Informality and Motherhood Penalties over the Life Cycle

Karolina Andrade
(Yale)

Diogo Britto
(Bocconi)

Lucas Finamor
(EESP-FGV)

Alexandre Fonseca
(Federal Revenue)

Fernando Mattar
(Blue Cross Blue Shield)

Breno Sampaio
(UFPE)

Gabriel Ulyssea
(UCL)

March 18, 2026

Abstract

Developing countries are characterized by substantial gender inequality and large labor market informality. The latter is particularly pronounced among women and increases substantially after childbirth. While informality may help mothers remain attached to the labor market by offering flexibility, informality is also associated with lower earnings, limited human capital accumulation, and no social insurance—potentially exacerbating career penalties over the life cycle. This paper studies motherhood penalties over women's life cycle, how they are shaped by informality, and how public policies could reduce them. We build a novel dataset linking over 40 million births in Brazil to both parents' employment outcomes to document large formal employment penalties of around 40%, half of which materializes before childbirth. We develop and estimate a partial equilibrium life cycle model of women's fertility and labor supply choices across self-employment, formal and informal wage employment. The counterfactual exercises indicate that motherhood penalties are long-lived, persisting for more than 30 years. Policies that increase formal employment flexibility could substantially reduce lifetime employment penalties, eliminate the shift to informal employment after motherhood, and reduce lifetime consumption losses by more than half.

Keywords: Informality, gender inequality, motherhood penalties

1 Introduction

Despite significant progress in the past decades, large gender gaps persist in the labor market (e.g. Blau and Kahn, 2017, Goldin, 2014, Heath et al., 2024, Petrongolo and Ronchi, 2020). Childbearing plays a central role (Cortés and Pan, 2023), with "motherhood penalties" explaining a significant fraction of gender inequality worldwide (Kleven et al., 2025). This is particularly acute in developing countries—especially in Latin America— which have some of the world's largest motherhood penalties and substantial gender gaps in labor supply and earnings (Berniell et al., 2023, Heath et al., 2024, Kleven et al., 2025).

Developing countries are also characterized by high levels of labor market informality, which are substantial among women and increase sharply after childbirth (Berniell et al., 2023). On the one hand, informal jobs offer flexibility that can help women with childcare responsibilities remain attached to the labor market (Berniell et al., 2021). On the other hand, informal employment is associated with lower earnings, limited human capital accumulation, and provides no social benefits or insurance (Ulyssea, 2020). Consequently, even if informality mitigates employment penalties after childbirth, it may exacerbate labor market penalties over the life cycle. Understanding how informality shapes motherhood penalties—and the potential tradeoffs between short- and long-run consequences for women's careers—is therefore central to understanding gender inequality in developing countries and to designing effective policies to reduce it.

This paper fills this gap by making two main contributions. First, we leverage new data on more than 40 million births in Brazil that we link to both parents and their employment outcomes, welfare use, and sociodemographic characteristics. These unique data, typically unavailable in developing countries, allow us to precisely estimate motherhood penalties using event study specifications combined with matching. We estimate large formal employment penalties, of around 40%, half of which materializes *before* childbirth. Second, motivated by these high motherhood penalties we develop and estimate a partial equilibrium, life cycle model to analyze women's fertility and labour supply choices across self-employment, formal and informal wage employment, and their consequences for women's human capital accumulation, life cycle earnings, employment, and consumption. We use the estimated model to quantify the main mechanisms behind these large motherhood penalties, as well as to study potential policies that could reduce them.

Our analysis relies on five main administrative data sources. First, the Brazilian Person Registry (*Cadastro de Pessoa Fisica*), which is maintained by the Brazilian Tax Authority (RFB) and covers the entire population. All individuals are identified by their national individual ID (the CPF), and the registry contains the full name, gender, date of birth,

mother's full name, and the history of addresses of more than 250 million individuals. Second, we use the Treasury's Dependents Claims registry, which despite its richness only covers individuals in the upper part of the income distribution who file tax returns and claim dependents. We thus use a third administrative registry, the *Cadastro Unico* (CadUnico). It includes all individuals who have ever been beneficiary of, or are eligible to, Federal Government's social programs, constituting a near-census of low socioeconomic status (SES) households. It includes individuals' and parents' full name, date of birth, sociodemographic and household characteristics of more than 110 million individuals also identified by CPF. Fourth, we use the School Census, which is filled by all public and private schools in the country, and provides complete parent-child links for all students enrolled in 2014. Fifth, we use the RAIS dataset, a linked employer-employee registry administered by the Federal Government covering the universe of formal workers and firms in Brazil from 2002 to 2019.

We complement these administrative data sources with the National Household Survey (PNAD), a repeated cross-section representative of the Brazilian population. The PNAD contains information on labor market outcomes (including informal employment) and fertility. We also use the "PNAD Continua" (PNAD-C), which is a rotating panel survey that follows households for five quarters and contains complete information on labor market outcomes thus allowing us to compute transition rates between all labor market states. Finally, we use the Decennial Population Census and the Origin-Destination Survey from the state of São Paulo to provide some motivat ing evidence.

We use these data to estimate the motherhood penalty on formal employment and earnings using an event study estimator with matching. Throughout the paper we focus on low-skill women (with up to completed high school), as this is the group for whom the penalties are the largest (e.g. more educated women can afford private childcare services). We compare parents who had their first birth in t with those who had it in $t + 4$, matching on observable sociodemographic characteristics (we perform robustness checks with parents who have their first birth two and three years after). We estimate a substantial motherhood penalty in formal employment outcomes: after birth, mothers' outcomes (presence in the formal sector, months formally employed and formal wages) decline more than 40% relative to the matched control group of women with no births. These gaps narrow down relatively fast compared to richer countries, such as Denmark, but a gap of around 20% remains after 3 years. Throughout the paper, we find no effects for fathers.

Interestingly, around half of this gap materializes *before* birth, during pregnancy. We then separate women into two groups—those employed 12 months prior to birth and those who are not—and show that the "pregnancy effect" is entirely driven by the latter. Moreover, most of the 60% motherhood penalty this group experiences occurs during pregnancy, which

is a novel piece of evidence. In contrast, women who were employed one year prior birth only start experiencing the motherhood penalty after the 4-months minimum mandatory, employer-provided maternity leave expires.

Motivated by these facts, we develop a partial equilibrium life cycle model in which women make decisions on three dimensions: (i) fertility—trying to conceive; (ii) labor market participation; and (iii) conditional on participation, whether to work as self-employed, formal or informal wage employment in full- or part-time regimes. Labour market frictions are modeled in the form of involuntary layoffs from formal and informal jobs, and stochastic arrival rates of employment opportunities, for both employed and unemployed individuals. We allow for sector-specific (i.e. self-employment, formal and informal wage employment) fixed costs of working in terms of time (e.g. commuting times), availability of part-time jobs, returns to experience, and child care costs. The latter is motivated by the fact that, conditional on hours worked, different sectors could imply different childcare costs if they offer different degrees of flexibility in the allocation of the workload. The model takes into account family demographics—marital status, number of children and age of youngest kid—and the history of formal and general work experience accumulated over their life cycle. Finally, it incorporates different shocks affecting women's decisions: conception, marriage/divorce, labor market participation, job offers, separations, and earnings.

