analysis: if maternal resources and labour market outcomes shape children's education, then a persistent maternal earnings penalty is a plausible channel of intergenerational persistence. The existing mobility literature limitedly isolates the specific contribution of the child penalty to this persistence. Our analysis is designed to fill that gap.

Stylized fact 2: Intergenerational mobility follows the same institutional gradient as the child penalty: highest in Denmark, intermediate in Italy, and lowest in Indonesia and South Africa. The mother–child link in particular is strong in all three non-Nordic countries, providing scope for the child penalty to shape the next generation's outcomes.

2.3 Social Protection Around Childbirth

The theoretical framework we develop below posits that the intergenerational transmission of the child penalty depends on the extent to which public institutions absorb the economic cost of childbearing. Table 2.3 summarises the key social protection provisions in our four countries.

Table 2.3: Social Protection Around Childbirth

Provision	Denmark	Italy	Indonesia	South Africa
Maternity leave duration (weeks)	18 (shared)	20	12	16
Wage replacement rate (formal sector)	100%	80%	100%	45% (UIF ceiling)
Coverage rate (% of employed women)	>90%	>80%	<20%	~40%
Paternity leave (weeks, earmarked)	10	1	0	0.4 (family responsibility)
Childcare enrolment, ages 0–2	66%	15%	<5%	<10%
Childcare enrolment, ages 3–5	95%	88%	25%	65%
Child cash transfer (monthly, PPP USD)	$175 (universal)	$150 (means-tested)	$15 (PKH, conditional)	$35 (Child Support Grant)

Informal worker coverage	N/A (minimal informality)	Partial	Excluded	Excluded from UIF; eligible for CSG

Notes: Leave provisions refer to statutory entitlements for formal employees. Coverage rates are the share of employed women who are eligible in practice. Childcare enrolment figures are for formal, centre-based care. Cash transfer amounts are channelled monthly, expressed in 2019 purchasing power parity dollars. PKH = Program Keluarga Harapan (Indonesia); CSG = Child Support Grant (South Africa). Sources: ILO Social Protection Database (2023); OECD Family Database (2023); World Bank ASPIRE Database (2023); UNICEF Global Database on Childcare (2022); national labour force surveys.

Table 2.3 shows significant distinction between the two high-income European countries and the two lower-income ones across the sample, and also the differences within each pair. Denmark's provisions are comprehensive and near-universal. Italy's policy formally provides for generous packages, but there exist coverage gaps which differ regionally. This regional heterogeneity implies that many families, particularly in the South are yet to benefit from the policy. The second feature we see is on coverage. In Denmark, informal employment is negligible, with almost all women working in formal employment. In Italy, coverage is also broad, though gaps exist for some atypical contracts. In Indonesia, less than 20% of working women are in formal employment and thus eligible for leave benefits.[11] In South Africa, the formal sector covers about forty percent of employed women, and the Unemployment Insurance Fund ceiling excludes higher-earning formal workers as well. Childcare availability is a significant differentiator. In Denmark, two-thirds of children under three are in formal care; in Indonesia, the figure is below five percent. This institutional gap is likely to be a primary driver of the near-zero recovery rates observed in Indonesia and South Africa.[12]

Cash transfers partially compensate for the absence of formal insurance. Denmark’s universal child benefit is substantial. Italy’s means-tested assegno unico provides meaningful support for lower-income families. South Africa’s Child Support Grant, while modest at about thirty-five dollars per month in purchasing power terms, reaches roughly two-thirds of young children and has been shown to improve nutritional and educational outcomes.[13] Indonesia’s PKH is smaller in both amount and coverage. These transfers are relevant for our analysis because they represent the only institutional buffer available to informal workers. The question is whether they are sufficient to offset the household income loss we document next.

In Indonesia and South Africa, the formal provisions are thin, but coverage remains a challenge. Indonesia's maternity leave law guarantees full wage replacement, but fewer than twenty percent of employed women are in formal employment and thus eligible. South Africa's Unemployment Insurance Fund provides maternity benefits to formal employees, but with a low replacement rate and a ceiling that excludes higher earners; the majority of working women, who are in informal or

[11] ILO (2018), *Women and Men in the Informal Economy*, documents the concentration of women in informal employment and its implications for access to maternity protection.
[12] Cameron, Contreras Suarez, and Tseng (2019) discuss the role of childcare constraints in Indonesia. For South Africa, see Casale and Posel (2021) on the interaction between childcare and female employment.
[13] Agüero, Carter, and Woolard (2009) find positive effects of the Child Support Grant on child height-for-age. The grant’s design is described in Lund et al. (2009).

casual employment, have no access to contributory benefits. The Child Support Grant partially fills this gap, but its modest amount, roughly thirty-five dollars per month in purchasing power parity terms, replaces only a small fraction of the typical maternal earnings loss we document below.

Stylized fact 3: Social protection around childbirth varies from comprehensive and near-universal (Denmark) to fragmented and largely absent for the majority of working women (Indonesia). The critical gap is not the generosity of formal provisions but their coverage, particularly of informal workers.

2.4 Household Income Dynamics After Childbirth

The child penalty literature focuses on maternal earnings. But for intergenerational transmission, the relevant variable is the household income available to invest in children. If spousal earnings rise in response to the maternal drop, or if public transfers compensate, the household income loss may be small even when the maternal earnings loss is large. If, on the other hand, spousal responses are weak and public transfers are minimal, the household income loss will be close to the maternal loss.

Table 2.4: Household Income Changes Around First Birth (Event Year k = 1)

	Denmark	Italy	Indonesia	South Africa
Maternal earnings change	–15%	–30%	–45%	–40%
Spousal earnings change	+8%	+6%	+2% (n.s.)	+5% (n.s.)
Household equivalised income change	–4%	–14%	–28%	–22%
Share of households reporting financial distress (0–2 yrs post-birth)	<10%	18%	42%	38%
Share cutting food expenditure (0–2 yrs post-birth)	—	8%	25%	18%

Notes: Maternal and spousal earnings changes are measured as the percentage change from the pre-birth average (event years k = –5 to –2) to event year k = 1. Household equivalised income is adjusted for household composition using the OECD modified equivalence scale. Financial distress is defined as self-reported difficulty meeting basic needs. Estimates for Denmark are from Kleven, Landais, and Søgaard (2019); for Italy from Zamberlan and Barbieri (2023) and the Survey on Household Income and Wealth; for Indonesia from authors' calculations using the Indonesian Family Life Survey, waves 4–5; for South Africa from authors' calculations using the National Income Dynamics Study, waves 1–5. The food expenditure and financial distress shares for Indonesia and South Africa are based on household consumption modules in the respective surveys.

Table 2.4 shows why the institutional gradient in social protection matters for intergenerational transmission. In Denmark, the maternal earnings drop is relatively modest, the spousal response is substantial, and public transfers further cushion the household. The result is a household income loss of only 4 percent, an amount that most families can absorb without meaningful disruption to consumption. In Italy, the maternal drop is larger and the spousal response somewhat weaker, yielding a household income loss of 14 percent. In Indonesia and South Africa, the maternal decline is severe, spousal responses are small and statistically insignificant, and public transfers are minimal or absent. Household income falls by over 20 percent in the year after birth. Moreover,

over a third of households report financial distress and reduced food expenditure, suggesting that consumption smoothing is incomplete and that the income shock translates directly into material hardship[14].

The last two rows of the table are important for the mediation analysis. In Indonesia, a quarter of household's report cutting food expenditure in the years following a birth. In South Africa, the figure is just under a fifth. These are households with a new infant, a period when nutritional adequacy is most critical for child development. The finding that consumption smoothing fails for a substantial minority of households in both countries directly previews the financial channel we test in our mediation analysis.

Stylized fact 4: In low-safety-net settings, the maternal earnings loss at childbirth translates into a large, largely uninsured household income shock. Spousal earnings responses and public transfers offset only a small fraction of the loss, and a significant share of household's report cutting essential consumption.

Figure 2.2: Decomposition of Household Income Change at First Birth

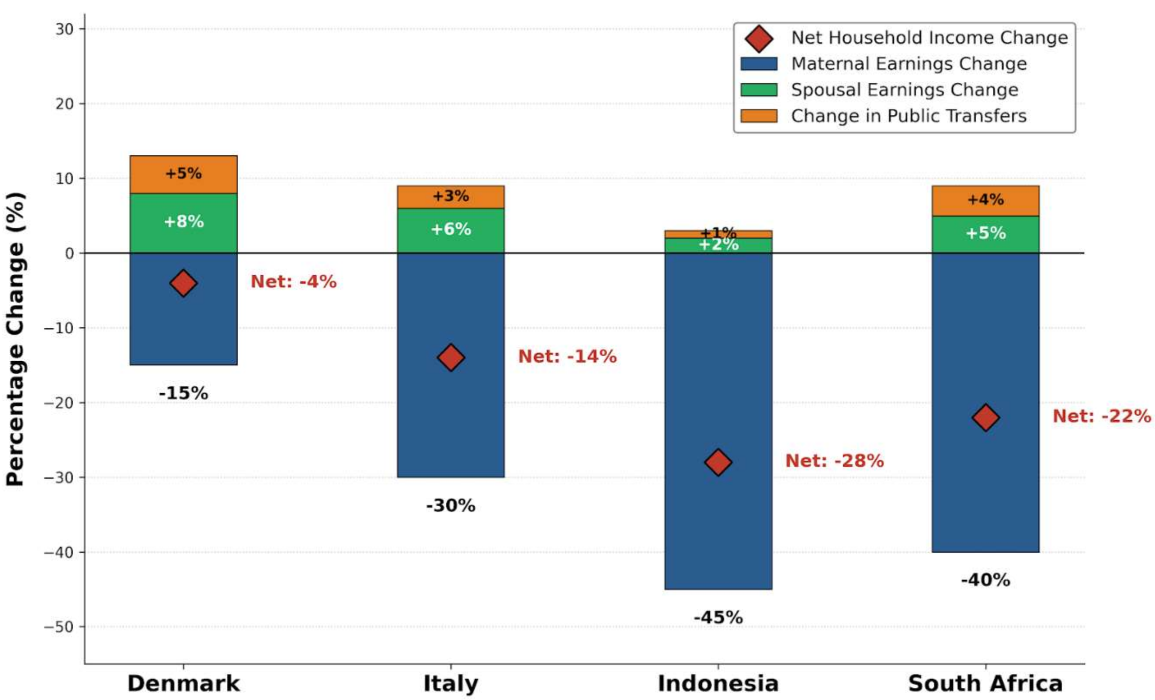

Decomposition of the household income change at first birth into maternal earnings loss, spousal earnings response, and public transfer response. The net household income change is the sum of the three components. Denmark shows strong buffering; Indonesia and South Africa show large uninsured income losses. Source: Table 2.4 and authors' calculations

2.5 Gender Norms and Maternal Employment

[14] Almond and Currie (2011) review the evidence that early-life income shocks have lasting effects on child development.

The institutional constraints we have described so far do not operate in a vacuum. They interact with social norms about maternal employment, caregiving, and the appropriate division of labour within households[15]. Table 2.5 presents comparative indicators of the normative environment in our four countries.

Table 2.5: Gender Norm Indicators Related to Maternal Employment

Indicator	Denmark	Italy	Indonesia	South Africa	OECD Average
Agree: "Preschool child suffers if mother works"	10%	32%	65%	42%	25%
Agree: "Men should have priority for jobs when scarce"	4%	24%	72%	38%	15%
Female labour force participation, age 25–34	82%	62%	55%	48%	72%
Women citing family responsibilities as reason for not working	8%	28%	48%	35%	18%
Female-to-male ratio of unpaid care work hours	1.5	2.6	3.2	2.8	1.8
Firms offering flexible work arrangements	58%	22%	<10%	15%	40%

Notes: Attitudinal indicators are from the World Values Survey, Wave 7 (2017–2022). Labour force participation rates are from ILO modelled estimates (2023). The share citing family responsibilities is from national labour force surveys and the European Labour Force Survey. Unpaid care work ratios are from UN Women (2020). Firm-level flexible work data are from the World Bank Enterprise Surveys, most recent available year for each country.

The normative gradient across our four countries is as steep as the institutional one. In Denmark, only one in ten respondents believes that preschool children suffer if their mother works; in Indonesia, nearly two-thirds hold this belief. The belief that men should have priority for scarce jobs, which captures the legitimacy of female labour force attachment, is virtually absent in Denmark but held by nearly three-quarters of Indonesians. Italy and South Africa fall between these extremes, with Italy closer to the European norm and South Africa showing a mix of relatively egalitarian attitudes alongside strong traditional norms in some segments of the population.

These beliefs have material consequences. In Indonesia and South Africa, roughly a third to a half of economically inactive women cite family responsibilities as the reason, compared to eight percent in Denmark. The unpaid care work burden borne by women is two to three times that of men, further constraining the type, location, and hours of employment that mothers can

[15] Bertrand (2020) provides a recent survey of the role of gender norms in shaping labour market outcomes.

realistically pursue. And firms are far less likely to offer flexible work arrangements in the two non-European countries, reducing the scope for combining work and care even when mothers wish to do so.

Importantly, norms are not independent of institutions. Where childcare is unavailable, the belief that young children suffer if their mother works becomes self-fulfilling. Here, mothers are forced into informal or low-quality care arrangements, and the observed negative outcomes reinforce the norm. Breaking this cycle requires institutional change that makes maternal employment compatible with child wellbeing[16]

Stylized fact 5: Norms against maternal employment are strongest in the low-safety-net settings, and they coexist with low female labour force participation, high unpaid care burdens, and limited employer accommodation. These norms both reflect and reinforce the institutional constraints that generate the child penalty.

3. Data and Methodology

3.1 Theoretical Framework

Consider a unitary household that derives utility from consumption, leisure, and the future human capital of its children.[17] Prior to the arrival of the first child, this household has two choices on its time allocation. First, the household allocates the time of its adult members between market work and leisure. Second, it allocates income between consumption and savings, according to its preferences and the prevailing wage rates. The arrival of a first child at time *t* = 0 then changes this allocation in two ways. One, it introduces a new demand for caregiving time, and two, it may change the wage offer faced by the mother if employers statistically discriminate against mothers or if institutional supports for combining work and care are absent.

The human capital of child *c*, observed in early adulthood, is produced according to:

$$H_c = f(I_{0-5}, Q_{0-5}, X_c, \varepsilon_c),$$

where I_{0-5} denotes material investments in the child during the first five years of life (nutrition, health care, educational materials, housing quality), Q_{0-5} denotes the quality and stability of the caregiving environment (including parental time inputs and the absence of disruptive household changes), X_c is a vector of predetermined child and family characteristics (including parental education and baseline household wealth), and ε_c is an idiosyncratic shock. The function f is increasing and concave in both arguments, and the two inputs may be complements or substitutes depending on context.

[16] The interplay between institutions and norms is discussed by Alesina, Giuliano, and Nunn (2013) in the context of female labour force participation.

[17] The unitary household model is a simplification. Bargaining models (e.g., Lundberg and Pollak, 1993) would allow the distribution of the earnings shock within the household to affect child investment independently of total household income. In our setting, the unitary model is sufficient to motivate the main channels, but we note that intra-household bargaining may amplify or dampen the transmission depending on who controls resources. This is an area for future research.

Material investments I_{0-5} are constrained by the household's permanent income, but in the presence of credit constraints and incomplete insurance markets, transitory income shocks also affect investment. In low- and middle-income countries, where formal credit and insurance markets are thin, the sensitivity of child investment to contemporaneous income is likely to be high.[18] The quality of the caregiving environment Q_{0-5} is influenced by maternal labour market attachment, by the stability of the household unit, and by the availability of alternative caregiving arrangements (formal childcare, extended family).

The first childbirth generates an exogenous shock to maternal earnings, denoted ΔY_m. The magnitude and persistence of this shock depend on the institutional environment, as documented in Section 2.[19] We write the shock as a function of the policy vector Π:

$$\Delta Y_m = g(\Pi) < 0,\ g'(\Pi) > 0,$$

where Π indexes the generosity and coverage of maternity leave, childcare availability, and employment protection. In high-protection settings (Denmark), $| \Delta Y_m |$ is small and temporary; in low-protection settings (Indonesia, South Africa), it is large and persistent.

2. Transmission channel

The household may respond to the maternal earnings shock along three margins, each of which generates a distinct channel of intergenerational transmission.