We estimate the model using a two-step, simulated method of moments (SMM) estimator. In the first step, we set or estimate some parameters outside the model, such as the discount rate and coefficient of risk aversion. In the second step, we take those as given and estimate a SMM using a large vector of data moments, including the event study estimates, transition rates, employment and fertility patterns by age. We use 120 moments to estimate 34 parameters in the second stage. The model fits the data remarkably well, both with targeted and untargeted moments.

We use the estimated model to perform counterfactual simulations with the goal of better understanding the determinants of the motherhood penalties we observe in the data, the role of informality, and the short- versus long-run consequences of childbirth for women's labor market outcomes. The counterfactual exercises indicate that sector-specific childcare costs are the main source of persistent motherhood penalties. Lowering these costs considerably reduces participation losses over the life cycle and eliminates shifts towards informality conditional on participation. By contrast, other childcare-related frictions, such as pregnancy-related search costs and the utility cost of working with a newborn, primarily affect short-term outcomes, while long-term outcomes are unchanged.

We then explore counterfactuals that increase the flexibility of formal contracts to match that of informal and self-employment arrangements. Greater flexibility in the formal sector—

through changes in childcare costs, fixed work costs, and part-time availability—attenuates the motherhood penalty, particularly along the intensive margin (i.e. employment allocations conditional on participation). While expanding part-time availability reduces the penalty, fully aligning formal-sector flexibility with that of self-employment, especially through sector-specific childcare costs, generates substantially larger welfare gains for women.

We contribute to the literature on motherhood and child penalties (Adda et al., 2017, Kleven et al., 2025, 2019, 2021, Kuziemko et al., 2020), and on gender gaps in labor supply and earnings (Blau and Kahn, 2017, Cortés and Pan, 2023, Goldin, 2014, Petrongolo and Ronchi, 2020). We provide precise reduced-form estimates for a large developing country using linked administrative data, which allow us to uncover an important "pregnancy penalty" that accounts for nearly half of the overall motherhood penalty we estimate. We combine rigorous empirical evidence with a quantitative analysis from a novel life-cycle model featuring labor supply and fertility decisions, extending the framework of Adda et al. (2017) to account for the presence of informal jobs and self-employment—key features of developing countries' labor markets—as well as heterogenous child care costs across sectors.

We also contribute to the informality literature, in particular the body of literature that analyzes the determinants of workers' sector choice (see Ulyssea, 2020, Ulyssea et al., 2025, for reviews). This literature has mostly focused on static frameworks with no life cycle considerations. Important exceptions are Finamor (2025), Joubert (2015), Joubert and Todd (2024), to which we bring the first model, to the best of our knowledge, that features women's decisions regarding fertility and labor supply, emphasizing differences in human capital accumulation and flexibility across sectors. We build on the evidence in Berniell et al. (2021) and Leyva (2025), but we go beyond their short-run analyses, and improve upon the evidence from pseudo-panel designs to estimate the motherhood penalties (Berniell et al., 2023, 2024).

The remaining of the paper is organized as follows. Section 2 describes the data sets used and how we leverage different administrative data sources to link parents to births. Section 3 shows the empirical strategy and reduced-form estimates of the child penalties for parents in the formal sector in Brazil. Sections 4, 5, and 6 present the structural model, its estimation, results and model fit, respectively. In Section 7 we discuss the counterfactual results.

2 Data

This section describes all the data sources used in both the reduced-form and structural analyses. Throughout the paper, we define formal workers as those who hold private sector jobs with a signed labor contract ("carteira de trabalho assinada"). We exclude public

sector and military personnel from this category. Informal employees are those who have an employer but no formal contract, while self-employed do not have an employer-employee relationship.

Linking parents to births We use three administrative data sources to link parents to births. First, the Brazilian Person Registry (*Cadastro de Pessoa Fisica*), which is maintained by the Brazilian Tax Authority (RFB) and covers the entire population. All individuals are identified by their national individual ID (the CPF), and the registry contains the full name, gender, date of birth, mother's full name, and the history of addresses of more than 250 million individuals identified by a CPF. Additionally, the Tax Authority also maintains the administrative record of all dependents claims in tax returns filled between 2006 and 2019. This dataset provides a straightforward CPF-CPF pair of tax-filler parent and dependent child (similar to Chetty et al., 2014), but it only captures individuals in the upper part of the income distribution who file tax returns and claim dependents.

Second, we use the Single Registry of Social Programs (*CadUnico*) which is an administrative registry that centralizes information on all actual and potential beneficiaries of federal social programs, constituting a near-census of low socioeconomic status (SES) households in Brazil spanning from 2011 to 2020. It contains household-level information such as composition and characteristics of the dwelling, and individual-level information on all household members including age, gender, schooling, full name, year of birth, race, addresses, mother's and father's full names, and income profiles of more than 110 million individuals identified by CPF.[1]

The third one is the School Census, which is mandatory and hence filled by all public and private schools in the country. It contains detailed information on students and schools, which can be tracked with a unique student and school identifier over the years. In 2014, the census provides parent-child links for all students enrolled in both private and public schools, including children's full name, birth date, municipality of birth, and both parents' full names. These data allow us to link children's student ID to their parent's ID, enabling us to link children and parents throughout several years in the School Census.

We link children to their mothers and fathers separately and take the intersection of both linked datasets as our main sample. For mothers, the first step uses the Brazilian Person Registry that contains one's mother's name. If the mother can be uniquely identified by

[1] All individuals in *CadUnico* must provide an ID number, which can be the CPF or the CNIS (*Cadastro Nacional de Informações Sociais*), an alternative individual identifier used by the federal government for beneficiaries of social programs. When only the latter is available, we were able to create a cross-walk between both using alternative data sources, so attrition due to lack of individual identifiers is negligible.

her name, we use it to link the child's identifier to her mother's.[2] For fathers, we first use *CadUnico* implementing this same procedure, linking children to uniquely-named fathers (around 50% of the population). For both mothers and fathers, we additionally match children to males and females that have their parents' full name and live in the same address. Because we use the *CadUnico* to identify fathers, this process yields CPF-CPF pairs of children and fathers from the middle-lower part of the income distribution. We thus complement these with information from adult males who claimed their children (aged 0-24) as tax dependents between 2006 and 2019, which directly yields CPF-CPF pairs from the upper part of the income distribution. When the tax filler person is a female, we do the same to identify additional mother-children links not included in the first part of the procedure. Finally, we obtain additional parents-children links directly available from the School Census.

Formal employment outcomes We use the *Relação Anual de Informações Sociais* (RAIS), a linked employer-employee registry administered by the Federal Government covering the universe of formal workers and firms in Brazil. It contains individual information on each job spell, such as the contracts' starting and ending dates, earnings, the reason for termination, and detailed demographic characteristics such as date of birth, race, and education. We construct a panel with the complete formal employment history of individuals identified by their CPF from 2002 to 2019.

Survey Data To complement the administrative data discussed above, we use the National Household Survey (PNAD), a repeated cross-section collected by the National Bureau of Statistics (IBGE) that is representative of the Brazilian population. It contains complete information on individuals' labor market outcomes—including informal and self employment—as well as information on women's fertility, such as the number of children ever had and age of the youngest child. We complement the PNAD with its more recent panel version ("PNAD Continua"—PNAD-C), which is a rotating panel that follows households for five quarters and also contains complete information on labor market outcomes, although not on fertility.

Sample Across data sets, we focus on individuals who live in urban areas, working in the private sector, and with at most completed high school. We restrict parents to be aged 17

[2]The share of individuals with a unique name in the country is large because Brazilians typically carry one or more surnames from both their parents. These individuals are easily identified in the person registry, which includes names and individual identifiers. We adopt some general conventions and standardize all names and addresses across different administrative registries to avoid common misspellings and inconsistencies. Using this procedure, **?** show that individuals with unique names do not strongly differ from the overall population along several characteristics.

or older. Women are between 17 and 45 years old and born between 1970 and 1996, while men are born between 1950 and 1996. The sample is restricted to individuals whose first child was born in 2012 or 2013, along with their control groups, which are defined to be four years apart in treatment timing. These year restrictions allow us to compare treated and control individuals within similar birth cohorts and exclude periods affected by the COVID-19 pandemic. For these cohorts, we observe labor market outcomes for five years around childbirth, covering the period from 2010 to 2017.

These restrictions give us 2.9 million first births matched to both parents, from which we take a random sample to estimate the motherhood penalties in the following section. Table 1 describes our sample and already highlights substantial gender differences. Men are older on average, as they tend to have children later in life. The average number of children is similar across genders, at about 1.2, which is below Brazil's overall fertility rate since individuals are observed roughly four years after their first birth and have not yet completed fertility. In terms of education and labor market outcomes, women are more likely to have completed high school but are less likely to be formally employed in the year prior to childbirth and more likely to be registered in the Cadastro Único, which reflects their greater reliance on social programs.

3 Motherhood Penalties in the Formal Labor Market

We estimate motherhood penalties in the formal labor market using our individual-level data on family-links and employment covering the 2002-2019 period, combined with a difference-in-differences estimator with matching. We define the treatment group as all parents who had their first child in 2012 or 2013, while the control group is formed by parents who had their first child four years later (2016 and 2017).[3]

Each treated parent is matched with a control parent who: (i) has the same sex; (ii) has their first child four years later; (iii) was born in the same year; (iv) lives in the same city; (v) has the same presence in the formal sector 12 months prior to the treated parent's childbirth; (vi) has the same presence in *CadUnico* 12 months prior to the treated parent's childbirth; and (vii) has the same schooling level. When a treated parent is matched with multiple controls, one unit is randomly selected.

We estimate the following regression on the sample of treated and matched control parents, separately for mothers and fathers:

[3]This specification allows us to track the effects over a longer horizon. We conduct robustness checks using parents who gave birth two and three years later as alternative controls, and the results remain unchanged.

$$Y_{it} = \alpha + \gamma Treat_i + \sum_{t=-P}^{T} \gamma_t \left(Treat_i \times Time_t\right) + \sum_{t=-P}^{T} \lambda_t Time_t + \epsilon_{it} \tag{1}$$

where i denotes individuals, and $Treat_i$ is a dummy indicating the treatment group. $Time_t$ denotes dummies for months since child birth, which we can define very precisely because we have the exact dates of birth in our data, and $t = 0$ indicates the month of birth. The coefficients $\{\gamma_{-24}, ..., \gamma_{40}\}$ thus identify dynamic treatment effects, $t = -12$ is the baseline omitted period, and $\{\gamma_{-P}, ..., \gamma_{-11}\}$ estimate anticipation effects. Finally, λ_t time fixed-effects absorb aggregate, time-varying shocks.

Results Figure 1 shows the effect of child birth on formal labor market participation, as estimated from Equation 1. Treatment effects are rescaled by the average outcome for the control group, and each dot shows than the employment effect relative to the matched control group in the same period. For fathers, coefficients are precisely estimated around zero in the baseline period. Effects remain negligible during the first 18 months after childbirth, with a modest decline of about 5% in the second and third years. For women, baseline coefficients are slightly negative, consistent with some anticipatory exit from formal employment prior to childbirth, although the magnitude remains below 5%. Once treated women become pregnant, we observe a sizeable decline in employment, reaching -20% at childbirth and falling further to about -45% by the end of the first year. Employment gradually recovers after the first year but remains at -30%, corresponding to a gap of employment of roughly 12 percentage points relative to the control group.

Figures A.1 and A.2 present the results for months worked and formal-sector labor earnings, showing very similar patterns. Additionally, Figure A.3 presents the same analysis for alternative control groups composed of individuals who became parents two and three years after the treated group. The results are remarkably similar across specifications, reinforcing the robustness of our findings.

To examine heterogeneity in treatment effects by initial labor market attachment, we split the sample of women according to women's employment status at $t = -12$, and the results are presented in Figure 2. Women who are not formally employed may be out of the labor force, unemployed, working informally, or self-employed. In the short run, the effects vary significantly for the two groups. For formally-employed women, employment remains stable until four months after childbirth, which corresponds to the exact duration of maternity leave. Once employment protection ends, employment falls by roughly 35%. In contrast, for women who were not formally-employed at baseline, negative effects begin during pregnancy.

Most of the decline occurs before childbirth: employment is already about -50% one period before childbirth and declines to around -60% within six months. Interestingly, for both groups, employment recovery begins after 12 months at similar rates. Nevertheless, by the end of the period, employment among initially non-employed women remains about 40% below the control group, compared with roughly 25% for those initially employed.

4 Life Cycle Model

4.1 Setup

The model is a lifecycle model of fertility and occupation decisions for women. Time is discrete and each period lasts six months. We model decisions from the age 18 up to 59, at which point the decision is static.[4] We adopt a partial equilibrium framework in which labor demand is assumed to be exogenous and agents behave independently.

In each period, the individual decides whether to try conceiving or not. Conditional on deciding to conceive, there is uncertainty on whether she will become pregnant or not. If she does, in the following period she will give birth—pregnancy lasts for one period or 6 months—and thus babies are born in the second period after decisions are made. Conditional on their conception choice, the woman decides whether to participate in the labor market. If she decided not to participate, she stays out of the labor force and no further choices are made this period. If she enters, she may receive offers and decide whether to work as a formal employee, informal employee or self-employed. If no job is accepted, she may be unemployed. Some jobs offer the possibility of part-time work, and therefore women working on these jobs may choose to work either full or part-time. Formal and informal work accumulate different types of experience, which have differential returns in each sector.

4.2 State space and choices

At the beginning of each period, a vector (Ω) of state variables defines the relevant information for all the choices to be made in that period. The state space includes: age (t), marital status (m), number of children (n), age of the youngest child (a), last-employment status (j^-), formal experience (e) and general experience (g).

[4]This is motivated by the fact that, in Brazil, for most of the sample period, a female worker who has contributed to the social security system for a minimum of 15 years is eligible for retirement at age 60, whereas below this age the worker is only eligible if she has a minimum of 30 years of contributions.

$$\Omega = (t, m, n, a, j^{-}, e, g) \tag{2}$$

Women have three choices to make at every period: (i) whether to try to conceive (b), (ii) employment status (j), which encompasses both the labor market participation choice and occupation choice in case of participation, and (iii) whether to work full or part-time (p).

4.3 Labor market

The state variable j summarizes the labor market position, combining participation status and occupational category: (O) out of the labor force, (U) unemployed, (F^0) formal employed without part-time work, (F^1) formal employed with part-time work, (I^0) informal employed without part-time work, (I^1) informal employed with part-time work, (S^1) self-employed with part-time work. Self-employed individuals always have the option to work part-time, that is, there is no S^0. Women in jobs that offer part-time options (F^1, I^1, S^1) may choose to work part-time ($p = 1$) or full-time ($p = 0$), while those in jobs without this option (F^0, I^0) can only work full-time ($p = 0$).