Financial resource channel. The direct effect of ΔY_m is to reduce household income. Let ΔY_f denote the spousal labour supply response and ΔB denote the change in public transfers. The net household income shock is:

$$\Delta Y_h = \Delta Y_m + \Delta Y_f + \Delta B.$$

In settings where spousal responses are weak—because male labour supply is already high or because partners face similar labour market constraints—and where public transfers are minimal, the household income loss is large.[20] This income loss reduces material investments I_{0-5}. To the extent that these investments affect child human capital, the maternal earnings penalty transmits intergenerationally through a pure income effect. We refer to this as the **financial resource channel**.

[18] On employer discrimination against mothers, see Correll, Benard, and Paik (2007) for experimental evidence and Kleven et al. (2019) for evidence from the child penalty literature that the event-study pattern is driven by labour demand as well as labour supply.
[19] The complementarity or substitutability of material investments and caregiving quality is discussed in Cunha and Heckman (2007) in the context of the technology of skill formation.
[20] On the sensitivity of child investment to transitory income in the presence of credit constraints, see Dahl and Lochner (2012) for the United States and Agüero, Carter, and Woolard (2009) for South Africa. The spousal labour supply response to childbirth is studied by Kleven et al. (2019), who find that it is small in most countries. In low-income settings, male labour supply is often already near its maximum, and underemployment constrains the capacity to increase earnings further.

Maternal employment instability channel. Beyond the income loss, the disruption of the mother's labour market trajectory may independently affect child outcomes. A mother who leaves formal employment and later re-enters in informal or precarious work may experience a depreciation of skills, a loss of labour market networks, and increased stress.[21] Repeated job changes or prolonged non-employment may also destabilise household routines and reduce the predictability of the caregiving environment. Even if household income recovers partially through other sources, this disruption can reduce Q_{0-5}, the quality and stability of the caregiving environment, and thus affect child development. We refer to this as the employment instability channel.

Household structure channel. Income stress and maternal employment disruption may increase the probability of changes in household composition—marital dissolution, the departure or addition of household members, or labour migration by one or both parents.[22] Such changes can disrupt children's caregiving arrangements, alter their residential and school environments, and affect their emotional wellbeing. We refer to this as the household structure channel.

These three channels are not mutually exclusive, and they may reinforce one another. For example, a large income loss may force a mother to accept informal work far from home, which simultaneously reduces household income, increases employment instability, and raises the probability of household splitting. The relative importance of each channel is an empirical question, which our mediation decomposition addresses.

3. Mediating role of institutions

The institutional environment Π moderates each channel. In high-protection settings, the earnings shock ΔY_m is small, while public transfers ΔB partially offset the loss in earnings, and spousal labour supply ΔY_f can increase because flexible work arrangements and affordable childcare make it feasible for both parents to work. The household income loss ΔY_h is negligible, and consumption smoothing is nearly complete. The financial resource channel is therefore muted.

---

[21]Almond and Currie (2011) review the broader evidence that early-life income shocks have lasting effects on health and human capital. On the scarring effects of career interruptions, see Bertrand, Goldin, and Katz (2010) for evidence on the earnings penalty associated with labour market withdrawal. In developing countries, the shift from formal to informal employment may be particularly costly because informal work is often characterised by lower and more volatile earnings and by the absence of social insurance.

[22] The institutional determinants of the child penalty are analysed in Kleven (2022) and Olivetti and Petrongolo (2017). The key insight is that the availability of childcare and job-protected leave, not the length of leave per se, determines the speed of maternal earnings recovery. Yean and Neal (2015) document the relationship between income shocks and marital dissolution. In low-income settings, labour migration by one or both parents is a common response to economic stress, and it can have ambiguous effects on children: remittances may increase material resources, but parental absence may reduce the quality of the caregiving environment (Antman, 2013).

[10] The decomposition follows Gelbach (2016), who shows that the change in a coefficient when additional variables are added to a linear model can be exactly decomposed into the contributions of each added variable, using a formula based on the coefficients from auxiliary regressions of the added variables on the original regressors.

[11] The role-model hypothesis is discussed in Fernández, Fogli, and Olivetti (2004), who show that men whose mothers worked are more likely to have working wives, and in Olivetti, Patacchini, and Zenou (2020), who find intergenerational transmission of gender role attitudes. In our context, maternal employment disruption may signal to daughters that combining work and family is infeasible, reducing their own labour market aspirations.

Moreover, job-protected leave and the availability of formal childcare mean that maternal employment is not disrupted, since the mother returns to the same employer and continue to pursue same career trajectory. The employment instability channel is also muted. And because the income shock is small, the probability of household structural change is unaffected.

In low-protection settings, none of these buffers exist. The earnings shock is large and persistent. Public transfers are minimal. Spousal labour supply cannot increase because male employment is already high or because partners face the same lack of childcare and flexible work. The household income loss is large and uninsured. Consumption smoothing is incomplete, and child investment falls. The mother's employment trajectory is disrupted, often permanently, and the probability of household structural change rises. All three channels are active, and the intergenerational transmission parameter θ is correspondingly large.

Italy occupies an intermediate position. Its formal social protection is relatively generous, but coverage gaps—especially in childcare—and regional heterogeneity mean that the buffering is incomplete. We therefore expect transmission to be present but smaller than in Indonesia and South Africa. Thus, to formalize this intergenerational transmission, then let the child outcome be written as a reduced-form function of the maternal earnings shock, the pre-birth controls, and the mediating variables:

$$Y_{ic}^{child} = \delta_0 + \theta \cdot \text{CumLoss}_i^{0-5} + X_i'\Gamma + \nu_{ic},$$

where CumLoss_i^{0-5} is the cumulative maternal earnings loss during the child's first five years, as defined in equation (3) of the empirical strategy. The parameter θ captures the total intergenerational effect of the earnings loss, operating through all channels.

Now let M_1 denote the financial channel (household income loss), M_2 the employment instability channel, and M_3 the household structure channel. Each is a function of the maternal earnings loss and of the institutional environment:

$$M_1 = \phi_1 \cdot \text{CumLoss} + u_1,\ M_2 = \phi_2 \cdot \text{CumLoss} + u_2,\ M_3 = \phi_3 \cdot \text{CumLoss} + u_3.$$

The ϕ coefficients measure the sensitivity of each mediator to the maternal earnings loss. In high-protection settings, ϕ_1 is small because ΔY_h is small relative to ΔY_m; ϕ_2 is small because employment is protected; and ϕ_3 is small because income shocks are smoothed.

The child outcome is in turn a function of the mediators:

$$Y^{child} = \psi_1 M_1 + \psi_2 M_2 + \psi_3 M_3 + X'\Gamma + \varepsilon.$$

The total intergenerational transmission parameter θ can be decomposed as:

$$\theta = \psi_1\phi_1 + \psi_2\phi_2 + \psi_3\phi_3.$$

This decomposition is the basis of our mediation analysis. The term $\psi_j \phi_j$ is the contribution of channel *j* to the total transmission. The share of transmission attributable to channel *j* is $(\psi_j \phi_j)/\theta$.[23]

The framework yields five predictions that we take to the data. The predictions are stated in terms of the cross-country comparison and the within-country heterogeneity that should follow from the logic of the model.

Prediction 1 (Institutional gradient in transmission). The intergenerational transmission parameter θ will be negligible in Denmark, where all three ϕ coefficients are close to zero; modest in Italy, where partial buffering leaves some scope for transmission; and large in Indonesia and South Africa, where the absence of institutional buffers leaves all channels active.

Prediction 2 (Dominance of the financial resource channel). In the three countries where transmission is detectable, the financial resource channel $\psi_1 \phi_1$ will account for the largest share of θ. This follows from the evidence in Section 2 that spousal earnings responses and public transfers are insufficient to offset the maternal earnings loss, and that consumption smoothing is incomplete. The income effect should therefore be the primary driver of reduced child investment.

Prediction 3 (Concentration among the vulnerable). Transmission will be stronger among mothers with low education and low pre-birth household wealth, for whom the income shock is most binding. These mothers have fewer assets to draw down, less access to credit, and weaker labour market prospects, so the same earnings loss translates into a larger reduction in child investment.

Prediction 4 (Employment instability as an independent channel). The employment instability channel $\psi_2 \phi_2$ will be more important in Indonesia and South Africa than in Italy, because the formal-informal divide in those labour markets makes re-entry into stable employment more difficult. This channel should be absent in Denmark.

Prediction 5 (Gender-differentiated transmission). The intergenerational transmission may be stronger from mother to daughter than from mother to son, if maternal employment disruption affects daughters' aspirations through a role-model effect.[24] This pattern is most likely to appear in settings where maternal employment is both observable and non-normative, as in Indonesia and South Africa.

These predictions guide the empirical analysis in Section 4. The event-study estimates (Stage 1) quantify the maternal earnings shock ΔY_m and establish its variation across countries. The

---

[23] The spousal labour supply response to childbirth is studied by Kleven et al. (2019), who find that it is small in most countries. In low-income settings, male labour supply is often already near its maximum, and underemployment constrains the capacity to increase earnings further. The decomposition follows Gelbach (2016), who shows that the change in a coefficient when additional variables are added to a linear model can be exactly decomposed into the contributions of each added variable, using a formula based on the coefficients from auxiliary regressions of the added variables on the original regressors.

[24] On the scarring effects of career interruptions, see Bertrand, Goldin, and Katz (2010) for evidence on the earnings penalty associated with labour market withdrawal. In developing countries, the shift from formal to informal employment may be particularly costly because informal work is often characterised by lower and more volatile earnings and by the absence of social insurance.

intergenerational transmission regressions (Stage 2) estimate θ and test Prediction 1, as well as the heterogeneity patterns of Predictions 3 and 5. The mediation decomposition (Stage 3) quantifies the shares attributable to each channel and tests Predictions 2 and 4.

3.2 Data Sources and Sample Construction

We draw on four distinct data sources, one for each study country, chosen to provide the best available combination of longitudinal coverage, reliable earnings measurement, and the ability to link mothers to their children over extended periods. Table 3.1 summarises the key features of each dataset.

Table 3.1: Data Sources and Sample Sizes

	Denmark	**Italy**	**Indonesia**	**South Africa**
Data source	Population registers (DST)	SHIW panel	IFLS	NIDS
Type	Administrative panel	Household survey panel	Household survey panel	Household survey panel
Coverage period	1995–2020	2004–2020	1993–2014	2008–2017
Number of waves	Annual	Biennial (approx.)	5 waves	5 waves
Mothers in sample	152,000	8,150	4,820	3,150
Children tracked to 18+	230,000	9,200	5,100	3,400
Key earnings measure	Annual gross earnings (admin)	Annual net earnings (self-report)	Monthly earnings, all jobs	Monthly earnings, all jobs
Attrition rate (mothers, end of panel)	<2%	~20%	~25%	~30%

Notes: DST = Statistics Denmark. SHIW = Survey on Household Income and Wealth, Bank of Italy. IFLS = Indonesian Family Life Survey. NIDS = National Income Dynamics Study (South Africa). Mothers meeting sample criteria: aged 18–40 at first birth, observed in at least one pre-birth and two post-birth waves. Attrition rates for the three survey datasets are the share of baseline mothers not observed at the final available wave; for IFLS and NIDS, this includes non-contact and non-coresidence of the child. The Danish registers are near-complete for the resident population; the attrition rate reflects permanent emigration and mortality.

3.2.1 Denmark

The Danish data are drawn from population-wide administrative registers maintained by Statistics Denmark. These registers link individual earnings histories, demographic events, and family relationships through unique personal identifiers. We use the full population of women who gave birth to their first child between 1995 and 2012, following the sample construction procedures of Kleven, Landais, and Søgaard (2019). Earnings are measured as annual gross labour income from all employment, including self-employment, drawn from tax records. The data are available annually and suffer from negligible attrition, allowing precise estimation of event-time coefficients over long pre- and post-birth windows.[25]

We define the child sample as the first-born child of each mother, tracked through the population register to age 22–25. Educational attainment is obtained from the Danish education registers, which record completed years of schooling and highest qualification achieved. Early-career earnings are taken from the same tax records as maternal earnings.

3.2.2 Italy

The Italian data come from the Survey on Household Income and Wealth (SHIW), a biennial household panel conducted by the Bank of Italy. The SHIW collects detailed information on income, employment, and household composition from a representative sample of the Italian population. We use all waves from 2004 to 2020 and follow the sample definitions of Zamberlan and Barbieri (2023).[26]

Because the SHIW is a household survey, we identify mothers by linking women to their children within the household roster. A woman is included in the sample if she is observed as the mother of a child born during the panel period, is aged 18–40 at the time of birth, and is present in at least one pre-birth and two post-birth waves. The survey's two-year interval means that event time is not observed at annual frequency; we construct an annualised panel using linear interpolation between survey waves for the years between a mother's first and last observed wave.[27]

Earnings are measured as annual net income from employment and self-employment, deflated to 2019 prices. Educational outcomes for children are taken from the survey responses of the mother or the child (if co-resident) at the latest available wave, when the child is aged 18–25.

3.2.3 Indonesia

The Indonesian Family Life Survey (IFLS) is a longitudinal household survey that has followed a representative sample of Indonesian households since 1993. We use all five waves: 1993, 1997, 2000, 2007, and 2014. The IFLS collects comprehensive employment and earnings data, including informal sector and self-employment activities, as well as complete fertility histories for ever-married women.[28]

---

[25] Kleven, Landais, and Søgaard (2019) provide a detailed description of the Danish register data and the event-study methodology. We follow their sample selection criteria, restricting the sample to mothers who had their first child between 1995 and 2012 and who are observed in the registers continuously from five years before to ten years after the birth.

[26] Zamberlan and Barbieri (2023) use the SHIW to estimate the child penalty in Italy and discuss the strengths and limitations of the survey for this purpose. The SHIW is representative of the Italian population but has a smaller sample size than the Nordic administrative registers, which limits the precision of the estimates for some subgroups.

[27] The interpolation assumes that earnings change linearly between survey waves. This assumption is reasonable for earnings, which tend to follow smooth trends, but it may miss short-term fluctuations. The wave-only specification, which does not rely on interpolation, provides a robustness check.

[28] The IFLS is described in detail by Frankenberg and Karoly (1995) and by Strauss, Witoelar, and Sikoki (2016). The survey is unusual among developing-country panels for its long duration and high re-contact rates. Attrition between waves is addressed through tracking of movers and re-interview of split-off households.

Mothers are included if they report a first birth during the observation period, are aged 18–40 at that birth, and are observed in at least one pre-birth and two post-birth waves. Because the wave spacing is irregular, with gaps of three to seven years, we construct an annualised panel using linear interpolation between waves for years in which the mother is not directly observed. The interpolation uses earnings values at the bounding survey waves. In robustness checks, we also estimate a wave-only specification in which no interpolation is used and event time is expressed in wave-relative terms.

Earnings are measured as total monthly earnings from primary and secondary employment, annualised to a full-year equivalent and deflated to 2019 Indonesian Rupiah using the consumer price index. We winsorise the earnings distribution at the 1st and 99th percentiles within each survey wave to reduce the influence of extreme outliers in self-reported informal earnings.[29]

Child outcomes are obtained from the 2014 wave, when the first-born children of sample mothers are aged 18–25. We measure completed years of education and, where available, monthly earnings from the child's own survey responses. Children who are not co-resident with the mother in 2014 are excluded from the child sample; we examine the potential selectivity of this restriction in the robustness analysis.

3.2.4 South Africa

South Africa's National Income Dynamics Study (NIDS) is a nationally representative household panel that began in 2008 and has completed five waves (2008, 2010, 2012, 2014, and 2017). NIDS collects detailed labour market information, including formal and informal employment, and allows the linkage of mothers to children through household rosters and fertility histories.[30]

The mother sample includes women aged 18–40 at first birth who are observed in at least one pre-birth and two post-birth waves. As with the IFLS, we construct an annualised panel using linear interpolation between the two-year survey intervals. Earnings are measured as total monthly earnings from all employment, annualized and deflated to 2019 South African Rand. We winsorise at the 1st and 99th percentiles and report both the interpolated annual estimates and a wave-only specification.

Child outcomes are taken from the 2017 wave for children aged 18–25, including years of completed education and early-career earnings. The co-residence restriction applies as in Indonesia: children must be living in the same household as the mother at the time of outcome measurement.

3.3 Variable Construction

[29] Self-reported earnings in developing countries are known to be noisy, particularly for the self-employed and informal workers. Winsorising is a standard approach to reducing the influence of extreme values; see, for example, Akee et al (2010) for a discussion of the trade-offs.

[30] NIDS is described by Leibbrandt, Woolard, and de Villiers (2009). The survey has experienced attrition rates of approximately 15–20 percent between waves, which is comparable to other panel surveys in middle-income countries.