Individuals receive job offers stochastically, and, when employed, face exogenous separation shocks. These job arrival rates and separation rates capture the labor market frictions in our model. If the woman enters the labor market, she may receive offers from the formal, informal, and self-employment sector. The probability of receiving an offer to j' depends on her current status j. We will denote this probability $\lambda_{j,j'}$. For paid employment, there is a probability that a new job offer includes the option of part-time work. Whenever she receives a job offer from a given sector with a given part-time option, she keeps the job until it is either exogenously destroyed, which occurs with probability δ_j, or she quits. The participation decision is made before any offer arrives. The woman therefore compares the value of being out of the labor force with the expected value of entering, taking into account the distribution of potential offers across sectors.[5] Base arrival rates for formal and

[5]In terms of labor market dynamics, there is no difference for the (exogenous) arrival or destruction between full and part-time jobs. That is, $\lambda_{j^0,j'} = \lambda_{j^1,j'}$ and $\delta_{j^0} = \delta_{j^1}$, for all j. This still allows for endogenous quit and switches across sectors to be different for full and part-time jobs, given agents' optimal choices.

self-employment are adjusted for an experience channel, such as

$$\begin{aligned}\lambda_{j,F} &= \lambda^0_{j,F}\left(\frac{1}{1+\omega_F \exp(-e)}\right)\\ \lambda_{j,S} &= \lambda^0_{j,S}\left(\frac{1}{1+\omega_S \exp(-(e+g))}\right)\end{aligned} \tag{3}$$

That is, arrival rates are adjusted by a term that is increasing in experience. This captures the idea that less experienced workers face a lower probability of receiving job offers or becoming self-employed. The exponential specification reflects the fact that experience has larger effects on job-finding rates at low levels of experience, with more muted effects thereafter. It also captures that only formal experience (e) is rewarded in the formal sector.

When working in the formal sector, women accumulate formal experience. When working as informal employees or self-employed, women accumulate general experience. When working part-time in either sector, women accumulate a fraction $(1-\rho_e)$ of the full-time equivalent experience accumulation. That is,

$$\begin{aligned}e_{t+1} &= e_t + \mathbb{1}[j_t = F]\,(1-\rho_e p_t)\\ g_{t+1} &= g_t + \mathbb{1}[j_t \in (I,S)]\,(1-\rho_e p_t)\,.\end{aligned} \tag{4}$$

Labor earnings are given by a quadratic in experience and an additive sector productivity shock. The formal sector only rewards formal experience, while the other sectors reward total experience, formal and the general experiences. When working part-time, women lose a fraction ρ_w of their labor earnings. At each point in time, there is a sector-specific productivity shock given by ζ^j_t.

$$w_t(j,e,g,p;\zeta_t) = \begin{cases} exp\Big(r^F_0 + r^F_1(e) + r^F_2(e)^2 + \zeta^F_t\Big)(1-\rho_w p), & \text{if } j = F\\ exp\Big(r^I_0 + r^I_1(e+g) + r^I_2(e+g)^2 + \zeta^I_t\Big)(1-\rho_w p), & \text{if } j = I\\ exp\Big(r^S_0 + r^S_1(e+g) + r^S_2(e+g)^2 + \zeta^S_t\Big)(1-\rho_w p), & \text{if } j = S\\ 0, & \text{if } j = O, U. \end{cases} \tag{5}$$

4.4 Net income and consumption

Net income is given by individuals paying taxes, receiving net transfers from their partners, and receiving (net) government transfers. If the woman is married, she receives net transfer

Y^h from her husband. If she works in the formal sector, her final wage is αw, where $\alpha < 1$ could be given by tax and social security contributions and $\alpha > 1$ by government transfer programs. If she is unemployed and was previously formally employed, she is entitled to unemployment benefits Y^u. If the does not have any additional source of income she receives a government transfer Y^0. Therefore, her income is given by:

$$Y_t = \max(w_t(1 + \alpha \mathbb{1}[j = F]) + Y^m m + Y^u \mathbb{1}[j = U \text{ and } j^- = F], Y^0). \tag{6}$$

We assume that consumption is determined by income and abstract from savings and borrowing decisions due to data limitations. This assumption is supported by evidence from expenditure surveys showing that low-skilled women's households have very low levels of savings.

4.5 Leisure

Women are endowed with one unit of leisure at every period. If she does not participate in the labor force, she enjoys the full leisure ($\ell = 1$). If she enters the labor force, she loses some utility depending on her labor market choices. If she stays unemployed she loses h^U of leisure time, due to job search. If she is working, she pays a sector-specific fixed participation cost (h^j) and the hours, depending on full time (h^{FT}), or part time (h^{PT}). If she has a young child and she works, she loses hours dedicated to childcare costs. Which depends on the age of the youngest child and the sector:

$$g(a, j) = \begin{cases} 0, & \text{if } j \in [O, U] \\ \left(\frac{\bar{A}-a}{\bar{A}+\frac{1}{\pi}a}\right)\theta(1 - \theta_I \mathbb{1}\{j = [I]\} - \theta_S \mathbb{1}\{j = [I]\}), & \text{if } j \in [F, I, S] \end{cases} \tag{7}$$

Therefore the expression for leisure will be given by:

$$\ell = 1 - h^U \mathbb{1}\{j = U\} - \Big(h^j + h^{FT}(1 - p) + h^{PT} p + g(a, j)\Big)\mathbb{1}\{j \in [F, I, S]\}. \tag{8}$$

4.6 Pregnancy and Children

Children enter the model in several different ways, which we layout here. Children bring joy to parents, here we take a simplified version and assume all the benefits are enjoyed at the time of conception, when women derive the fertility shock ζ_b. To capture the fact that most women tend not to have their children too close together, we implement a spacing utility cost. This is a parameter that women must pay, in case they want to have their children with less than 1.5 years of age difference. This cost is given by γ, and it is fully paid if

the woman is trying to conceive when the baby is just born (1st semester), and it declines linearly to zero when the age of the youngest child is 3 semesters (1.5 years). After that, there is no spacing cost.

Children bring sector-specific childcare costs ($g(a, j)$), which captures only how it is differentially costly to work and adjust hours depending on the age of the youngest child and sector. That is why we term sector-specific childcare cost, and it is normalized to zero when woman is not working. Having positive costs for women not working would be equivalent to simultaneously change the child intercept at the utility function.

Participation cost for women with babies. We have a utility cost (ϕ) which is paid in case women with babies (0-12 months) decide to participate in the labor market. This is to capture how difficult it is for women to go back to work with babies. This cost is paid regardless of the sector of employment and it is not differential across sectors.

Pregnant women may suffer a lower probability of receiving job offers, we allow pregnancy to affect the job offer arrival rates, which are depressed by a factor κ. This can capture labor market discrimination, anticipation of future discrimination by women, or lower search efforts. All these channels are observationally equivalent, and we do not attempt to disentangle them. We call this term pregnancy search costs. It is the same cost for all sectors.

Maternity leave and employment protection. We follow the legislation and to capture the maternity leave and employment protection for formal workers. That means that women that were formally employed when the baby is born are entitle do maternity leave (4 months) and job protection (6 months).

4.7 Shocks

There are three categories of shocks in this economy: preference, labor market, and demographic shocks. In the firs class, preference shocks, at every period women observe their preferences for being out of the labor force (ζ_o) and for conceiving (ζ_b). The distribution of these shocks may depend on the state space. The second group of shocks are those related to the labor market, they include productivity shocks that enter the wage determination (ζ_w^j), and the exogenous destruction, arrival of new offers, and arrival of part-time jobs (summarized on the vector ζ_λ). Lastly, the demographic shocks govern the marital status transitions and pregnancy success (summarized by ζ_d). The vector ζ collects all these shocks. We denote μ_k and σ_k, respectively the mean and standard deviation of shock k. To capture the fact that married women are more likely to have children and that most families have two children we allow the fertility shock to have a different mean by marital status and by whether the woman has already two children or not. That is $\mu_{\zeta_b}(m, n) = \mu_{\zeta_b} + \mu_{\zeta_{b,m}} m + \mu_{\zeta_{b,n}} \mathbb{1}\{n \geq 2\}$.

4.8 Flow utilities

The per-period flow of utility has four main components: instantaneous utility from consumption and leisure, captured by a CRRA function, utility from kids, utility from conception and being out of the labor force, and disutility of transitioning from being out of the labor force to working.

$$\begin{aligned} u(\Omega; b, j, p; \zeta) = & \frac{(c^{\nu}\ell^{(1-\nu)})^{1-\sigma}}{1-\sigma} + \zeta_b b + \zeta_o \mathbb{1}\{j = O\} - \chi \mathbb{1}\{j \neq O, j^{-} = O\} \\ & - \phi \mathbb{1}\{j \neq O, a <= 1\} - \gamma \left(1 - \frac{a}{1.5}\right) \mathbb{1}\{b = 1, a <= 1.5\} \end{aligned} \tag{9}$$

4.9 Within-period timing

Figure 3 presents an schematic representation of the within-period timing of the model. At the beginning of the period, the woman know all her demographic status (age, marital status, number of children and age of youngest child), and remaining state variables (employment status and experience). She first observes her conception preference shock and decides whether to try to conceive ($b = 0$ or $b = 1$).[6] After this decision, she observes the participation shock, and decides whether to enter the labor market or not. If she does not enter, no further choices are made.

If she enters, her previous employment status plays a crucial role. If she was unemployed, she may receive new offers from F, I, and S. If the does and accepts, she moves to this sector. If the new offers have the possibility of part-time work, she decides whether to work full or part-time. If she receives no offer or does not accept any, she stays unemployed. If she is employed at the beginning of the period, the problem is similar, with two exceptions. She may have her job exogenously destroyed at the beginning of the period, which moves her instantaneously to unemployment, and she will face the labor market as an unemployed person. Otherwise she may receive job offers again, but in this case, if no offer is accepted she may keep her old job. If the current job does not offer the possibility of part-time work, she also has the option of receiving a new job offer from the same sector, with the part-time possibility. After observing all the shocks and possibilities (including de productivity shocks), she then makes the decision where to work (j) and whether to work part-time or full-time if the job has this option. The only exception is for self-employment, to capture uncertainty on self-employed activities, the productivity shock for self-employment is only revealed after making the choice of becoming self-employed.

[6] If the woman is pregnant, or with a child that is just born, or after 45 years of age, she does no hat fertility decisions and $b = 0$.

Then she enjoys consumption and leisure for this period, observes the demographics shocks and move to the next period. The demographic shocks affects next period state variables but not the current period flow utility.

4.10 State space dynamics

We can now specify how the state space evolves. Age (t) evolves deterministically in semesters. If a woman has a successful pregnancy, a is set to -1 to indicate pregnancy, and the following period to $a = 0$ to be birth of the child. After that, a evolves stochastically until the child is 18 years of age, which is the maximum age recorded. A woman is pregnant if, in the previous period, she decided to conceive ($b = 1$) and her pregnancy is successful. This happens with an exogenous probability that depends on the mothers' age. Specifically, the probability is given by:

$$\Pr(a_{t+1} = -1 \mid b = 1) = \frac{\bar{T}_{b,1} - t}{\bar{T}_{b,2} - t}. \tag{10}$$

If the pregnancy is successful, when the birth occurs, n is also increased by one.

The probabilities of entering and exiting a marital/cohabiting relationship are given, respectively, by:

$$\Pr\left(m_{t+1} \neq 0 \mid m_t = 0\right) = \left[1 + \exp\left(\gamma_0 + \gamma_1 t + \gamma_2 t^2\right)\right]^{-1} \tag{11}$$

$$\Pr\left(m_{t+1} = 0 \mid m_t \neq 0\right) = \gamma_3 \tag{12}$$

with $m_t = 0$ representing singlehood. Experience vectors update depending on sector of employment and hours of work as presented by equation 4.

4.11 Model Solution

The model is solved numerically using backward recursion from the final decision period back to the first. We start the model with women at age 18, taken as given the joint distribution of marriage, number of children, age of the youngest child, and sector of employment at age 18 in PNAD.

5 Estimation

This section describes the estimation procedure for the empirical model specified in the previous section. We implement a two-step estimation strategy, where in the first step we

set some parameters outside the model and in the second step we estimate the remaining parameters using the method of simulated moments.

5.1 Parameters estimated outside the model

In the first step, we set some parameters to institutional values or to external calibrations from the literature. We also estimate some parameters externally. These include mostly parameters relating to parts of the model that are assumed to be exogenous.

From the preference parameters, we set the discount rate β to 0.94, corresponding to an annual discount rate of 0.97 and we set the CRRA coefficient σ to 3. These are values typically used in the literature, and in the middle-range of most values used.

We set the intercept of the wage functions to the observed average wages for individuals at age 18 in the three sectors. We do the same for the variance of the wage shock, setting it to the residual variance of the wage, after controlling for age, separately for each sector. For the part-time wage penalty, we set it to 57%, which is equivalent to the median part-time hours (25 hours), compared to the regular contract of 44 hours. We also set the proportion of experience accumulated in part-time jobs to be 0.5, meaning that working one semester in a part-time job gives half the experience of working one semester in a full-time job.

For the fixed cost of working, we set it to be the observed commuting time from individuals working as formal employees, informal employees, and self-employment in the population Census data. These numbers are 1.75 weekly hours for the formal sector, 1.46 hours for the informal sector and 0.79 hours for self-employment. We set the number of hours of work as the observed hours for full-time (44h weekly hours) and the median hours for part-time jobs (25h weekly hours). For unemployment we use 10h per week, to capture job search effort.

As maternity leave in Brazil lasts for four months and the period of the model is six months, we assume that women that are in maternity leave loss only one-third of the full leisure cost of working. The same applies to the utility cost of working with babies. Women in maternity leave only pay one-third of this cost. Regarding the conception shock, we normalize its variance to one.