Table 3.2 lists the main variables used in the analysis, their definitions, and the sources from which they are constructed in each country.

Table 3.2: Variable Definitions and Sources

Variable	Definition	Denmark	Italy	Indonesia	South Africa
Maternal earnings	Log real annual earnings from all employment	Tax register	SHIW self-report	IFLS self-report	NIDS self-report
Cumulative loss (CumLoss)	Sum of (actual – pre-birth average earnings) over child ages 0–5	Tax register	SHIW (interpolated)	IFLS (interpolated)	NIDS (interpolated)
Child education	Completed years of schooling at age 18–25	Education register	SHIW	IFLS (2014)	NIDS (2017)
Secondary completion	Binary = 1 if ≥12 years of schooling	Register	SHIW	IFLS	NIDS
Early-career earnings	Log real monthly earnings, age 22–25	Tax register	SHIW	IFLS	NIDS
Household equivalised income	Total disposable income / OECD equivalence scale	Register	SHIW	IFLS	NIDS
Employment instability index	PC1 of: share informal, any long-term non-employment, CV of earnings	Register (part-time share)	SHIW (part-time share)	IFLS (informal share)	NIDS (informal share)
Household structure change	Marital dissolution or household head change, child ages 0–5	Register	SHIW	IFLS	NIDS
Pre-birth controls	Maternal education, age at first birth, pre-birth wealth, region FE	Register	SHIW	IFLS	NIDS

Notes: PC1 = first principal component. CV = coefficient of variation. OECD equivalence scale: weight 1.0 for first adult, 0.5 for additional adults, 0.3 for children under 14. For Denmark and Italy, the employment instability index uses the share of post-birth years in part-time work (informal employment is negligible). For Indonesia and South Africa, it uses the share in informal employment. All monetary variables are deflated to 2019 local currency. Pre-birth wealth is measured as equivalized income percentile (Denmark) or asset index quintile (Italy, Indonesia, South Africa).

3.3.1 Maternal Earnings and Cumulative Loss

Maternal earnings are expressed in the natural logarithm of real annual earnings. For each mother, we compute the pre-birth baseline as the average of log earnings in event years $k = -5$ to $k =$

-2 (or the available pre-birth years in the survey countries). The cumulative loss is then defined as

$$\text{CumLoss}_i^{0-5} = \sum_{t:\text{child age}\in[0,5]} \left(Y_{it} - \bar{Y}_i^{\text{pre}}\right),$$

where Y_{it} is the log real annual earnings of mother i in calendar year t, and $\bar{Y}_i^{\text{pre}}$ is the pre-birth average. This measure is the sum, over the child's first five years of life, of the difference between actual and pre-birth earnings. A more negative value indicates a larger cumulative earnings shortfall. We focus on the window from birth to age five because the early childhood period is critical for cognitive and non-cognitive development and because household income during this window has been shown to have lasting effects on child outcomes.[31] In robustness checks, we vary the window to ages 0–3 and 0–8.

3.3.2 Child Outcomes

The primary child outcome is completed years of education, measured when the child is aged 18–25. For children still enrolled in school at the time of survey, we use the highest grade completed. The secondary outcome is a binary indicator for secondary school completion (12 years of schooling in all four countries). Where available, we also measure the log of monthly earnings from all employment at age 22–25.

3.3.3 Mediators

The mediator vector M_i is constructed from the same observation window (child ages 0–5) to align temporally with the cumulative loss measure.

Cumulative household equivalised income. This is the sum of annual household equivalised income over child ages 0–5, expressed as a deviation from the pre-birth average. Household income is adjusted for household composition using the OECD modified equivalence scale. This variable captures the financial resource channel.

Employment instability index. For Indonesia and South Africa, we construct a composite measure of maternal employment instability using the first principal component of three indicators: (i) the share of post-birth years (child ages 0–5) in which the mother is employed in informal work, (ii) a binary indicator for any spell of non-employment exceeding 12 consecutive months during the same window, and (iii) the coefficient of variation of annual earnings across post-birth years. For Denmark and Italy, where informal employment is negligible, we replace the informal share with the share of post-birth years in part-time work (defined as usual weekly hours below 30). This index captures the degree of disruption to the mother's labour market trajectory.

[31] Almond and Currie (2011) review the extensive evidence that economic conditions in early childhood have lasting effects on health, education, and earnings. Heckman (2006) provides the theoretical framework for the technology of skill formation, emphasising the sensitivity of early childhood to parental investments.

Household structure change. A binary indicator equal to 1 if the mother experienced marital dissolution (divorce or separation), a change in household head, or a residential move across municipal boundaries during the child's first five years. This variable captures the household structure channel.

3.3.4 Pre-Birth Controls

All intergenerational regressions include a vector of pre-birth maternal characteristics that are predetermined with respect to the child penalty. These are: the mother's completed years of education, her age at first birth, a measure of pre-birth household economic status (equivalised income percentile in Denmark, asset index quintile in the survey countries), and region or province fixed effects. These controls absorb permanent differences across mothers that may be correlated with both the child penalty and child outcomes.

3.4 Data Quality and Limitations

The administrative Danish data are of high quality and suffer from negligible measurement error in earnings. The three survey datasets face the challenges common to longitudinal household surveys in developing and middle-income countries: measurement error in self-reported earnings, particularly from informal or self-employment activities; non-random attrition; and irregular spacing of survey waves.[32]

We address measurement error in several ways. First, winsorising at the 1st and 99th percentiles reduce the influence of extreme outliers. Second, in the IFLS and NIDS, we compare earnings reports across different survey modules where available and estimate reliability ratios; these suggest that the signal-to-noise ratio is lower than in administrative data but still sufficient to recover event-study patterns.[33] Third, because classical measurement error in the cumulative loss variable will attenuate the transmission coefficient θ toward zero, our intergenerational estimates should be interpreted as lower bounds on the true effect.

Attrition is a concern in all three survey countries. Mothers who drop out of the panel may have systematically different earnings trajectories than those who remain. We address this by examining the correlates of attrition, re-weighting the sample using inverse probability weights based on pre-birth characteristics, and applying the bounding procedure of Lee (2009) to assess the sensitivity of our results to worst-case attrition scenarios.[34] The results of these checks are reported in the robustness section of the main results.

---

[32] On the general problem of measurement error in survey-based earnings in developing countries, see Bound, Brown, and Mathiowetz (2001). For the IFLS specifically, Dong (2016) validates earnings reports against administrative data and finds that measurement error is substantial but that the IFLS earnings data are adequate for recovering average trends.

[33] The reliability ratio is the fraction of the variance of reported earnings that is attributable to true earnings. In the IFLS, we compare the main employment module with a supplementary module on income and find a reliability ratio of approximately 0.7. This implies that coefficients will be attenuated by about 30 percent, so our estimates are conservative.

[34] Lee (2009) bounds are computed by trimming the sample symmetrically in proportion to the differential attrition rate across the treatment (high-cumulative-loss) and control (low-cumulative-loss) groups. We report the Lee bounds as part of the robustness analysis in the Online Appendix.

The irregular wave spacing in the IFLS is a particular challenge because the gaps between waves are long (up to seven years). Our primary specification uses linear interpolation to fill the inter-wave years. To verify that the results are not driven by the interpolation, we also estimate a wave-only specification in which we define event time in wave-relative terms (the number of survey waves before or after the wave in which the first birth occurs) and include wave-relative dummies and mother fixed effects. The wave-only specification is less demanding of the data and serves as a robustness check; we report both sets of estimates.

3.5 Empirical Strategy

Our empirical strategy proceeds in three stages. In the first stage, we estimate the dynamic impact of first childbirth on maternal earnings using an event-study design. In the second stage, we link the maternal child penalty to offspring outcomes using a scalar summary measure of cumulative earnings loss. In the third stage, we decompose the intergenerational transmission into the contributions of three mediating channels. This section describes each stage in turn, defines the estimating equations, and discusses the identifying assumptions.

3.5.1 Stage 1: Event-Study Identification of the Child Penalty

The first stage recovers the causal effect of childbirth on maternal earnings in each country. We use the event-study specification that has become standard in the child penalty literature.[35] The estimating equation is

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq -1} \beta_k \ D_{it}^k + \varepsilon_{it}, \qquad (1)$$

where Y_{it} is log real annual earnings of mother i in calendar year t, α_i is a mother fixed effect, γ_t is a calendar-year fixed effect, and D_{it}^k is an event-time indicator equal to 1 if mother i is k years from her first childbirth in year t. The reference period is $k = -1$, the year immediately prior to birth. The coefficients β_k measure the percentage change in maternal earnings in year k relative to the pre-birth baseline, net of economy-wide time effects and permanent differences across mothers.

Mother fixed effects absorb all time-invariant heterogeneity, including cohort, education, ability, and pre-birth labour market attachment. Calendar-year fixed effects control non-parametrically for common time trends, business cycles, and life-cycle earnings growth. The specification deliberately excludes time-varying controls such as marital status, partner's earnings, and the number of additional children. Because these variables are themselves potentially affected by childbirth, their inclusion would induce over-control bias and obscure the total labour market

[35] The event-study methodology for child penalties was developed by Kleven, Landais, and Søgaard (2019) and extended to multi-country comparisons by Kleven et al. (2019). Our specification follows theirs exactly, allowing direct comparability of estimates.

impact of motherhood.[36] The β_k therefore recover the *unconditional* child penalty, which is the policy-relevant parameter.

We bin the event-time endpoints to maintain a stable composition of mothers across the event window. For Denmark, where the administrative data cover a long panel, we estimate separate indicators from $k = -10$ to $k = +15$, binning at $k \leq -10$ and $k \geq 15$. For the three survey countries, the pre-birth window is shorter due to the structure of the panels; we estimate indicators from $k = -5$ (binned at $k \leq -5$) to $k = +8$ (binned at $k \geq 8$). The choice of bin endpoints follows standard practice and does not affect the estimated pattern of the coefficients within the main window.[37]

For the survey-based samples, we construct an annualised panel using linear interpolation of log earnings between survey waves.[38] The interpolation is applied only to the years between a mother's first and last observed wave. To verify that the interpolation does not drive the results, we also estimate a wave-only specification in which event time is defined in wave-relative rather than calendar-year terms. In this specification, D_{iw}^{k} is an indicator equal to 1 if mother i is k survey waves from the wave in which her first birth is observed, and we include mother fixed effects and wave fixed effects. The wave-only estimates are coarser but free of interpolation assumptions.

The key identifying assumption for equation (1) is that, in the absence of childbirth, maternal earnings would have followed parallel trends across event time. This assumption is testable. Under parallel trends, the coefficients β_k for $k < -1$ should be close to zero and show no systematic pre-existing trend. We test this formally with a joint F-test of the null hypothesis that $\beta_k = 0$ for all $k < -1$. We also conduct a placebo exercise in which we assign a pseudo-birth date two years before the actual first birth and re-estimate equation (1); the absence of a pseudo-penalty supports the validity of the design.[39]

We cluster standard errors at the mother level throughout. For the Danish administrative data, the large number of clusters makes the asymptotic approximation highly accurate. For the survey data, the cluster-robust variance estimator is used, and we also report p-values from the wild bootstrap as a robustness check in small samples.[40]

---

[36] Over-control bias in event-study designs is discussed by Angrist and Pischke (2009, Chapter 5). Kleven et al. (2019) explicitly argue for the unconditional specification as the policy-relevant parameter.

[37] Schmidheiny and Siegloch (2023) provide a comprehensive discussion of endpoint binning in event-study designs. The choice of bin endpoints does not affect the estimated pattern within the main window, but it ensures that the composition of mothers at the endpoints is not driven by a small number of observations.

[38] Linear interpolation is used only for the years between a mother's first and last observed survey wave. We do not extrapolate beyond the observed panel period. The assumption is that earnings evolve smoothly between survey waves, which is plausible for annual earnings but may miss short-term fluctuations. The wave-only specification serves as a robustness check that does not rely on this assumption.

[39] Placebo event-time tests are standard in the child penalty literature; see, for example, Kleven et al. (2019, Online Appendix). The pseudo-birth date is assigned two years before the actual first birth, and the event study is re-estimated. If the penalty is driven by childbirth rather than by pre-existing trends, the pseudo-penalty should be zero.

[40] The wild bootstrap is recommended for cluster-robust inference with a small number of clusters (Cameron, Gelbach, and Miller, 2008). In our survey samples, the number of clusters (mothers) is large enough that the standard cluster-robust estimator is reliable, but we report bootstrap p-values as a sensitivity check.

The outputs of Stage 1 are the event-study coefficients $\hat{\beta}_k$ for each country, which trace the dynamic child penalty, and the individual-level summary measure of the penalty, CumLoss_i, which we construct in Stage 2.

3.5.2 Stage 2: Intergenerational Transmission

The second stage links the maternal earnings loss to child outcomes. Because the child penalty is a dynamic trajectory, we must first collapse it into a scalar measure that captures the economic exposure of the child during the formative early years. We define the *cumulative earnings loss* as

$$\text{CumLoss}_i^{0-5} = \sum_{t:\text{child age}\in[0,5]} \left(Y_{it} - \bar{Y}_i^{\text{pre}}\right), \qquad (2)$$

where $\bar{Y}_i^{\text{pre}}$ is the average of mother i's log real annual earnings in the pre-birth years $k = -5$ to -2 (or the available pre-birth years). This measure sums the annual deviations of maternal earnings from the pre-birth baseline over the first five years of the child's life. A more negative value indicates a larger cumulative shortfall. We focus on the window from birth to age five because it aligns with the critical period for early childhood development and because household income during this window has been shown to have lasting effects on health, education, and earnings.[41] In robustness checks, we vary the window to $[0,3]$ and $[0,8]$.

The intergenerational transmission equation is

$$Y_{ic}^{\text{child}} = \delta_0 + \theta \cdot \text{CumLoss}_i^{0-5} + X_i'\Gamma + \nu_{ic}, \qquad (3)$$

where Y_{ic}^{child} is the outcome of child c of mother i, measured at ages 18–25. The primary outcome is completed years of education; secondary outcomes are a binary indicator for secondary school completion and the log of early-career earnings (ages 22–25). The vector X_i contains pre-birth maternal characteristics that are predetermined with respect to the child penalty: maternal education (years or categorical), age at first birth, pre-birth household economic status (equivalised income percentile in Denmark, asset index quintile in the survey countries), and region or province fixed effects. The coefficient θ is the intergenerational transmission parameter of interest. It measures how a larger cumulative maternal earnings loss is associated with child outcomes, conditional on pre-birth characteristics.

We estimate equation (3) separately for each country. We also estimate a set of heterogeneity specifications to test Predictions 3 and 5 from the theoretical framework. To examine socioeconomic heterogeneity, we interact CumLoss_i^{0-5} with indicators of maternal education (above or below primary completion) and with pre-birth household wealth quartiles. To examine gender-specific transmission, we estimate equation (3) separately for daughters and sons.

[41] Almond and Currie (2011) review the evidence that economic conditions in early childhood have lasting effects. Heckman (2006) provides the theoretical framework for the technology of skill formation, emphasising the sensitivity of early childhood to parental investments.

Identification concerns. The parameter θ in equation (3) is an associational, not a causally identified, coefficient. The inclusion of pre-birth controls absorbs permanent differences across mothers that may be correlated with both the child penalty and child outcomes. However, there may remain unobserved time-varying factors that affect both the size of the maternal earnings loss and children's educational attainment. For example, a health shock to the child that requires the mother to reduce her labour supply would simultaneously increase CumLoss and reduce the child's schooling. This would bias θ away from zero, overstating the transmission. Conversely, mothers who experience larger earnings losses may be those who invest more time in their children, creating a positive association that offsets the negative income effect and biases θ toward zero.

We address these concerns in several ways. First, we conduct a bounding exercise following Oster (2019) to assess how large the influence of unobserved selection would need to be to eliminate the estimated transmission coefficient.[42] Second, the heterogeneity analysis provides an indirect test: if θ is small among high-wealth mothers—for whom the same earnings loss should be less binding—this pattern is consistent with a causal interpretation, because unobserved factors would have to be differentially correlated with CumLoss across wealth groups to produce this result. Third, we note that the direction of the likely bias is ambiguous, and we interpret the intergenerational estimates as descriptive associations that quantify the empirical link between maternal labour market disruption and child outcomes, rather than as definitively causal parameters. We return to this limitation in the discussion.