For the institutional parameters, we set the value of the minimum transfer of the government to be the value of the conditional cash transfer program (Bolsa Familia), defined as the average real monthly benefit over the period 2004–2019. For the unemployment insurance, in reality the benefit is paid for a limited number of months and its value is determined as a non-linear function of the income of the worker prior to being laid-off, subject to a floor and a ceiling. To approximate the UI function, we compute the deflated averages of the benchmark levels of income specified by the labor legislation from 2002 to 2015 that are used to calculate

the value of the monthly installments. With these, we set the floor and ceiling and define the value paid based on these cut-offs. Over the period from 2002 to 2015, the Brazilian social security system would pay 4 or 5 monthly installments of UI to unemployed workers that had worked at least 12 or 24 months in a formal job, respectively, within the previous 36 months, provided that these workers had been fired without cause. As a simplification, in the model we assume that UI is always paid in 4.5 installments following a transition out of a formal job, regardless of how many periods the individual had been employed in the formal sector. Moreover, because the idiosyncratic shocks to income are assumed to be realized only after decisions are made and are not part of the state space, we also assume that UI is computed based on expected rather than realized income. Finally, we also set α to 1.3, based on the values computed by Haanwinckel and Soares (2021).

We exogenously define the marriage and divorce probabilities choosing parameters to fit some features of the data. To identify the probability of entering marriage/cohabitation as a function of age and the separation probability we compute the frequencies of marriage/cohabitation by age and educational level using data from the PNAD. Furthermore, from 2011 on, the PNAD provides information on whether a woman that is not married/cohabiting has ever lived with a partner or not. This information is crucial because using just the cross-sectional distribution of marriage/cohabitation to identify parameters of the transition process is not enough, as there would be a high and a low turnover process that would likely be equivalent. For this reason, we also include as targets the frequencies of single women that have ever cohabited or not with a partner, by age and educational level, which is enough to separate the desired low turnover process.[7]

Figure **??** presents the fit between the calibrated process and the empirical life-cycle profile of marriage/cohabitation. The overall fit of the profile is satisfactory. To further assess the validity of the calibration, we compare some statistics implied by this approximation with information on marriage and divorce obtained from the "Statistics of the Civil Registry" report by the IBGE. Note that this report concerns only legal marriages and divorces irrespective of cohabitation, while the calibrated process concerns any form of cohabitation spousal relationship. The calibrated process generates average ages at the start of the first relationship of 27 and 24 for women with and without high school education, respectively, and 25.5 overall. The 2002 and 2015 civil registry reports show that the average ages at first marriage in Brazil those years was 26.7 and 27, respectively. Moreover, the average duration of a marriage/cohabitation relationship implied by the model is around 13.4 years for both

[7]Marriage/cohabitation status is defined based on three variables in the data: the person's status within their family, the type of a family, and whether a person is living with a partner, (only present in the data from 2011 on).

educational levels, while the 2011 and 2015 civil registry reports both present an average duration of legal marriages of 15 years. Considering this, the model appears to provide a reasonable approximation to the dynamics of marriage, divorce and cohabitation.

5.2 Second step: Method of Simulated Moments

In the second step, conditional on the exogenous features described above, we estimate the remaining parameters of the model using the method of simulated moments (MSM), as proposed by McFadden (1989). At this stage we estimate a vector $\boldsymbol{\xi}$ of 34 parameters, that a mix of preference and technology parameters.

For a given $\boldsymbol{\xi}$, we can solve the model and obtain all the value and policy functions for individuals. We then simulate a number of complete life-cycle histories of endogenous decisions, states and outcomes. This gives us a simulated panel that we use to compute a vector of moments that can be directly compared to analogous empirical moments. Estimation then amounts to searching for the parameter vector that minimizes a distance function between the simulated and empirical moments. The MSM estimator is defined as:

$$\hat{\boldsymbol{\xi}} = \text{argmin}_{\boldsymbol{\xi}}[\mathbf{m}(\boldsymbol{\xi}) - \mathbf{m}_0]' W^{-1} [\mathbf{m}(\boldsymbol{\xi}) - \mathbf{m}_0] \tag{13}$$

where $\mathbf{m}(\boldsymbol{\xi})$ is the vector of simulated moments, $\mathbf{m}_0$ is the vector of estimated empirical moments and W is a positive semi-definite weighting matrix. We set W to be the average between the identity matrix and the diagonal matrix of the reciprocals of the variances of each sample moment.[8] This weighting matrix gives more weight to moments that are estimated with higher precision, while also giving some weight to moments that are estimated with lower precision, which might be informative for the estimation.

Since our model does not have a closed form analytical solution, a formal proof of identification is unfeasible. Here we discuss how some parameters are more informed by some moments than others there are chosen. Later, in order to validate this heuristic arguments we show the sensitivity matrix proposed by Andrews et al. (2017).

In the utility function, we estimate ν, which governs the relative importance of consumption and leisure in the utility function. This parameter will be more closely related to the moments related to the fraction of part-time work. Particularly from self-employment part-time, as it is the sector that does not offer any friction into part-time opportunities.

In terms of the shocks, we estimate the mean and the variance of the participation shock, and the mean of the conception shock (which varies by marital status). To inform these

[8] That is $W = 0.5I + 0.5W2$, where $W_2 = \lambda \text{diag}(\frac{1}{\hat{\sigma}^2_{m_1}}, \ldots, \frac{1}{\hat{\sigma}^2_{m_k}})$, where $\hat{\sigma}^2_{m_i}$ is the estimated variance of the i-th sample moment and λ is chosen to yield $tr(I) = tr(W_2)$.

parameters we add moments related to the overall participation rate, transitions in and out of participation, and the fertility rate by marital status. We also have moments of number of children by age of the mother. Another parameter that is influenced by the fertility moments is the utility gains from children (κ_2). Finally, another parameter that relates to the proportion of women out of the labor force by marital status is the net transfers of the husband to mothers (Y_m).

To inform the parameter governing the utility cost of spacing child birth are informed by the observed pattern of distance between the first and second child, for all women that have more than two children using the administrative records.

In terms of leisure, we estimate the parameters relating to the sector-specific childcare costs ($\pi, \theta, \theta_S, \theta_I$). There are three sets of moments that help identify these parameters. First the proportion of women in each employment status by age and, importantly, by age of the youngest child. Also the proportion of women working part-time in each sector, also by age of the youngest child. Finally, the reduced-form results showing the employment effects following childbirth for women with different previous employment status is also included. These same moments are also informative to estimate the utility cost of working with a baby (ϕ).

We estimate arrival and exogenous destruction for all employment status and they are highly informed by the observed transitions between all j and j' pairs. We also include as a moment the fraction of formal jobs that are destroyed. This does not coincide with the $j = F$ to $j' = U/O$ transition rate, as some unemployed individuals receive new offers and stays in the formal sector (or choose other sector). As we have this precise measure in the matched employer-employee data, we can included as moment. In this collection of parameters, we also have the probability of a given offer being part-time for the formal and informal sector, that are mostly informed by the part-time moments, importantly the transitions between proportion part-time by sector between mothers without kids and mothers with young babies. The patterns in and out of participation are also informative of the utility cost of joining the labor market (χ).

Returns to experience are estimate to match the life-cycle profile of wages by sector. For simplicity we only estimate one parameter for each sector, and assume that they quadratic with a peak at 20 years of experience. Adjustments on arrival rates by experience level are highly informed by the life-cycle pattern of employment by sector. At the same time, the pregnancy-search cost is mostly driven by the reduced-form estimates of employment effects following childbirth for women that were not formally employed.

6 Results

6.1 Parameter estimates

Table A.2 presents the parameter estimates of the model. We next discuss the key structural parameters and their relationship to children-related costs in the model. First, the estimates indicate that pregnancy lowers women's job-finding prospects: arrival rates for pregnant women are 39.5% lower than for non-pregnant women.[9]

Sector-specific childcare costs vary substantially across sectors. A woman with an infant who works in the formal sector faces a cost that amounts to 12.8 weekly hours, while women in informal jobs or self-employed have costs that are 44.0% and 75.7% lower. These costs decline with the age of the youngest child. When costs in formal jobs are about 12.8 weekly hours for babies, they are equivalent to about 7.9 weekly hours for a child aged 3 and 5.0 hours for a child aged 6. Additionaly, we estimate a large utility costs of working while caring for a baby equal to 0.383. This is equivalent to a women losing about 28.1% of consumption and leisure when working, regardless of the sector, while still taking care of an infant.

Another interesting result relates to the availability of part-time jobs. All self-employment opportunities offer part-time work, while approximately 15.4% of new informal jobs already come with a part-time option. This stands in sharp contrast to formal jobs, where only 1.4% of new opportunities are part-time. This parameter also governs the probability that a woman who stays in the same sector is offered a part-time opportunity. In particular, a women who remains in the informal sector for four quarters has a 48.8% probability of receiving a part-time opportunity, compared with only 5.4% in the formal sector.

6.2 Model fit

The estimated model is able to fit well the most important features of the data, such as the overall allocation and transition rates of women between labor market states and the life-cycle profile of fertility. Figure 4 shows the model fit for the proportion of women with no children, one, two, or three or more children by age. Figures A.4 and A.5 show, additionally, the model fit for the proportion of women with newborns by marital status and the distribution of years between the first and second child are born. Overall, the model provides an excellent fit to these targeted moments.

For the general patterns across sectors, the model is able to fit well the overall allocation of women across the three sectors over the life-cycle (Figure A.6), as well as the income

[9] Arrival rates for pregnant women are scaled by the parameter κ, estimated at 0.605.

profiles (Figure A.7). The transitions across all labor market status also display a very reasonable fit (Figure A.8).

In terms of employment patterns by fertility and age of the youngest child, we show how the model fits the data quite well. Figure 5 shows the model fit for the reduced-form estimates of the motherhood penalty in formal employment, aggregated in semesters to match the timing of the model. We can see how the model fits very closely the pregnancy period and the smaller drop when the child is born and formally employed women may enjoy maternity leave. After that, the model captures well the slow recovery of employment, slightly underestimating bot the drop at 6months-1year and the recovery speed. Figure A.9 shows the model fits well the reduced-form estimates, separately by baseline status, as well as the proportion of women by baseline status: formally employed or not (Figure A.10).

The model also fits well the proportion of women in each labor market status, by age of the youngest child. Figure 6 shows the results. The first panel shows the jump from around 30% of women that are out of the labor force and do not have kids, to more than 60% when the child is born. The model captures well both the jump and the slowly recovery, slightly overestimating the effect at the first year of the child. The following panels show the patterns for formal, informal, and self-employment.

Figure 7 shows the model fit for the proportion of women working part-time across the three employment options and overall. We can see how the formal sector has a very low proportion of part-time workers, and the pattern is flat with the age of the child. This is in sharp contrast with self-employment with a much higher proportion of par-time workers, which peaks when the children are born and slowly declines after that. Figure A.11 shows the model fit of proportion of women out of the labor force, and working part-time by marital status.

6.3 Model fit for untargeted moments

Up to now we have shown the model fit for the targeted moments, that is, the moments that are directly used in the estimation of the model. We can also look at how the model fits some untargeted moments. To do that we compute the proportion of women in each labor market status and also the proportion working part-time by the number of children in the data and in the model. Figure A.12 shows the results. The model presents the overall patterns of the data, with a few exceptions. For out of the labor force and for the formal sector, the data presents a stepper gradient over the number of children than the estimated model. It is worthy noting that the model does not distinguish between more than two children. That is, the state space of the model only considers none, one or two children, as the utility function

and constraints are the same for two or more children.

While the model targeted the average number of children by mother's age, it did so fitting at most three children. Figure A.13 shows the model fit for the distribution of number of children for women with complete fertility histories. While the model slightly underestimates the proportion of childless women, the overall pattern is well captured. We note how the model does not have infertile women, which could explain the underestimation of childless women.

7 Counterfactuals

In this section, we use the model to conduct counterfactual exercises that quantify the key drivers of the motherhood penalty and its life-cycle costs. To compute these costs, we compare women's choices in the baseline model to those in a counterfactual scenario in which women do not have children but continue to make decisions under the same preferences, as if unaware of their infertility. [10]

Figure 8 presents our estimates of the motherhood penalty in labor market participation, formal employment conditional on participation, and non-formal employment conditional on participation, which aggregates informal employment and self-employment. The black curve shows the motherhood penalties over more than 40 yers following the first birth. The effects are large and highly persistent, as it takes more than 30 years to fully dissipate. To provide further intuition, the yellow line plots the motherhood penalty for women who had only one child. The gap between the two curves can therefore be attributed to subsequent births, which reset the motherhood penalty.

In the model, there are three direct channels through which pregnancy and childcare costs affect women's choices: (i) pregnancy search costs that reduce the arrival rate of job offers for pregnant women, (ii) sector-specific childcare costs, and (iii) utility costs of working while caring for a newborn. We also implement three additional exercises focusing on the employment structure: (iv) increasing the availability of formal part-time opportunities, (v) making the formal sector as flexible as the informal sector, and (vi) making the formal sector as flexible as self-employment arrangements. In counterfactual (iv), we increase the share of formal jobs offering a part-time option (λ_f^{pt}) to 100%. In counterfactuals (v) and (vi), we allow the formal sector to mimic the other sectors in terms of the sector-specific childcare costs, fixed-costs of work, and share of part-time opportunities.

Figure 9 shows how the motherhood penalty changes across the six counterfactuals. We find that, although search costs during pregnancy contribute to the employment gap before

[10] Table 3 presents an alternative counterfactual in which infertility is fully anticipated by women.

childbirth, setting them to zero has little effect on the life-cycle profile or overall magnitude of the motherhood penalty. When the utility cost of working with a newborn is eliminated, labor force participation rises substantially in the first years following childbirth, while long-run participation and the intensive margin of informality remain unchanged. In contrast, sector-specific childcare costs have sizable long-run effects on participation and fully eliminate the intensive margin of informality.

Part (b) of Figure 9 plots the results for counterfactuals (iv)-(vi). Increasing part-time work availability in the formal sector to 100% reduces the motherhood penalty in all dimensions, with a modest effect on participation and a moderate effect on the intensive margin of informality. However, equalizing flexibility in the formal sector to that observed in the informal and self-employment has much larger effects on the intensive margin of informality.

To further analyze these results, Table 2 reports the baseline motherhood penalty and its value under each counterfactual for a broader set of outcomes: labor market participation, total formal participation, intensive formality, end-of-life formal experience, end-of-life total experience, average lifetime consumption, and average number of children by women with completed fertility.

All counterfactuals have sizable effects on fertility. Therefore, the implied reductions in the life-cycle motherhood penalty should be interpreted as conservative, since additional births contribute to larger penalties over the life cycle and every counterfactual points to an increase in the average number of children.

The estimates of the motherhood penalty above are computed by comparing the baseline economy to an alternative scenario in which women do not have children. In that scenario, as women do not anticipate that they will remain childless, the estimates capture only the behavioral changes that follow pregnancy and childbirth. If we account for the anticipatory effect, womens' behavior would further change. Table 3 compares these two scenarios and the results indicate that the motherhood penalty would be even larger once anticipatory responses are taken into account: women who expect to remain childless are more likely to participate in the labor force and to sort into formal employment.