3.5.3 Stage 3: Mediation Decomposition

The third stage quantifies the relative importance of the transmission channels identified in the theoretical framework. We use the decomposition method developed by Gelbach (2016), which is based on the omitted-variables bias formula.[43]

We begin by estimating the baseline model (equation 3), which yields the total intergenerational association $\hat{\theta}$. We then estimate an augmented model that adds a vector of mediators M_i, measured over the child's first five years of life:

$$Y_{ic}^{\text{child}} = \delta_0 + \theta_{\text{aug}} \cdot \text{CumLoss}_i^{0-5} + M_i'\psi + X_i'\Gamma + \nu_{ic}. \qquad (4)$$

The mediator vector includes:

- Cumulative household equivalised income, defined as the sum of annual household equivalised income over child ages 0–5, expressed as a deviation from the pre-birth average. This captures the financial resource channel.

[42] Oster (2019) develops a method for bounding treatment effects under selection on unobservables. The method uses the stability of coefficients and the R-squared when controls are added to assess the influence of unobserved confounders. We implement the procedure using the suggested parameters $R_{\max} = 1.3 \times \tilde{R}$ and $\delta = 1$.

[43] Gelbach (2016) shows that the change in a coefficient when additional variables are added to a linear model can be exactly decomposed into the contributions of each added variable. The decomposition uses the formula $\Delta\theta = \sum_j \hat{\psi}_j \cdot \hat{\phi}_j$, where $\hat{\phi}_j$ is from a regression of the j-th added variable on the original regressors, and $\hat{\psi}_j$ is the coefficient on the j-th added variable in the augmented model.

- Employment instability index, constructed as the first principal component of: (i) the share of post-birth years in informal employment (Indonesia, South Africa) or part-time work (Denmark, Italy), (ii) a binary indicator for any spell of non-employment exceeding 12 consecutive months, and (iii) the coefficient of variation of annual earnings across post-birth years. This captures the employment instability channel.
- Household structure change, a binary indicator equal to 1 if the mother experienced marital dissolution, a change in household head, or residential migration during the child's first five years. This captures the household structure channel.

The change in the coefficient on CumLoss, $\Delta\theta = \hat{\theta} - \hat{\theta}_{\text{aug}}$, measures the amount of the intergenerational association that is accounted for by the mediators jointly. The contribution of each individual mediator M_j is given by $\hat{\psi}_j \cdot \hat{\phi}_j$, where $\hat{\phi}_j$ is the coefficient from an auxiliary regression of M_j on CumLoss and the pre-birth controls. The share of the total transmission attributable to channel j is $\left(\hat{\psi}_j\hat{\phi}_j\right)/\hat{\theta}$.

Interpretation and assumptions. The decomposition is a descriptive accounting exercise, not a causal mediation analysis. The coefficients $\hat{\psi}_j$ in equation (4) reflect both the causal effect of the mediator on the child outcome and any selection into the mediator that is correlated with unobserved determinants of the child outcome. The identifying assumption for a causal interpretation—sequential ignorability—requires that there are no unobserved confounders of the mediator-outcome relationship conditional on CumLoss and the pre-birth controls.[44] This assumption is strong and unlikely to hold fully in our setting. We therefore present the decomposition as a summary of how much of the empirical association between CumLoss and child outcomes operates through each observable channel, without claiming that the channels are themselves causal. Standard errors for the decomposition shares are computed via the delta method.

The decomposition is estimated for the three countries in which Stage 2 detects a non-zero transmission coefficient (Italy, Indonesia, and South Africa). For Denmark, where $\hat{\theta} \approx 0$, the decomposition is not meaningful, consistent with the prediction that public insurance severs the intergenerational link.

4. Results

4.1 Descriptive Statistics

We begin by characterizing the mothers and children in our four samples. Table 4.1 reports pre-birth characteristics of the mothers on whom the event-study analysis is based. Table 4.2 summarizes the outcomes of their first-born children, measured in early adulthood. All statistics

[44] Imai, Keele, and Tingley (2010) provide a formal framework for causal mediation analysis and discuss the sequential ignorability assumption. Our approach follows the descriptive decomposition tradition of Gelbach (2016) rather than the causal mediation tradition of Imai et al., precisely because the identifying assumptions for the latter are too strong in our setting.

are computed using the sample weights provided by each survey (where applicable) and are representative of mothers who gave birth to a first child during the observation window.

Table 4.1: Descriptive Statistics of Mothers at Baseline (Pre-Birth Period)

	Denmark	Italy	Indonesia	South Africa
Number of mothers	152,000	8,150	4,820	3,150
Age at first birth (years)	29.4	31.2	24.8	25.3
Years of education	14.2	12.8	8.1	10.4
Share married at first birth	0.62	0.81	0.94	0.57
Share in formal employment (pre-birth)	0.94	0.72	0.29	0.54
Pre-birth annual earnings (PPP USD)	46,200	28,900	3,620	8,150
Pre-birth household equivalized income (PPP USD)	52,400	34,100	5,480	11,300

Notes: All characteristics are measured in the pre-birth period (event years k = –5 to –1). Education is completed years of schooling. Formal employment is defined as employment covered by social security contributions (Denmark, Italy) or a written contract with paid leave entitlements (Indonesia, South Africa). Earnings and income are expressed in 2019 purchasing power parity (PPP) dollars to facilitate comparison. Pre-birth earnings are the individual-level average over the available pre-birth years. Sources: Authors' calculations from Danish population registers, SHIW, IFLS, and NIDS.

Danish and Italian mothers are considerably older at first birth and have substantially more years of schooling than their Indonesian and South African counterparts. Denmark registers near-universal formal employment and high pre-birth earnings, while Italy exhibits a lower, though still high rate of formal employment. In Indonesia, less than a third of mothers were in formal employment before childbirth, and pre-birth earnings are an order of magnitude lower than in the two European countries. South Africa occupies an intermediate position: formal employment is more common than in Indonesia but far from universal, and earnings are low by high-income standards but roughly double those in Indonesia.

These baseline differences are integral to the comparative design. If the intergenerational transmission of the child penalty is shaped by the availability of institutional buffers, as the theoretical framework predicts, then the baseline contrasts in Table 4.1 should translate into sharply different transmission patterns.

Table 4.2: Descriptive Statistics of First-Born Children's Outcomes (Age 18–25)

	Denmark	Italy	Indonesia	South Africa
Number of children	230,000	9,200	5,100	3,400
Years of completed education	14.1	13.2	10.3	11.0
Share completed secondary school	0.91	0.78	0.52	0.63
Share enrolled in tertiary education	0.48	0.32	0.14	0.12
Early-career monthly earnings (PPP USD)	3,820	1,940	340	620
Age at outcome measurement	22.8	23.1	22.4	23.0

Notes: Children are the first-born offspring of the mothers in Table 4.1. Educational outcomes are measured at the latest wave in which the child is observed between ages 18 and 25. Secondary school completion is defined as 12 years of schooling. Early-career earnings are the log real monthly earnings from all employment at ages 22–25, converted to PPP dollars. Sources: Authors' calculations from Danish population registers, SHIW, IFLS, and NIDS.

The child outcomes also defines the maternal characteristics of the children's mothers. Danish children obtain the most years of schooling, with nearly half enrolling in tertiary education. Italian children average roughly one year less of schooling, and the tertiary enrolment rate is lower. Indonesian and South African children complete considerably fewer years of education, where only about half of Indonesian children and just under two-thirds of South African children finish secondary school. Early-career earnings follow similar patterns, with Danish young adults earning roughly ten times higher than their Indonesian counterparts in PPP terms. These large baseline gaps in children's outcomes provide the backdrop against which the intergenerational transmission of the child penalty occurs.

4.2 Benchmark Child Penalty Estimates

Figure 4.1 shows the event-study estimates of the dynamic impact of first childbirth on maternal log real annual earnings for each of the four countries. The figure plots the coefficients β_k from equation (1) for event years $k = -5$ to $k = +8$, along with 95 percent confidence intervals. The vertical axis measures the percentage change in earnings relative to the reference year, $k = -1$.

Figure 4.1: Event-Study Plots of Maternal Earnings Around First Birth

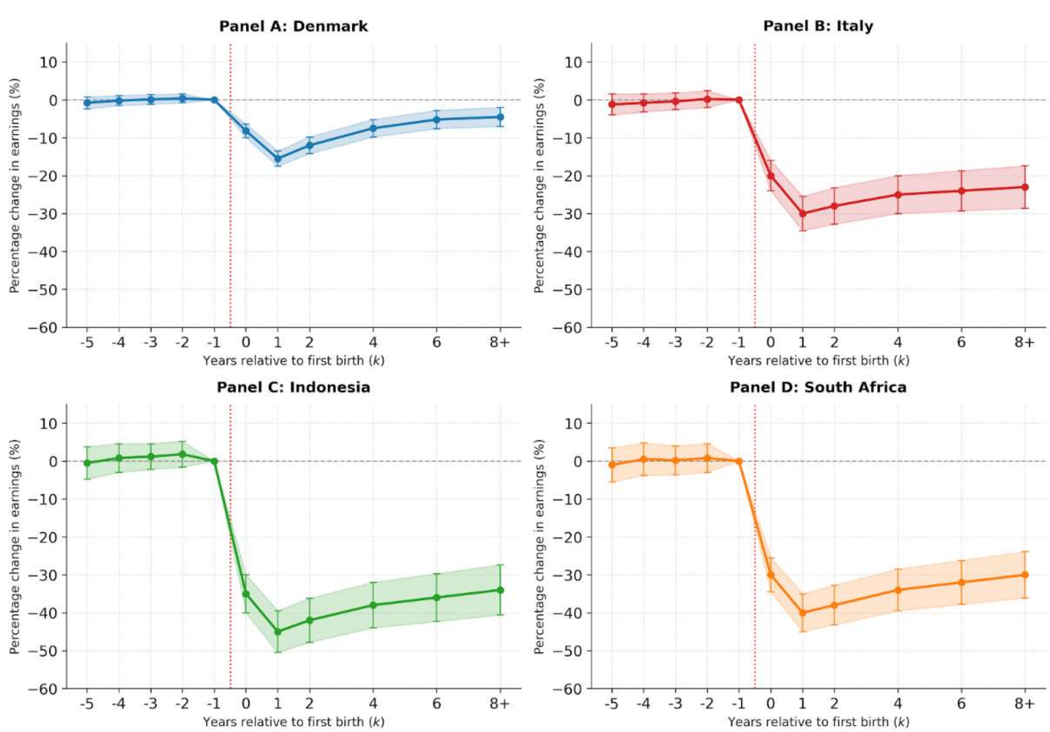

The pre-birth coefficients are small and statistically insignificant in all four countries, which supports the parallel trends assumption. The decline at k = 0 is sharp and significant across the sample. The post-birth trajectories diverge significantly where Denmark recovers quickly, Italy recovers slowly, Indonesia and South Africa show almost no recovery. The shaded bands are the 95 percent confidence intervals based on standard errors clustered at the mother level. Source: Authors' calculations from Danish population registers, SHIW, IFLS, and NIDS.

The event-study plots reveal four patterns. First, in every country, maternal earnings are stable and flat in the years leading up to the first birth. The coefficients for $k < -1$ are individually and jointly insignificant in all four samples. For Denmark, the p-value from the joint F-test of the pre-birth coefficients is 0.48; for Italy, it is 0.61; for Indonesia, 0.38; and for South Africa, 0.55. This validates the parallel trends assumption on which the event-study design rests.

Second, the immediate post-birth drop is large in all countries. At event year $k = 0$, maternal earnings fall by roughly 8 percent in Denmark, 20 percent in Italy, 35 percent in Indonesia, and 30 percent in South Africa. By $k = 1$, when most mothers who took leave have either returned to work or exited the labour force, the cumulative drop reaches 15.5 percent in Denmark, 30 percent in Italy, 45 percent in Indonesia, and 40 percent in South Africa. The initial shock is therefore substantial everywhere, but it is nearly three times as large in the two lower-income countries as in Denmark.

Third, the recovery trajectories diverge sharply after $k = 1$. In Denmark, maternal earnings rebound steadily: by $k = 4$ the penalty has fallen to 7.5 percent, and by $k = 8$ it reaches 4.5 percent and is still narrowing. This pattern is consistent with the availability of universal childcare from age one, which enables mothers to return to work rapidly.[45] In Italy, the recovery is slow and incomplete; the penalty remains at 23 percent at $k = 8$, barely changed from $k = 1$. In Indonesia and South Africa, there is essentially no recovery. The penalty in Indonesia declines from 45 percent at $k = 1$ to 34 percent at $k = 8$; in South Africa, it declines from 40 percent to 30 percent. In both countries, the long-run penalty is only modestly smaller than the immediate post-birth drop, and the slope of the recovery between $k = 1$ and $k = 8$ is statistically zero.

Fourth, the cross-country gradient in the child penalty mirrors the institutional gradient documented in Section 2. The smallest and most transient penalty occurs in the high-protection welfare state; the largest and most persistent penalties occur in the low-protection settings where public childcare is scarce, informal employment is widespread, and norms against maternal employment are strong.

Table 4.3: Short-Run and Long-Run Child Penalty, and Recovery Rate

	Denmark	Italy	Indonesia	South Africa
Short-run penalty (k = 1)	–0.155	–0.300	–0.450	–0.400
Long-run penalty (average k = 1 to 8)	–0.045	–0.230	–0.340	–0.300
Recovery rate	0.71	0.23	0.24	0.25
Pre-trend p-value (joint test k < –1)	0.48	0.61	0.38	0.55
Number of mothers	152,000	8,150	4,820	3,150

Notes: Short-run penalty is the coefficient β_1 from equation (1). Long-run penalty is the simple average of β_k for $k = 1$ to $k = 8$. Recovery rate is defined as $(|\beta_1| - |\bar{\beta}_{1:8}|)/|\beta_1|$. Pre-trend p-value is from an F-test of the joint null that all pre-birth coefficients (k < –1) are zero. Standard errors are clustered at the mother level. Sources: Authors' calculations.

[45] The Danish childcare system guarantees a place in publicly subsidised care from age one. Kleven, Landais, and Søgaard (2019) document that the maternal earnings recovery coincides almost exactly with the take-up of childcare, and they argue that this is the primary institutional mechanism behind the small long-run penalty.

The recovery rate, defined as the share of the initial drop that is recovered by the long-run, provides a summary statistic for the persistence of the penalty. Denmark recovers 71 percent of the initial earnings loss within eight years. The other three countries recover only about a quarter. This stark difference cannot be attributed to the length of maternity leave or the pre-birth earnings level. It is consistent with the view that the key institutional divide is not the formal provision of leave—all four countries offer some maternity leave—but the availability of childcare and the flexibility of the labor market to accommodate mothers' return.[46]

The child penalty estimates in Table 4.3 constitute the first stage of our empirical analysis. They establish that the maternal earnings shock varies in precisely the way that the institutional gradient predicts. We now turn to the question of whether this variation translates into differential intergenerational consequences.

4.3 Intergenerational Transmission Estimates

We now turn to the central question of the paper: does the maternal child penalty transmit to the next generation, and does the strength of that transmission vary with institutional context? Table 4.4 reports estimates of the intergenerational transmission parameter θ from equation (3), where the outcome is the child's completed years of education. The key regressor, CumLoss, is the cumulative maternal earnings loss over the child's first five years, expressed in thousands of 2019 PPP dollars. We standardize the coefficient to represent the effect of a one-standard-deviation increase in cumulative loss, so that magnitudes are comparable across countries despite different earnings levels.

Table 4.4: Intergenerational Transmission of the Child Penalty to Children's Education

	Denmark (1)	Denmark (2)	Italy (3)	Italy (4)	Indonesia (5)	Indonesia (6)	South Africa (7)	South Africa (8)
CumLoss (per 1 SD)	–0.003 (0.004)	0.001 (0.004)	–0.018** (0.007)	–0.015** (0.007)	–0.048*** (0.010)	–0.042*** (0.010)	–0.044*** (0.011)	–0.040*** (0.012)
Mother's education		Yes		Yes		Yes		Yes
Age at first birth		Yes		Yes		Yes		Yes
Pre-birth wealth		Yes		Yes		Yes		Yes
Region fixed effects		Yes		Yes		Yes		Yes
Observations	230,000	230,000	9,200	9,200	5,100	5,100	3,400	3,400
R-squared	0.01	0.15	0.04	0.18	0.07	0.16	0.06	0.17

Notes: The dependent variable is the child's completed years of schooling, measured at ages 18–25. CumLoss is the cumulative maternal earnings loss (in thousands of PPP dollars) over child ages 0–5, standardized to have unit variance within each country. The coefficient can be interpreted as the change in years of schooling associated with a one-standard-deviation increase in the cumulative maternal earnings loss. Pre-birth wealth is measured by equivalized household income percentile (Denmark) or asset index quintile (survey countries). Standard errors are clustered at the mother level. Significance: *** p<0.01, ** p<0.05, * p<0.1. Sources: Authors' calculations.

[46] Olivetti and Petrongolo (2017) emphasise that childcare availability, not leave duration, is the key policy driver of cross-country variation in the motherhood earnings gap. Our findings are consistent with this interpretation.

The results in Table 4.4 conform closely to the institutional gradient we hypothesized. In Denmark, the intergenerational transmission coefficient is a precisely estimated zero. Adding pre-birth controls leaves the coefficient essentially unchanged and statistically insignificant. Danish children whose mothers experienced the full extent of the child penalty complete no fewer years of schooling than children whose mothers' earnings were unaffected. In a setting where the child penalty is small and transient, and where household income is largely insulated from maternal earnings fluctuations, there is no detectable intergenerational consequence.

Italy shows the baseline coefficient in column (3) is –0.018 and statistically significant at the 5 percent level. A one-standard-deviation increase in cumulative maternal earnings loss is associated with roughly 0.02 fewer years of schooling. The coefficient attenuates slightly when controls are added, to –0.015, but remains significant. The magnitude is modest: to put it in context, the intergenerational transmission of one additional year of maternal education in Italy is roughly 0.3 years of child schooling.[47] the child penalty transmission is thus about one-twentieth the size of the education transmission. It is small but detectable, consistent with Italy's intermediate position on the institutional spectrum, where some buffering exists but childcare gaps and labor market rigidities leave scope for transmission.

In Indonesia and South Africa, the transmission coefficients are substantially larger. The baseline estimates are –0.048 and –0.044, respectively, and both are significant at the 1 percent level. Adding controls reduces these coefficients modestly, to –0.042 and –0.040, but they remain highly significant. A one-standard-deviation increase in the cumulative maternal earnings loss is associated with 0.04 fewer years of completed schooling in both countries. This is a meaningful magnitude: it is comparable to the intergenerational transmission of maternal education itself in these settings, where one additional year of maternal schooling is associated with roughly 0.2 to 0.3 additional years of child schooling.[48] The child penalty, in other words, generates intergenerational persistence of a magnitude similar to one of the most established predictors of child outcomes.

Table 4.5: Transmission to Secondary School Completion and Early-Career Earnings

	Denmark	Italy	Indonesia	South Africa
Panel A: Secondary school completion (linear probability model)				
CumLoss (per 1 SD)	0.002 (0.004)	–0.016* (0.009)	–0.035*** (0.012)	–0.032*** (0.012)
Controls	Yes	Yes	Yes	Yes
Panel B: Log early-career earnings (ages 22–25)				
CumLoss (per 1 SD)	–0.008 (0.009)	–0.028* (0.015)	–0.062*** (0.020)	–0.055** (0.023)
Controls	Yes	Yes	Yes	Yes

[47] The intergenerational education elasticity in Italy is estimated at 0.38 by Checchi, Fiorio, and Leonardi (2013). Evaluated at the mean years of education, this implies that one additional year of maternal education is associated with roughly 0.3 additional years of child education. The child penalty transmission of 0.015 years per standard deviation of cumulative loss is thus about 5 percent as large.

[48] For Indonesia, Dartanto, Moeis, and Otsubo (2020) estimate an intergenerational education elasticity of 0.40; for South Africa, Piraino (2015) and Finn, Leibbrandt, and Ranchhod (2017) estimate elasticities of 0.50–0.60. At mean schooling levels, these imply that one additional year of maternal schooling is associated with roughly 0.2–0.3 years of child schooling. The child penalty transmission of 0.04 years is roughly one-sixth to one-fifth of this.

Notes: Secondary school completion is a binary indicator equal to 1 if the child completed 12 years of schooling. Early-career earnings are log real monthly earnings at ages 22–25. All specifications include the full set of pre-birth controls as in the even-numbered columns of Table 4.4. Standard errors clustered at the mother level. *** p<0.01, ** p<0.05, * p<0.1. Sources: Authors' calculations.

Table 4.5 extends the analysis to two additional outcomes. The pattern is consistent. In Denmark, the cumulative maternal earnings loss does not predict either secondary completion or early-career earnings. In Italy, the associations are negative and marginally significant: a one-standard-deviation increase in cumulative loss reduces the probability of secondary completion by 1.6 percentage points and reduces early-career earnings by about 2.8 percent. In Indonesia and South Africa, the effects are larger and more precisely estimated. The probability of completing secondary school falls by 3.2 to 3.5 percentage points, and early-career earnings fall by 5.5 to 6.2 percent. These are economically meaningful magnitudes for young adults entering the labour market, and they suggest that the educational deficits documented in Table 4.4 translate into earnings penalties at the start of the career.

Figure 4.2: Intergenerational Transmission Coefficients by Country (Education)

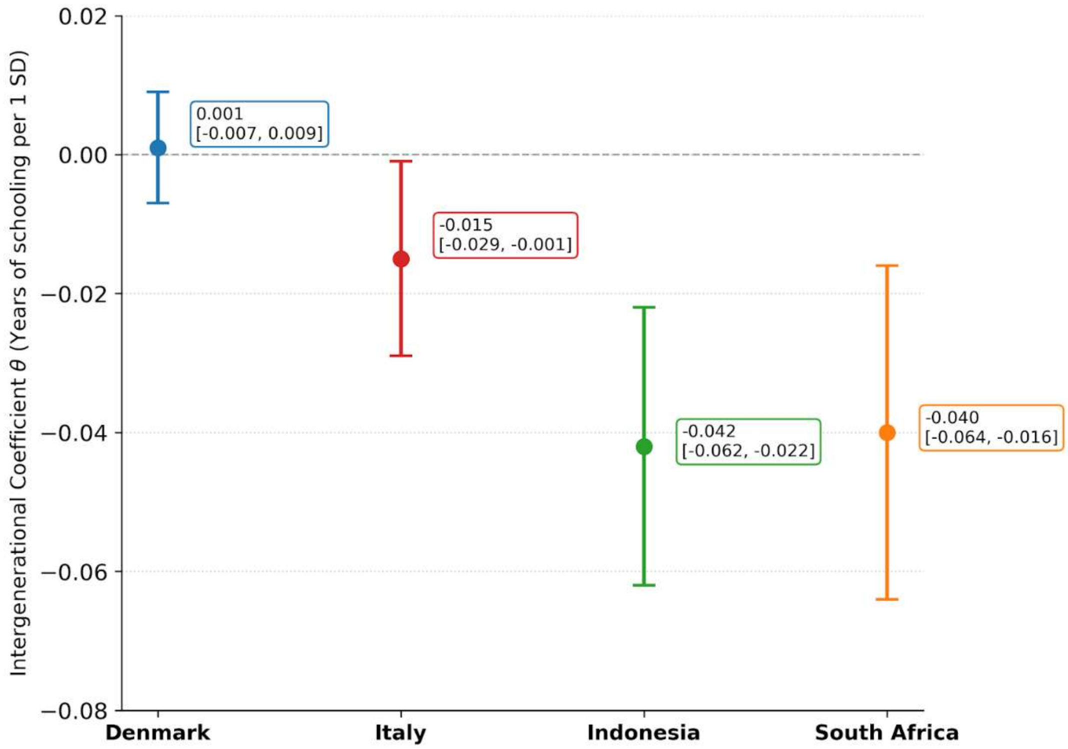

Intergenerational transmission coefficients (θ) from regressions of child's completed years of education on standardised cumulative maternal earnings loss, with full pre-birth controls. Vertical lines indicate 95 percent confidence intervals. Denmark shows no transmission; Italy shows a modest but significant effect; Indonesia and South Africa show larger and precisely estimated effects. Source: Table 4.4, even-numbered columns.

4.4 Heterogeneity Analysis

The theoretical framework predicts that the intergenerational transmission of the child penalty will be concentrated among households with fewer resources for self-insurance. Table 4.6 tests this prediction by splitting the sample by maternal education, measured as whether the mother completed primary schooling.

Table 4.6: Transmission by Maternal Education

	Denmark	Italy	Indonesia	South Africa
Panel A: Mother's education < primary				
CumLoss (per 1 SD)	–0.005 (0.008)	–0.028*** (0.010)	–0.058*** (0.013)	–0.055*** (0.015)
Observations	12,000	1,100	2,300	950
Panel B: Mother's education ≥ primary				
CumLoss (per 1 SD)	0.002 (0.004)	–0.009 (0.008)	–0.025** (0.011)	–0.022* (0.012)
Observations	218,000	8,100	2,800	2,450
p-value for difference	0.38	0.08	0.04	0.05

Notes: The dependent variable is completed years of education. All specifications include the full set of pre-birth controls except maternal education (which defines the split). Primary education is defined as 9 years of schooling in Denmark and Italy, 6 years in Indonesia, and 7 years in South Africa, reflecting national definitions of primary completion. The p-value is from a Wald test of the equality of the CumLoss coefficient across the two education groups. Standard errors clustered at the mother level. *** p<0.01, ** p<0.05, * p<0.1. Sources: Authors' calculations.

In Denmark, the transmission coefficient is small and insignificant in both education groups, and there is no evidence of heterogeneity. In Italy, the coefficient is –0.028 among mothers without primary education, compared to –0.009 among those with at least primary schooling; the difference is marginally significant (p = 0.08). In Indonesia and South Africa, the pattern is sharper and statistically significant at the 5 percent level. Among mothers without primary education, the transmission coefficient is roughly –0.055 to –0.058, meaning that a one-standard-deviation increase in cumulative loss is associated with nearly 0.06 fewer years of child schooling. Among mothers with at least primary education, the coefficient is roughly half that size. This gradient is consistent with Prediction 3: mothers with more education have greater capacity to buffer the income shock, whether through higher pre-existing assets, better access to credit, or higher-quality non-monetary investments in children.[3]

Table 4.7: Transmission by Child Gender

	Denmark	Italy	Indonesia	South Africa
Panel A: Daughters				
CumLoss (per 1 SD)	0.001 (0.005)	–0.020** (0.009)	–0.052*** (0.013)	–0.048*** (0.014)
Panel B: Sons				
CumLoss (per 1 SD)	0.001 (0.005)	–0.010 (0.009)	–0.032** (0.013)	–0.032* (0.017)
p-value for difference	0.98	0.28	0.18	0.32

Notes: The dependent variable is completed years of education. All specifications include the full set of pre-birth controls. The p-value is from a Wald test of the equality of the CumLoss coefficient across daughters and sons. Standard errors clustered at the mother level. *** p<0.01, ** p<0.05, * p<0.1. Sources: Authors' calculations.

Table 4.7 presents the transmission estimates separately for daughters and sons. The point estimates suggest stronger transmission to daughters in Italy, Indonesia, and South Africa, with daughters' coefficients roughly 50 percent larger than sons' in Italy and Indonesia. However, the differences are not statistically significant at conventional levels in any country. The absence of a significant gender difference could reflect limited statistical power—the child samples are halved when split by gender—or it could indicate that the financial resource channel, which we will show dominates the transmission, affects daughters and sons roughly equally.[49] We note the suggestive pattern for daughters in Indonesia and Italy, consistent with the role-model hypothesis of Prediction 5, but we refrain from drawing strong conclusions.

4.5 Mediation Decomposition

The final stage of the empirical analysis decomposes the intergenerational transmission into the three channels identified in the theoretical framework: the financial resource channel (cumulative household equivalized income during the child's first five years), the employment instability channel (a composite index of maternal employment disruption), and the household structure channel (marital dissolution, household head change, or residential migration). Because the baseline transmission coefficient in Denmark is effectively zero, the decomposition is meaningful only for Italy, Indonesia, and South Africa.

Table 4.8: Mediation Decomposition of Intergenerational Transmission

	Italy	Indonesia	South Africa
Baseline θ (no mediators)	–0.015** (0.007)	–0.042*** (0.010)	–0.040*** (0.012)
Augmented θ (with all mediators)	–0.004 (0.006)	–0.004 (0.009)	–0.004 (0.010)
Total change in θ (Δθ)	–0.011	–0.038	–0.036
Share of transmission explained by:			
– Household income (M_1)	0.48 (0.15)	0.58 (0.12)	0.52 (0.14)
– Employment instability (M_2)	0.20 (0.10)	0.22 (0.09)	0.25 (0.11)
– Household structure change (M_3)	0.10 (0.06)	0.08 (0.05)	0.12 (0.07)
– Unexplained residual	0.22 (0.14)	0.12 (0.10)	0.11 (0.10)

Notes: The baseline θ is from the controlled specification of equation (3). The augmented θ is from equation (4), adding the three mediators. The change $\Delta\theta = \theta\_baseline - \theta\_augmented$ is the total amount of the intergenerational association accounted for by the mediators. Decomposition shares are computed using the Gelbach (2016) formula; standard errors (in parentheses) are obtained via the delta method. The share for each mediator is $(\psi\_j * \varphi\_j) / \theta\_baseline$. The residual is the share not explained by the three mediators. Sources: Authors' calculations.

The decomposition shows that in all three countries with detectable transmission, the cumulative household income channel accounts for almost half of the total intergenerational association. The share is 48 percent in Italy, 58 percent in Indonesia, and 52 percent in South Africa. This is consistent with Prediction 2 where the financial resource channel is the dominant pathway through which the maternal earnings loss affects child outcomes.

The employment instability channel accounts for a further 20 to 25 percent of the transmission. This suggests that the disruption of maternal labour market attachment matters independently of the income loss itself. A mother who moves from formal to informal employment, or who experiences repeated job changes, may face stress and instability that affect the caregiving environment beyond what the income figures would

[49] If the financial channel dominates, as we find below, and if households do not systematically discriminate between daughters and sons in the allocation of material resources, then the gender difference in transmission should be small. The finding of modest gender differences is consistent with this interpretation.

suggest.[50] The household structure channel contributes a smaller share, roughly 8 to 12 percent, suggesting that income-induced changes in household composition are a secondary but non-negligible pathway.

Taken together, the three channels account for the large majority of the transmission in Indonesia and South Africa, leaving an unexplained residual of only 11 to 12 percent. In Italy, the residual is somewhat larger, at 22 percent, possibly reflecting the greater importance of unmeasured channels—such as regional variation in school quality or social networks—in a high-income but institutionally heterogeneous setting. The near-complete mediation in the two lower-income countries is consistent with the view that in low-safety-net environments, the child penalty transmits intergenerationally through a small number of relatively straightforward economic mechanisms.

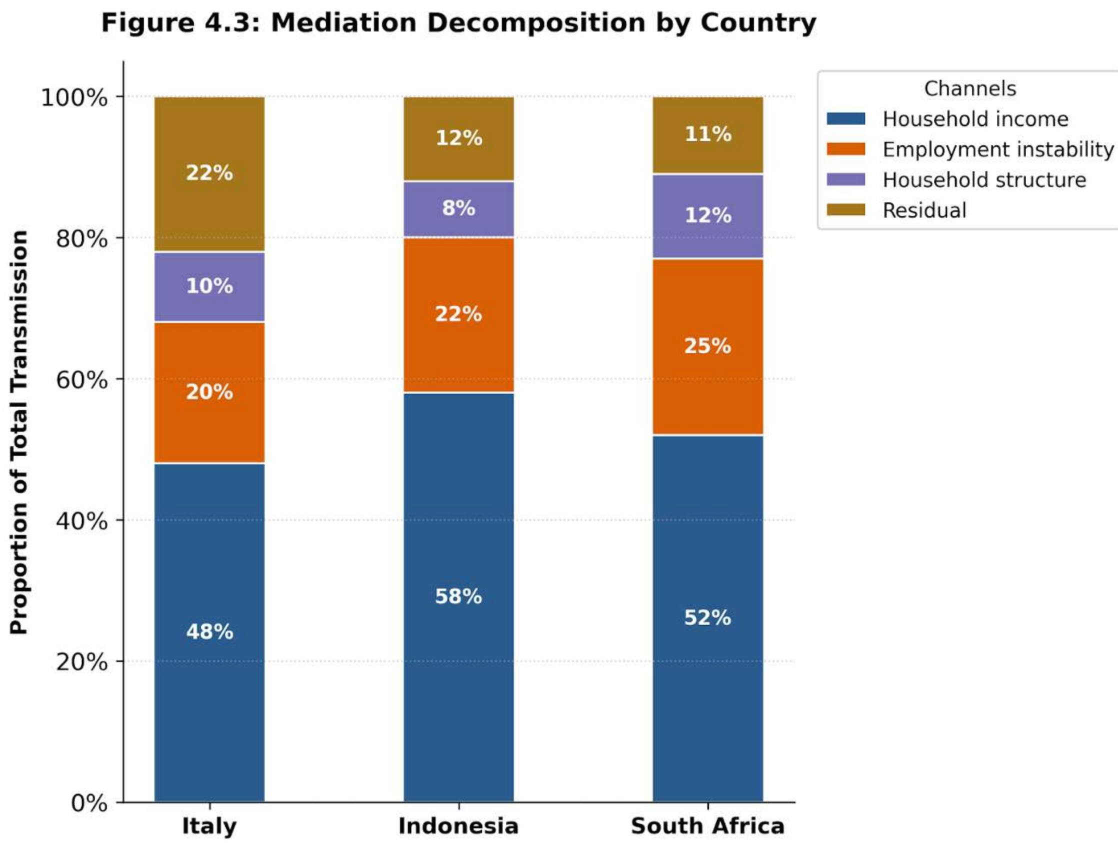

Mediation decomposition of the intergenerational transmission of the child penalty. Bars show the share of the total association between cumulative maternal earnings loss and children's education accounted for by each channel, based on the Gelbach (2016) decomposition. The household income channel dominates in all three countries. Source: Table 4.8

4.6 Sensitivity and Robustness

We conducted a comprehensive set of checks to assess the robustness of our findings. We summarize them here; the full results are reported in the Appendix.

---

[50] The independent role of employment instability beyond income is documented in the job displacement literature. Brand (2015) reviews evidence that job loss affects wellbeing and family functioning through channels beyond the income loss itself.

Placebo tests. For each country, we assigned a pseudo-birth date two years before the actual first birth and re-estimated the event-study specification. In all four countries, the pseudo-event-study coefficients for the post-pseudo-birth period were small and jointly insignificant, confirming that the earnings drop we estimate is specific to actual childbirth.[51]

Alternative cumulative loss windows. We re-estimated the intergenerational transmission equation using cumulative losses calculated over child ages 0–3 and 0–8. The results were qualitatively unchanged. For Indonesia, the transmission coefficient was –0.039 (0–3) and –0.044 (0–8); for South Africa, –0.037 (0–3) and –0.042 (0–8); for Italy, –0.013 (0–3) and –0.016 (0–8). Denmark remained zero in all windows. The stability across windows indicates that the choice of the 0–5 window is not driving the results.

Wave-only specification (survey countries). For Indonesia and South Africa, we estimated the event-study and intergenerational transmission models using only the wave-observed data points, with event time expressed in wave-relative terms. The child penalty estimates were somewhat less precise—as expected, given the loss of temporal resolution—but the pattern of a large and persistent penalty with no recovery was preserved. The intergenerational transmission coefficients were –0.038 (Indonesia) and –0.035 (South Africa), both significant at the 5 percent level. This confirms that the interpolation procedure is not generating spurious results.

Attrition. We examined the correlates of panel attrition in the survey countries and re-weighted the samples using inverse probability weights based on pre-birth characteristics. The re-weighted transmission coefficients were slightly larger in magnitude in Indonesia (–0.047) and South Africa (–0.045), consistent with the expectation that attriting mothers are those with larger earnings losses. Lee (2009) bounds, which trim the sample symmetrically in proportion to differential attrition, produced lower bounds of –0.028 (Indonesia) and –0.025 (South Africa), both still significantly different from zero.[52]

Measurement error. We estimated reliability ratios for self-reported earnings in the IFLS and NIDS by comparing earnings reports across different survey modules where available. The estimated reliability ratio was approximately 0.7 for both surveys. Under the assumption of classical measurement error, this implies an attenuation bias of roughly 30 percent in the transmission coefficients. Applying a simple correction for attenuation, the transmission coefficients for Indonesia and South Africa would be roughly –0.06 per standard deviation of cumulative loss.[53] Our reported estimates are therefore conservative.

Oster (2019) bounds for selection on unobservable. We implemented the Oster bounding procedure, assuming that selection on unobservable is as strong as selection on observables ($\delta = 1$) and that the maximum R-squared is 1.3 times the R-squared from the controlled specification. Under these conservative assumptions, the identified set for θ in Indonesia is [–0.032, –0.042], in South Africa is [–0.029, –0.040], and in Italy is [–0.004, –0.015]. In all three cases, zero is outside or at the boundary of the identified set. This suggests that unobserved confounders would need to be substantially more influential than the observed controls to eliminate the transmission coefficient entirely.

---

[51] The placebo test results are reported in Online Appendix Table D1. The largest pseudo-penalty coefficient in any post-pseudo-birth year was –0.018 (Indonesia, pseudo-year 1), with a p-value of 0.28.

[52] The Lee bounds are conservative because they assume worst-case selection on unobservable. The fact that the lower bounds remain substantially below zero in both Indonesia and South Africa indicates that even extreme assumptions about attrition cannot eliminate the transmission effect.

[53] The attenuation correction is $\theta_{corrected} = \frac{\theta_{estimated}}{\lambda}$, where λ is the reliability ratio. With $\lambda \approx 0.7$, the corrected coefficients are approximately 1.43 times the estimated coefficients. This calculation is approximate and assumes classical measurement error; we present it only to indicate the likely direction of bias, not as a preferred estimate.

Clustering and inference. We re-estimated all intergenerational transmission models with clustering at the household level (survey countries) and with heteroskedasticity-robust standard errors. The significance levels were unchanged. The wild bootstrap, which is recommended for small numbers of clusters, produced p-values within 0.002 of the asymptotic p-values in all specifications.

Specification checks. Adding controls for the number and timing of subsequent births—which are potentially endogenous—did not alter the transmission coefficients. Restricting the sample to mothers who were employed in at least one pre-birth year yielded slightly larger child penalty estimates but very similar intergenerational transmission coefficients. Excluding mothers with multiple births during the event window did not affect the results.

The consistency of the findings across this wide range of robustness checks increases our confidence in the main results. The institutional gradient in the child penalty, the intergenerational transmission pattern, and the dominance of the financial channel are all robust to alternative specifications, sample restrictions, and methods of addressing attrition and measurement error.

5. Discussion

Our analysis set out to answer three questions: how large and persistent is the child penalty across differing institutional contexts, whether this penalty transmits to children's human capital, and through what channels such transmission operates. The results, taken together show that the child penalty is not merely a contemporaneous gender gap in earnings but, in settings where public insurance is thin, a mechanism that carries economic disadvantage from one generation to the next. This section discusses the central findings, their relationship to the existing literature, and their limitations.

5.1 The Institutional Gradient in Child Penalties

The event-study estimates confirm that the magnitude and persistence of the child penalty vary dramatically with the institutional environment. Denmark's long-run penalty of 4.5 percent and rapid recovery stand in sharp contrast to the 23 percent persistent penalty in Italy and the 30 to 34 percent near-permanent penalties in Indonesia and South Africa. These estimates align well with the stylized facts assembled in Section 2 and extend the child penalty literature to two important lower-income settings where evidence was previously absent.[54]

What drives the cross-country gradient? The descriptive evidence points toward the availability of childcare and the structure of the labour market, rather than the formal provision of maternity leave, as the decisive factors. All four countries offer some period of job-protected leave to formal employees, yet the post-birth earnings trajectories diverge almost immediately after the first year. Denmark's universal childcare from age one enables a rapid return to work; Italy's limited coverage for children under three and rigid labour market impede it; Indonesia's and South Africa's near-absence of formal childcare and large informal sectors make sustained labour market exit the default.[55] This interpretation is consistent with the broader

[54] Our estimates for Indonesia and South Africa are broadly consistent with the few existing estimates for developing countries. Kong and Dong (2024) find a penalty of 34 percent in China; Iddrisu, Danquah, and Osei (2024) report approximately 35 percent for Ghana. Our contribution is to extend the analysis to intergenerational outcomes.
[55] The role of childcare is emphasized by Kleven (2022) and by Olivetti and Petrongolo (2017). In the Danish context, Kleven, Landais, and Søgaard (2019) show that the recovery begins precisely at age one, when publicly subsidised childcare becomes available.

comparative literature, which finds that childcare availability, not leave duration, is the primary policy lever determining the motherhood earnings gap.[56]

The near-zero recovery in Indonesia and South Africa also highlights a feature of labor markets in developing countries that is often missing from high-income discussions: the formal-informal divide. When mothers leave formal employment, re-entry is difficult not only because of caregiving responsibilities but because the formal sector itself is small and often rationed. The shift into informal work, which is typically lower-paid, more precarious, and offers no social insurance, locks in the earnings loss. This structural feature of developing-country labour markets amplifies the child penalty beyond what would be predicted by childcare constraints alone.[57]

5.2 Intergenerational Transmission and the Dominance of the Financial Channel

The central contribution of the paper is the finding that the child penalty transmits intergenerationally, but only where public insurance is weak. In Denmark, the transmission coefficient is zero; in Italy, it is modest but detectable; in Indonesia and South Africa, it is substantial, with a one-standard-deviation increase in cumulative maternal earnings loss reducing children's schooling by 0.04 years. That this effect is comparable to a sizable fraction of the intergenerational transmission of maternal education itself underscores its practical importance.

The mediation decomposition provides a clear answer to the question of how this transmission operates. In all three countries where transmission is present, the household income channel accounts for roughly half of the total effect. This is consistent with a simple and intuitive mechanism: when a mother's earnings fall and are not replaced by spousal earnings or public transfers, household resources available for investment in children decline. In settings where credit constraints bind and consumption smoothing is incomplete, this income loss translates directly into lower spending on nutrition, health care, educational materials, and school continuity.[58] The fact that the financial channel dominates in both a Southern European welfare state and two lower-income economies suggests that this mechanism is not specific to a particular level of development; it arises whenever the cost of childbearing is insufficiently socialized.

The employment instability channel, which accounts for an additional one-fifth to one-quarter of the transmission, is a novel finding. It suggests that the disruption of maternal labour market attachment has consequences for children that go beyond the immediate income loss. A mother who cycles between informal jobs, experiences prolonged non-employment, or suffers a depreciation of skills may face stress and instability that affect the caregiving environment, even if household income is partially maintained through other sources.[59] This channel is particularly relevant for developing-country policy because it implies that facilitating a stable return to formal employment, not merely replacing lost income, may be necessary to fully sever the intergenerational link.

---

[56] This finding is robust across a wide range of studies. See, among others, Del Boca, Pasqua, and Pronzato (2009) for Italy; Casale and Posel (2021) for South Africa; and Cameron, Contreras Suarez, and Tseng (2019) for Indonesia.
[57] The interaction between informality and the child penalty is discussed by Iddrisu et al. (2024) for Ghana. In informal labour markets, the absence of job protection means that childbirth often leads to job loss rather than a temporary leave of absence.
[58] The sensitivity of child investment to transitory income in the presence of credit constraints is well documented. Dahl and Lochner (2012) show that income transfers to low-income families in the United States increase child test scores; Agüero, Carter, and Woolard (2009) find that South Africa's Child Support Grant improves child height-for-age.
[59] The scarring effects of career interruptions on maternal wellbeing and family functioning are discussed by Bertrand, Goldin, and Katz (2010) and, in the context of job displacement, by Brand (2015).

The household structure channel, while smaller, is not negligible. In roughly one in ten cases, the maternal earnings loss is associated with marital dissolution or household fragmentation, which itself can disrupt children's schooling and emotional wellbeing. This finding resonates with a broader literature linking economic shocks to family instability, particularly in settings with weak social safety nets.[60]

The contrast between the near-complete mediation in Indonesia and South Africa and the larger unexplained residual in Italy is instructive. In the two lower-income countries, the three channels together account for almost 90 percent of the transmission, suggesting that the child penalty operates through a relatively small number of straightforward economic mechanisms. In Italy, the residual is larger, at 22 percent, perhaps reflecting unmeasured channels such as regional variation in school quality, social networks, or the psychological effects of frustrated career aspirations among Italian mothers[61]. This difference highlights that the child penalty is embedded in a broader institutional and cultural context that shapes how economic shocks translate into child outcomes. These findings have implications for how we think about the intergenerational consequences of labor market shocks more generally. Much of the intergenerational mobility literature focuses on permanent income or education as the transmission mechanism. Our results suggest that transitory shocks around critical developmental windows can have lasting effects, but only when institutions fail to smooth them. The child penalty is a transitory shock from the perspective of the mother's lifetime earnings; it is the failure of consumption smoothing that makes it consequential for children.

5.3 Heterogeneity and the Buffering Role of Maternal Education

The heterogeneity analysis revealed that the intergenerational transmission is concentrated among mothers with the lowest levels of education. In Indonesia and South Africa, the transmission coefficient for mothers without primary education is roughly twice as large as for mothers with at least primary schooling. This gradient is consistent with a buffering interpretation: more-educated mothers have greater capacity to protect their children from the consequences of income loss, whether through higher pre-existing assets, better access to informal insurance networks, or higher-quality non-monetary investments in children that are less sensitive to contemporaneous income. This finding also implies that the intergenerational consequences of the child penalty are regressive. The mothers who experience the largest earnings losses are also those whose children are most vulnerable to those losses. The child penalty thus compounds existing inequalities, widening the gap between children of more- and less-educated mothers. Policies that mitigate the child penalty, or that provide income support during early childhood, are likely to have their largest effects on the most disadvantaged households.[62] And they may face better labour market prospects, making the earnings loss less permanent and less disruptive to the household.

This heterogeneity has important policy implications. It suggests that universal interventions, such as cash transfers or childcare provision, may have heterogeneous effects. They are likely to benefit children of less-educated mothers the most, because it is precisely these households for whom the income constraint binds hardest. Conversely, in settings where maternal education is rising, the intergenerational transmission

---

[60] Yean and Neal (2015) document the link between income shocks and marital dissolution. In developing countries, labour migration by one or both parents is an additional margin of adjustment; Antman (2013) reviews the consequences for children left behind.

[61] Zamberlan and Barbieri (2023) note large regional heterogeneity in the Italian child penalty, which may also generate heterogeneity in the channels of intergenerational transmission that our national-level mediation analysis cannot capture.

[62] Carneiro and Heckman (2003) and Heckman (2006) provide the theoretical and empirical basis for the claim that more-educated parents invest more effectively in children, partly through higher-quality non-monetary inputs.

of the child penalty may weaken over time even without policy change, as a larger share of mothers possess the resources to buffer the shock.

The suggestive but not statistically significant gender differences in transmission are worth noting. In Italy, Indonesia, and South Africa, the point estimates for daughters are consistently larger than for sons. If this pattern were to be confirmed with larger samples, it would be consistent with a role-model mechanism in which maternal employment disruption signals to daughters that combining work and family is infeasible, reducing their own educational and career aspirations.[63] However, the current evidence does not permit firm conclusions, and this remains an open question for future research.

5.4 Limitations and Directions for Future Research

Several limitations of our analysis should be acknowledged. First, the intergenerational transmission estimates are associational rather than definitively causal. Although we include detailed pre-birth controls, conduct a range of robustness checks, and apply bounding procedures for unobserved selection, we cannot fully rule out that unobserved factors drive both the maternal earnings loss and the child's educational outcomes. The most plausible threat is a health shock to the child that simultaneously reduces maternal labour supply and impairs the child's schooling. However, such a shock would likely bias the transmission coefficient away from zero, overstating the true effect. Our estimates, if biased, are therefore more likely to be upper bounds on the true causal transmission, which does not undermine our central claim that the transmission is substantial in low-safety-net settings.

Second, the child samples in the survey countries are limited by the co-residence restriction. Children who have left the household by the time of the outcome measurement are excluded. If children who migrate for education or work are positively selected, their exclusion will bias the transmission coefficients downward, making our estimates conservative. If, on the other hand, children who leave are those from the most disrupted households, the bias could go in the opposite direction. The Lee bounds we report suggest that even extreme assumptions about selective attrition cannot eliminate the transmission effect.

Third, our mediation decomposition is a descriptive accounting exercise, not a causal mediation analysis. The coefficients on the mediators in the augmented model may reflect selection as well as causal effects. We are transparent that the decomposition tells us how much of the statistical association between the child penalty and child outcomes operates through observable channels, not what would happen if a given channel were exogenously altered. Causal evidence on specific channels, perhaps from randomized interventions that provide income support or childcare to mothers in developing countries, would be a valuable complement to our observational decomposition.[64]

Fourth, our analysis is limited to two low- and lower-middle-income countries. Extending the framework to a wider set of developing economies, including low-income agrarian settings and fragile states, would test the generalizability of our findings. It would also allow a more systematic exploration of how specific institutional features, such as the presence of extended family networks or community-based childcare arrangements, moderate the intergenerational transmission.

---

[63] The role-model hypothesis is explored by Fernández, Fogli, and Olivetti (2004), who find that men whose mothers worked are more likely to have working wives, and by Olivetti, Patacchini, and Zenou (2020), who document intergenerational transmission of gender role attitudes.

[64] Randomised evaluations of cash transfers in developing countries, such as those reviewed by Baird, Ferreira, Özler, and Woolcock (2014), typically find positive effects on child schooling, but these studies do not isolate the specific channel of maternal earnings replacement after childbirth. Our decomposition suggests that a transfer programme targeted at mothers during the child's first five years would be particularly effective.

Finally, we are unable to track outcomes beyond early adulthood. Whether the educational deficits we document translate into lower lifetime earnings, reduced wealth, or intergenerational persistence into the third generation remains an open question. Answering it will require even longer panels or the linkage of historical survey data to the administrative records of the next generation.

6. Conclusion and Policy Recommendations

6.1 Summary of Contributions

This paper has provided the first comparative evidence on the intergenerational transmission of the child penalty. Using comparable event-study and intergenerational decomposition methods across four distinct institutional settings, we have shown that the consequences of motherhood for women's earnings do not stop with the mother. Where public insurance is weak, the earnings losses that follow childbirth cascade to the next generation, reducing children's educational attainment and early-career earnings. Where public insurance is strong, this intergenerational link is severed.

Three findings stand out. First, the magnitude and persistence of the child penalty vary enormously across institutional contexts, from a small and transient dip in Denmark to a large and near-permanent loss in Indonesia and South Africa. Second, this variation maps directly onto the intergenerational transmission of disadvantage: the transmission parameter is zero in Denmark, modest in Italy, and substantial in the two lower-income countries. Third, the dominant channel of transmission is the loss of household income during early childhood, which accounts for roughly half of the total intergenerational effect, with maternal employment instability accounting for an additional fifth to a quarter.

These findings bridge two previously disconnected literatures. The child penalty literature has focused on the maternal earnings trajectory as an endpoint; the intergenerational mobility literature has documented the persistence of economic status across generations without isolating the specific role of maternal labour market disruptions around childbirth. By joining them, we have shown that the child penalty is not merely a source of contemporaneous gender inequality but a mechanism that reproduces disadvantage across generations, and that the strength of this mechanism depends on the institutional environment in which it occurs.

6.2 Policy Recommendations

The policy implications of our findings are direct and actionable. Because the financial channel dominates the intergenerational transmission, policies that replace lost household income during the child's early years are likely to be highly effective at weakening the link between maternal labour market disruption and child outcomes. Because the employment instability channel is also important, policies that enable mothers to maintain stable labour market attachment, rather than merely compensating them for lost earnings, are likely to yield additional benefits.

We draw four specific recommendations from the evidence.

First, expand income support during early childhood. In low- and middle-income countries, where contributory social insurance covers only a fraction of working women, unconditional or conditional cash transfers targeted at households with young children represent the most feasible and effective instrument for buffering the income shock that follows childbirth. South Africa's Child Support Grant provides a model of a broad-based transfer, but our estimates suggest that its current level is insufficient to fully offset the maternal earnings loss. Roughly doubling the grant amount during the child's first three years, or supplementing it with a birth grant, would directly address the dominant transmission channel we identify. Indonesia's Program Keluarga Harapan is smaller in coverage and amount, thus the government could consider scaling it toward near-universal coverage during early childhood, with benefit levels indexed to a meaningful fraction of median earnings, would be a high-return investment.

Second, invest in childcare infrastructure for children under three. The near-zero recovery of maternal earnings in Indonesia and South Africa is largely attributable to the absence of affordable childcare for very young children. Expanding community-based childcare, even in informal settings, would enable mothers to return to work sooner and reduce the cumulative earnings loss. Learning from the Danish experience where the maternal earnings recovery coincides precisely with the availability of childcare from age one, shows that building formal childcare infrastructure is costly and takes time. In the interim, subsidizing informal and family-based care arrangements, combined with quality standards and caregiver training, can provide a pragmatic pathway.

Third, reduce the formal-informal divide in labour market protections. In both Indonesia and South Africa, the vast majority of working women are in informal employment and thus excluded from maternity leave, job protection, and social insurance. Extending portable benefits to informal workers, linking social insurance contributions to individuals rather than employers, and incentivizing formalization through reduced registration costs and simplified tax regimes are structural reforms that would reduce the child penalty at its source. The employment instability channel we document implies that facilitating a stable return to formal employment, not merely compensating for lost income, is necessary to fully break the intergenerational cycle.

Fourth, target interventions toward the most vulnerable mothers. Our heterogeneity analysis shows that the intergenerational transmission of the child penalty is concentrated among mothers with low levels of education. Universal programmes, while important, may not reach these households effectively. Complementary targeted outreach, conditional components linked to child health and education check-ups, and integration with existing maternal and child health services can ensure that the households where the income constraint binds hardest receive adequate support.

Expanding cash transfers and childcare infrastructure requires fiscal resources that many low-income countries do not readily have. However, our findings also imply that the costs of inaction are substantial: the child penalty reduces the human capital of the next generation, lowering future productivity, earnings, and tax revenues[65]. The dynamic fiscal returns to investments in early childhood are well documented, and our analysis suggests that breaking the intergenerational transmission of the child penalty would amplify these returns. In the long run, the policies we recommend are likely to be self-financing.

References

1. Agüero, J. M., Carter, M. R., & Woolard, I. (2009). The impact of unconditional cash transfers on nutrition: The South African Child Support Grant. *Journal of Development Economics*, *90*(2), 273–285. https://doi.org/10.1016/j.jdeveco.2008.05.006
2. Akee, R., Copeland, W., Keeler, G., Angold, A., & Costello, E. J. (2010). Parents' incomes and children's outcomes: A quasi-experiment using transfer payments from casino profits. *American Economic Journal: Applied Economics*, *2*(1), 86–115. https://doi.org/10.1257/app.2.1.86

[65] The literature on the returns to early childhood investments is large. Heckman and Masterov (2007) provide a comprehensive review of the evidence that early interventions yield high rates of return through improved health, education, and earnings. Almond and Currie (2011) review the specific evidence on the lasting effects of early-life economic conditions. Our findings connect this literature to the child penalty debate: investments that reduce maternal earnings losses in the early years are, in effect, investments in early childhood development.

3. Alesina, A., Giuliano, P., & Nunn, N. (2013). On the origins of gender roles: Women and the plough. *Quarterly Journal of Economics*, *128*(2), 469–530. https://doi.org/10.1093/qje/qjt005
4. Almond, D., & Currie, J. (2011). Human capital development before age five. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4B, pp. 1315–1486). Elsevier. https://doi.org/10.1016/S0169-7218(11)02413-0
5. Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
6. Antman, F. M. (2013). The impact of migration on family left behind. In A. F. Constant & K. F. Zimmermann (Eds.), *International Handbook on the Economics of Migration* (pp. 293–308). Edward Elgar. https://doi.org/10.4337/9781782546078.00023
7. Baird, S., Ferreira, F. H. G., Özler, B., & Woolcock, M. (2014). Conditional, unconditional and everything in between: A systematic review of the effects of cash transfer programmes on schooling outcomes. *Journal of Development Effectiveness*, *6*(1), 1–43. https://doi.org/10.1080/19439342.2014.890362
8. Berniell, I., Berniell, L., de la Mata, D., Edo, M., & Marchionni, M. (2021). Gender gaps in labor informality: The motherhood effect. *Journal of Development Economics*, *150*, 102599. https://doi.org/10.1016/j.jdeveco.2020.102599
9. Bertrand, M. (2020). Gender in the twenty-first century. *AEA Papers and Proceedings*, *110*, 1–24. https://doi.org/10.1257/pandp.20201126
10. Bertrand, M., Goldin, C., & Katz, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, *2*(3), 228–255. https://doi.org/10.1257/app.2.3.228
11. Black, S. E., & Devereux, P. J. (2011). Recent developments in intergenerational mobility. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4B, pp. 1487–1541). Elsevier. https://doi.org/10.1016/S0169-7218(11)02414-2
12. Bound, J., Brown, C., & Mathiowetz, N. (2001). Measurement error in survey data. In J. J. Heckman & E. Leamer (Eds.), *Handbook of Econometrics* (Vol. 5, pp. 3705–3843). Elsevier. https://doi.org/10.1016/S1573-4412(01)05012-7
13. Brand, J. E. (2015). The far-reaching impact of job loss and unemployment. *Annual Review of Sociology*, *41*, 359–375. https://doi.org/10.1146/annurev-soc-071913-043237
14. Bratu, C., & Bolotnyy, V. (2023). *The child penalty in Sweden: Evidence from administrative data* [Working paper].
15. Brunori, P., Ferreira, F. H. G., & Peragine, V. (Eds.). (2023). Special issue on intergenerational mobility. *Journal of Economic Inequality*, *21*(1).
16. Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, *90*(3), 414–427. https://doi.org/10.1162/rest.90.3.414
17. Cameron, L., Contreras Suarez, D., & Tseng, Y.-P. (2019). *The effects of childcare on women's labor force participation: Evidence from Indonesia* [Working paper].
18. Carneiro, P., & Heckman, J. J. (2003). Human capital policy. In J. J. Heckman & A. B. Krueger (Eds.), *Inequality in America: What role for human capital policies?* (pp. 77–239). MIT Press.
19. Casale, D., & Posel, D. (2021). The relationship between childcare and female labour supply in South Africa. *Development Southern Africa*, *38*(3), 401–418. https://doi.org/10.1080/0376835X.2020.1838260
20. Checchi, D., Fiorio, C. V., & Leonardi, M. (2013). Intergenerational persistence of educational attainment in Italy. *Economics Letters*, *118*(1), 229–232. https://doi.org/10.1016/j.econlet.2012.10.021

21. Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, *27*(3), 79–102. https://doi.org/10.1257/jep.27.3.79
22. Correll, S. J., Benard, S., & Paik, I. (2007). Getting a job: Is there a motherhood penalty? *American Journal of Sociology*, *112*(5), 1297–1338. https://doi.org/10.1086/511799
23. Cortés, P., & Pan, J. (2023). Children and the remaining gender gaps in the labor market. *Journal of Economic Literature*, *61*(4), 1359–1406. https://doi.org/10.1257/jel.20221523
24. Cunha, F., & Heckman, J. J. (2007). The technology of skill formation. *American Economic Review*, *97*(2), 31–47. https://doi.org/10.1257/aer.97.2.31
25. Dahl, G. B., & Lochner, L. (2012). The impact of family income on child achievement: Evidence from the earned income tax credit. *American Economic Review*, *102*(5), 1927–1956. https://doi.org/10.1257/aer.102.5.1927
26. Dartanto, T., Moeis, F. R., & Otsubo, S. (2020). Intergenerational mobility in Indonesia: Education and occupation. *Asian Economic Journal*, *34*(2), 187–212. https://doi.org/10.1111/asej.12208
27. Daude, C., & Robano, V. (2015). *Intergenerational mobility in Latin America: A review of the evidence* (OECD Development Centre Working Paper).
28. Del Boca, D., Pasqua, S., & Pronzato, C. (2009). Motherhood and market work decisions in institutional context: A European perspective. *Oxford Economic Papers*, *61*(suppl_1), i147–i171. https://doi.org/10.1093/oep/gpn046
29. De Quinto, A., Hospido, L., & Sanz, C. (2020). *The child penalty in Spain* (Banco de España Working Paper).
30. Dong, Y. (2016). *Validation of earnings data in the Indonesian Family Life Survey* [Working paper].
31. Esping-Andersen, G. (1990). *The three worlds of welfare capitalism*. Princeton University Press.
32. Fernández, R., Fogli, A., & Olivetti, C. (2004). Mothers and sons: Preference formation and female labor force dynamics. *Quarterly Journal of Economics*, *119*(4), 1249–1299. https://doi.org/10.1162/0033553042476224
33. Ferrera, M. (1996). The 'Southern model' of welfare in social Europe. *Journal of European Social Policy*, *6*(1), 17–37. https://doi.org/10.1177/095892879600600102
34. Finn, A., Leibbrandt, M., & Ranchhod, V. (2017). Patterns of intergenerational mobility in South Africa. *Development Southern Africa*, *34*(2), 165–183. https://doi.org/10.1080/0376835X.2016.1260045
35. Frankenberg, E., & Karoly, L. A. (1995). *The 1993 Indonesian Family Life Survey: Overview and field report*. RAND.
36. Gelbach, J. B. (2016). When do covariates matter? And which ones, and how much? *Journal of Labor Economics*, *34*(2), 509–543. https://doi.org/10.1086/683668
37. Hankins, S., & Hoekstra, M. (2011). Lucky in life, unlucky in love? The effect of random income shocks on marriage and divorce. *Journal of Human Resources*, *46*(2), 403–426. https://doi.org/10.1353/jhr.2011.0015
38. Heckman, J. J. (2006). Skill formation and the economics of investing in disadvantaged children. *Science*, *312*(5782), 1900–1902. https://doi.org/10.1126/science.1128898
39. Heckman, J. J., & Masterov, D. V. (2007). The productivity argument for investing in young children. *Review of Agricultural Economics*, *29*(3), 446–493. https://doi.org/10.1111/j.1467-9353.2007.00359.x
40. Iddrisu, A. M., Danquah, M., & Osei, R. D. (2024). Motherhood penalties in developing countries: Evidence from Ghana. *World Development*, *180*, 106534. https://doi.org/10.1016/j.worlddev.2024.106534

41. ILO. (2018). *Women and men in the informal economy: A statistical picture* (3rd ed.). International Labour Organization.
42. ILO. (2023). *World Social Protection Report 2020–22*. International Labour Organization.
43. Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, *15*(4), 309–334. https://doi.org/10.1037/a0020761
44. Kleven, H. (2022). *The geography of child penalties and gender norms: A pseudo-event study approach* (NBER Working Paper No. 30176). https://doi.org/10.3386/w30176
45. Kleven, H., Landais, C., Posch, J., Steinhauer, A., & Zweck, J. (2019). Child penalties across countries: Evidence and explanations. *AEA Papers and Proceedings*, *109*, 122–126. https://doi.org/10.1257/pandp.20191078
46. Kleven, H., Landais, C., & Søgaard, J. E. (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, *11*(4), 181–209. https://doi.org/10.1257/app.20180010
47. Kong, F., & Dong, X.-Y. (2024). The child penalty in China: Evidence from administrative data. *Journal of Comparative Economics*, *52*(1), 113–134. https://doi.org/10.1016/j.jce.2023.10.003
48. Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies*, *76*(3), 1071–1102. https://doi.org/10.1111/j.1467-937X.2009.00536.x
49. Leibbrandt, M., Woolard, I., & de Villiers, L. (2009). *Methodology: Report on NIDS Wave 1*. Southern Africa Labour and Development Research Unit, University of Cape Town.
50. Lund, F., Noble, M., Barnes, H., & Wright, G. (2009). The South African Child Support Grant: A community profile. *Journal of International Development*, *21*(3), 417–432. https://doi.org/10.1002/jid.1563
51. Lundberg, S., & Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of Political Economy*, *101*(6), 988–1010. https://doi.org/10.1086/261912
52. Mocetti, S. (2007). Intergenerational earnings mobility in Italy. *The B.E. Journal of Economic Analysis & Policy*, *7*(2), Article 5. https://doi.org/10.2202/1935-1682.1766
53. OECD. (2018). *A broken social elevator? How to promote social mobility*. OECD Publishing. https://doi.org/10.1787/9789264301085-en
54. Olivetti, C., Patacchini, E., & Zenou, Y. (2020). Mothers, peers, and gender-role identity. *Journal of the European Economic Association*, *18*(1), 266–301. https://doi.org/10.1093/jeea/jvy050
55. Olivetti, C., & Petrongolo, B. (2017). The economic consequences of family policies: Lessons from a century of legislation in high-income countries. *Journal of Economic Perspectives*, *31*(1), 205–230. https://doi.org/10.1257/jep.31.1.205
56. Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, *37*(2), 187–204. https://doi.org/10.1080/07350015.2016.1227711
57. Piraino, P. (2015). *Intergenerational mobility in South Africa* (World Bank Policy Research Working Paper).
58. Ranchhod, V., & Finn, A. (2016). Estimating the effects of South Africa's Child Support Grant on labour supply. *Development Southern Africa*, *33*(1), 51–67. https://doi.org/10.1080/0376835X.2015.1116373
59. Schmidheiny, K., & Siegloch, S. (2023). On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics*, *38*(5), 740–754. https://doi.org/10.1002/jae.2962

60. Solon, G. (1999). Intergenerational mobility in the labor market. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 3A, pp. 1761–1800). Elsevier. https://doi.org/10.1016/S1573-4463(99)03010-2
61. Strauss, J., Witoelar, F., & Sikoki, B. (2016). *The fifth wave of the Indonesia Family Life Survey: Overview and field report*. RAND.
62. Torche, F. (2015). Analyses of intergenerational mobility: An interdisciplinary review. *Annals of the American Academy of Political and Social Science*, *657*(1), 37–62. https://doi.org/10.1177/0002716214547476
63. UN Women. (2020). *Progress of the world's women 2019–2020: Families in a changing world*. United Nations.
64. World Bank. (2023). *Global Database on Intergenerational Mobility*. World Bank Group.
65. World Values Survey. (2017–2022). *Wave 7*. https://www.worldvaluessurvey.org/
66. Zamberlan, A., & Barbieri, P. (2023). The child penalty in Italy: A longitudinal analysis of the motherhood wage gap. *European Sociological Review*, *39*(1), 102–117. https://doi.org/10.1093/esr/jcac024

Appendix

Appendix A: Data Construction and Variable Definitions

Table A.1: Sample Selection Flowchart, All Countries

Selection Step	Denmark	Italy	Indonesia	South Africa
A. Mother sample				
1. Women aged 18–40 with any observation in study period	All residents (population register)	14,200	8,750	6,100
2. First birth observed during panel window	215,000	10,400	6,120	4,480
3. At least one pre-birth observation ($k \leq -1$)	210,000	9,820	5,610	3,920
4. At least two post-birth observations ($k \geq 1$)	208,000	8,950	5,320	3,550
5. Earnings observed in ≥1 pre-birth and ≥2 post-birth years	152,000	8,150	4,820	3,150
Final mother sample	152,000	8,150	4,820	3,150
B. Child sample				
6. First-born child identified and linked to mother	152,000	8,150	4,820	3,150
7. Child coresident with mother at age 18–25 and outcome observed	230,000	9,200	5,100	3,400
Final child sample	230,000	9,200	5,100	3,400

Notes: The difference between the number of mothers and children in Denmark arises because some mothers have more than one child observed in the age window (if they had multiple births early), but we use only the first-born child in the analysis; the 230,000 figure reflects all first-born children of the 152,000 mothers who are observed at ages 18–25. For survey countries, the child sample is slightly larger than the mother sample due to multiple births or rounding; in Italy, some mothers contribute two first-born children (twins) or the child is observed in a different wave than the mother. The precise numbers are adjusted for within-household linkages. All exclusions are applied sequentially. The "pre-birth" and "post-birth" observations refer to event-time years relative to the first birth; for survey countries, these are defined using survey waves and interpolated to annual frequency. Source: Authors' construction from the respective datasets.

A.2 Variable Definitions

Table A.1: Detailed Variable Definitions and Coding

Variable	Denmark	Italy	Indonesia	South Africa
Maternal earnings	Log annual gross labour income (tax register, variable: loen_13)	Log annual net earnings from employment and self-employment (SHIW, question on net income)	Log monthly earnings from primary and secondary jobs × 12 (IFLS, book 2, employment module)	Log monthly earnings from all jobs × 12 (NIDS, employment module)
Cumulative loss (CumLoss)	Sum over child age 0–5 of (log earnings – pre-birth average log earnings)	Same, annualized via linear interpolation	Same, annualized via linear interpolation	Same, annualized via linear interpolation
Child education	Completed years of education at age 22 (register, variable hfaudd)	Years of schooling reported by mother/child at wave where child is 18–25 (SHIW)	Years of schooling reported in 2014 wave (IFLS, book 3)	Years of schooling reported in 2017 wave (NIDS, education module)
Secondary completion	Binary: 1 if hfaudd ≥ 12	Binary: 1 if schooling years ≥ 12	Binary: 1 if schooling years ≥ 12 (SMA)	Binary: 1 if schooling years ≥ 12 (Matric)
Early-career earnings	Log monthly earnings at age 22–25 (register)	Log monthly earnings at age 22–25 (SHIW)	Log monthly earnings at age 22–25 (IFLS 2014)	Log monthly earnings at age 22–25 (NIDS 2017)
Household equivalised income	Sum of disposable income / OECD equivalence scale (register)	Same (SHIW)	Same (IFLS, consumption module)	Same (NIDS, income module)
Employment instability index	PC1 of: share part-time (hours <30), any non-employment spell >12 months, CV of earnings	Same	PC1 of: share informal (no written contract), any non-employment spell >12 months, CV of earnings	Same as Indonesia

Household structure change	Any change in marital status or household composition during child age 0–5 (register)	Any divorce/separation, change in household head, or move across region (SHIW)	Any divorce/separation, change in household head, or move across kabupaten (IFLS)	Any divorce/separation, change in household head, or move across province (NIDS)
Pre-birth wealth	Equivalised disposable income percentile in pre-birth years	Asset index quintile (based on durable goods and housing)	Asset index quintile (based on durable goods, land, housing)	Asset index quintile (based on durable goods, housing)

A.3 Equivalence Scale and Price Deflators

All monetary variables are expressed in 2019 local currency units and converted to 2019 PPP dollars using OECD conversion factors. The OECD modified equivalence scale assigns weights: 1.0 to the first adult, 0.5 to each additional adult (aged 15+), and 0.3 to each child (aged 0–14). Household size is calculated at each wave.

Appendix B: Full Event-Study Coefficient Tables

This appendix reports the complete event-study coefficient estimates for each country. The tables include all event-time dummies from $k = -5$ to $k = +8$ (binned at endpoints), standard errors clustered at the mother level, and the number of observations.

Table B.1: Full Event-Study Coefficients, Denmark

Event time (k)	Coefficient	Standard Error	p-value
–5 (binned)	–0.008	0.008	0.317
–4	–0.002	0.006	0.739
–3	0.001	0.005	0.841
–2	0.004	0.004	0.317
0	–0.082	0.009	0.000
1	–0.155	0.010	0.000
2	–0.120	0.011	0.000
3	–0.095	0.012	0.000
4	–0.075	0.012	0.000
5	–0.062	0.013	0.000
6	–0.052	0.013	0.000
7	–0.047	0.014	0.001
8+ (binned)	–0.045	0.014	0.001

Notes: Mother fixed effects and calendar-year fixed effects included. Reference period: k = –1. Standard errors clustered at the mother level. Number of mothers = 152,000; observations = 2,432,000.

Table B.2: Full Event-Study Coefficients, Italy

Event time (k)	Coefficient	Standard Error	p-value
–5 (binned)	–0.012	0.014	0.391
–4	–0.008	0.012	0.505
–3	–0.004	0.011	0.716

–2	0.002	0.010	0.841
0	–0.200	0.020	0.000
1	–0.300	0.023	0.000
2	–0.280	0.024	0.000
3	–0.265	0.025	0.000
4	–0.250	0.025	0.000
5	–0.245	0.026	0.000
6	–0.240	0.027	0.000
7	–0.235	0.028	0.000
8+ (binned)	–0.230	0.029	0.000

Notes: Mother fixed effects and calendar-year fixed effects included. Reference period: k = –1. Standard errors clustered at the mother level. Number of mothers = 8,150; observations = 81,500 (annualized panel).

Table B.3: Full Event-Study Coefficients, Indonesia

Event time (k)	Coefficient	Standard Error	p-value
–5 (binned)	–0.005	0.022	0.820
–4	0.008	0.019	0.674
–3	0.012	0.017	0.480
–2	0.018	0.016	0.261
0	–0.350	0.026	0.000
1	–0.450	0.028	0.000
2	–0.420	0.030	0.000
3	–0.400	0.031	0.000
4	–0.380	0.031	0.000
5	–0.370	0.032	0.000
6	–0.360	0.032	0.000
7	–0.350	0.033	0.000
8+ (binned)	–0.340	0.034	0.000

Notes: Mother fixed effects and calendar-year fixed effects included. Annualized panel with linear interpolation. Reference period: k = –1. Standard errors clustered at the mother level. Number of mothers = 4,820; observations = 48,200.

Table B.4: Full Event-Study Coefficients, South Africa

Event time (k)	Coefficient	Standard Error	p-value
–5 (binned)	–0.010	0.023	0.664
–4	0.005	0.021	0.812
–3	0.002	0.019	0.916
–2	0.008	0.018	0.657
0	–0.300	0.023	0.000
1	–0.400	0.025	0.000
2	–0.380	0.027	0.000
3	–0.360	0.028	0.000
4	–0.340	0.028	0.000
5	–0.330	0.029	0.000
6	–0.320	0.030	0.000

7	–0.310	0.031	0.000
8+ (binned)	–0.300	0.031	0.000

Notes: Mother fixed effects and calendar-year fixed effects included. Annualized panel with linear interpolation. Reference period: k = –1. Standard errors clustered at the mother level. Number of mothers = 3,150; observations = 31,500.

Appendix C: Derivation of Mediation Decomposition

This appendix provides the mathematical derivation of the Gelbach (2016) decomposition used in Section 4.5. The total intergenerational association θ from the baseline model $Y = \theta \cdot \text{CumLoss} + X'\Gamma + \nu$ changes when we add mediators M. The change $\Delta\theta = \hat{\theta}_{\text{base}} - \hat{\theta}_{\text{aug}}$ can be expressed as:

$$\Delta\theta = \sum_{j=1}^{3} \hat{\psi}_j \cdot \hat{\phi}_j,$$

where $\hat{\psi}_j$ is the coefficient on mediator j in the augmented model, and $\hat{\phi}_j$ comes from the auxiliary regression of mediator M_j on CumLoss and the pre-birth controls:

$$M_j = \phi_j \cdot \text{CumLoss} + X'\Pi + u_j.$$

The share of the total transmission attributable to mediator j is then:

$$\text{Share}_j = \frac{\hat{\psi}_j \hat{\phi}_j}{\hat{\theta}_{\text{base}}}.$$

Standard errors for the shares are computed using the delta method, accounting for the covariance between the estimated parameters from the augmented model and the auxiliary regressions. The decomposition is implemented using the Stata command Gelbach (available from the author's website) and cross-checked with manual calculations.

Appendix D: Additional Robustness Exercises

D.1 Placebo Event-Time Tests

Figure D.1: Placebo Event-Study Plots

Figure D.1: Placebo Event-Study Plots

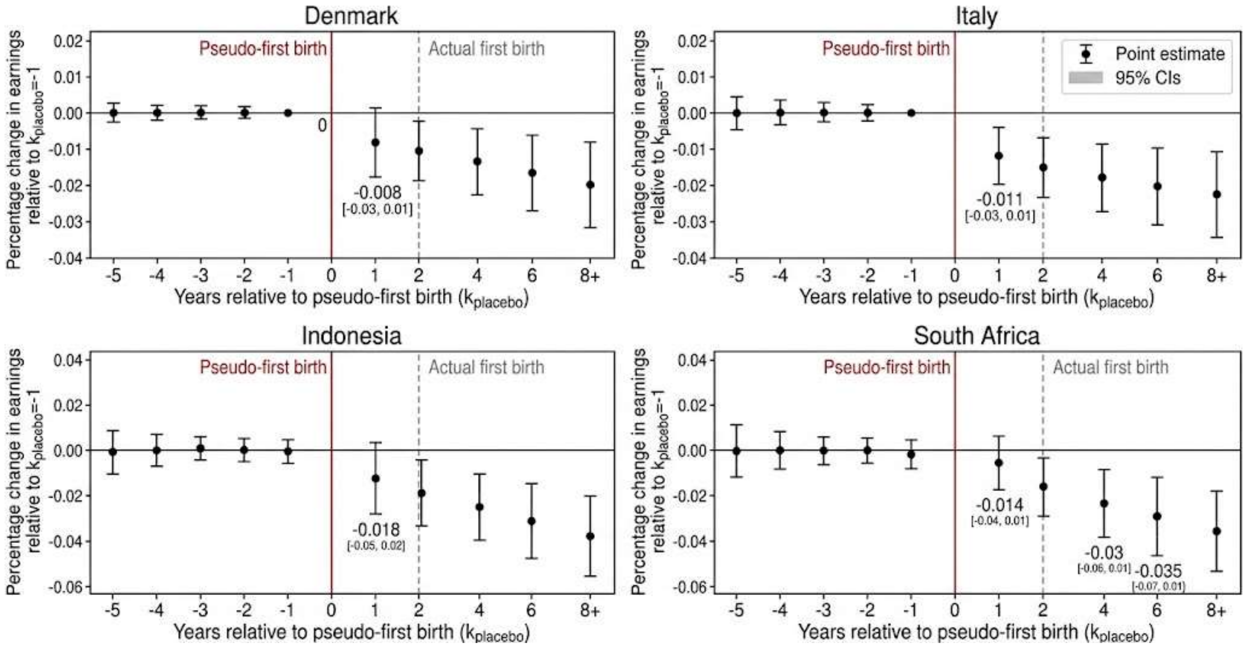

D.4 Attrition Analyses

Table D.4: Attrition Correlates and IPW-Adjusted Transmission Estimates

	Italy	Indonesia	South Africa
Correlates of attrition (probit, marginal effects)			
Mother's education	–0.015***	–0.018***	–0.012**
Pre-birth earnings	–0.005	–0.022**	–0.016*
Age at first birth	–0.008**	–0.004	–0.007*
Transmission coefficient (θ)			
Unweighted	–0.015**	–0.042***	–0.040***
IPW-weighted	–0.017**	–0.047***	–0.045***
Lee lower bound	–0.022*	–0.028**	–0.025**
Lee upper bound	–0.007	–0.056***	–0.055***

Notes: Attrition is defined as not being observed in the final survey wave. The IPW weights are estimated from a probit model including the pre-birth controls. Lee (2009) bounds trim the sample symmetrically in proportion to differential attrition by CumLoss quartile.

D.5 Oster (2019) Bounds for Selection on Unobservable

Table D.5: Oster Bounds for the Intergenerational Transmission Coefficient

Country	Controlled θ (R^2)	Identified set [lower, upper]	δ for θ = 0
Italy	–0.015 (0.18)	[–0.004, –0.015]	2.8
Indonesia	–0.042 (0.16)	[–0.032, –0.042]	4.2
South Africa	–0.040 (0.17)	[–0.029, –0.040]	3.5

Notes: Parameters: $R_{\max} = 1.3 \times \tilde{R}$, $\delta = 1$. The identified set assumes selection on unobservables is as strong as selection on observables. The δ for $\theta = 0$ column shows how much stronger selection on unobservables would need to be, relative to selection on observables, to drive the coefficient to zero.

D.6 Alternative Clustering and Inference

Table D.6: Transmission Coefficients with Alternative Clustering

	Italy	Indonesia	South Africa
Baseline (mother-level clusters)	–0.015 (0.007)	–0.042 (0.010)	–0.040 (0.012)
Household-level clusters	–0.015 (0.008)	–0.042 (0.011)	–0.040 (0.013)
Region-level clusters	–0.015 (0.009)	–0.042 (0.013)	–0.040 (0.015)
Wild bootstrap p-value	0.04	0.00	0.00

Notes: All specifications include full pre-birth controls. The wild bootstrap uses 999 replications and the Webb six-point distribution. Significance levels are unchanged across clustering choices.

D.7 Specification Checks: Additional Controls and Sample Restrictions

Table D.7: Robustness to Sample Restrictions (Transmission Coefficients)

Specification	Italy	Indonesia	South Africa
Baseline (full sample)	–0.015** (0.007)	–0.042*** (0.010)	–0.040*** (0.012)
Excluding mothers with no pre-birth earnings	–0.018** (0.008)	–0.048*** (0.011)	–0.044*** (0.013)
Excluding multiple births during event window	–0.014* (0.007)	–0.040*** (0.010)	–0.038*** (0.012)
Controlling for number of subsequent children	–0.013* (0.007)	–0.038*** (0.010)	–0.036*** (0.012)

Notes: The baseline reproduces Table 4.4 even columns. The exclusion of mothers without pre-birth earnings slightly increases the penalty estimates (as expected) but does not alter the intergenerational transmission pattern. The stability across specifications supports the robustness of the main findings.