References

Adda, J., Dustmann, C. and Stevens, K. (2017). The career costs of children, *Journal of Political Economy* **125**(2): 293–337.

Andrews, I., Gentzkow, M. and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments, *The Quarterly Journal of Economics* **132**(4): 1553–1592.

Berniell, I., Berniell, L., De la Mata, D., Edo, M. and Marchionni, M. (2021). Gender gaps in labor informality: The motherhood effect, *Journal of Development Economics* **150**.

Berniell, I., Berniell, L., De La Mata, D., Edo, M. and Marchionni, M. (2023). Motherhood and flexible jobs: Evidence from latin american countries, *World Development* **167**: 106225.

Berniell, I., Fernández, R. and Krutikova, S. (2024). Gender inequality in latin america and the caribbean, *Technical report*, National Bureau of Economic Research.

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations, *Journal of economic literature* **55**(3): 789–865.

Chetty, R., Hendren, N., Kline, P. and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states, *The Quarterly Journal of Economics* **129**(4): 1553–1623.

Cortés, P. and Pan, J. (2023). Children and the remaining gender gaps in the labor market, *Journal of Economic Literature* **61**(4): 1359–1409.

Finamor, L. (2025). Labor market informality, risk, and insurance.

Goldin, C. (2014). A grand gender convergence: Its last chapter, *American economic review* **104**(4): 1091–1119.

Haanwinckel, D. and Soares, R. R. (2021). Workforce composition, productivity, and labour regulations in a compensating differentials theory of informality, *The Review of Economic Studies* **88**(6): 2970–3010.

Heath, R., Bernhardt, A., Borker, G., Fitzpatrick, A., Keats, A., McKelway, M., Menzel, A., Molina, T. and Sharma, G. (2024). Female labour force participation, *VoxDevLit* **11**(1): 1–43.

Joubert, C. (2015). Pension design with a large informal labor market: Evidence from chile, *International Economic Review* **56**(2): 673–694.

Joubert, C. and Todd, P. E. (2024). Gender pension gaps in a private retirement accounts system: A dynamic model of household labor supply and savings, *Journal of Econometrics* **243**(1-2): 105337.

Kleven, H., Landais, C. and Leite-Mariante, G. (2025). The child penalty atlas, *Review of Economic Studies* **92**(5): 3174–3207.

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

Kleven, H., Landais, C. and Søgaard, J. E. (2021). Does biology drive child penalties? evidence from biological and adoptive families, *American Economic Review: Insights* **3**(2): 183–198.

Kuziemko, I., Pan, J., Shen, J. and Washington, E. (2020). The mommy effect: Do women anticipate the employment effects of motherhood?, *Technical report*, National Bureau of Economic Research.

Leyva, A. M. (2025). Job flexibility and informality.

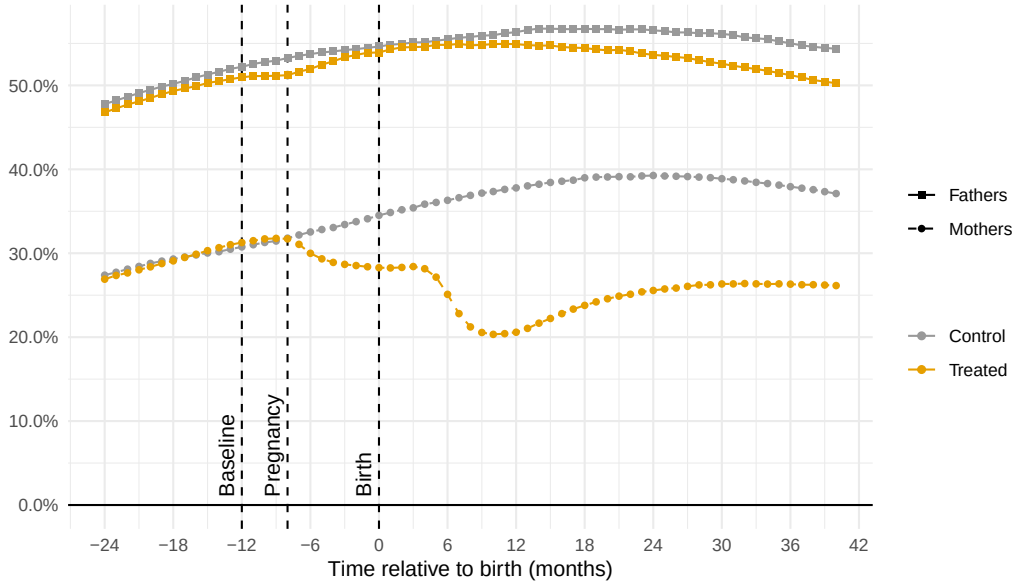
McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration, *Econometrica* **57**(5): 995–1026.

Petrongolo, B. and Ronchi, M. (2020). Gender gaps and the structure of local labor markets, *Labour Economics* **64**: 101819.

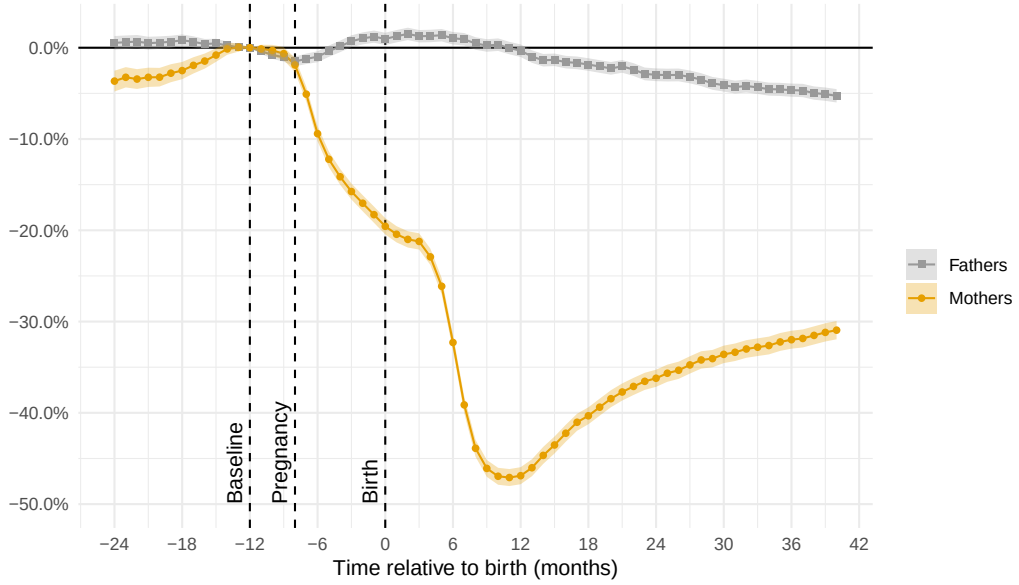
Ulyssea, G. (2020). Informality: Causes and consequences for development, *Annual Review of Economics* **12**: 525–546.

Ulyssea, G., Bobba, M., Gadenne, L. and Harari, M. (2025). Informality, *VoxDevLit* .

Figure 1: Motherhood Penalty in Formal Employment



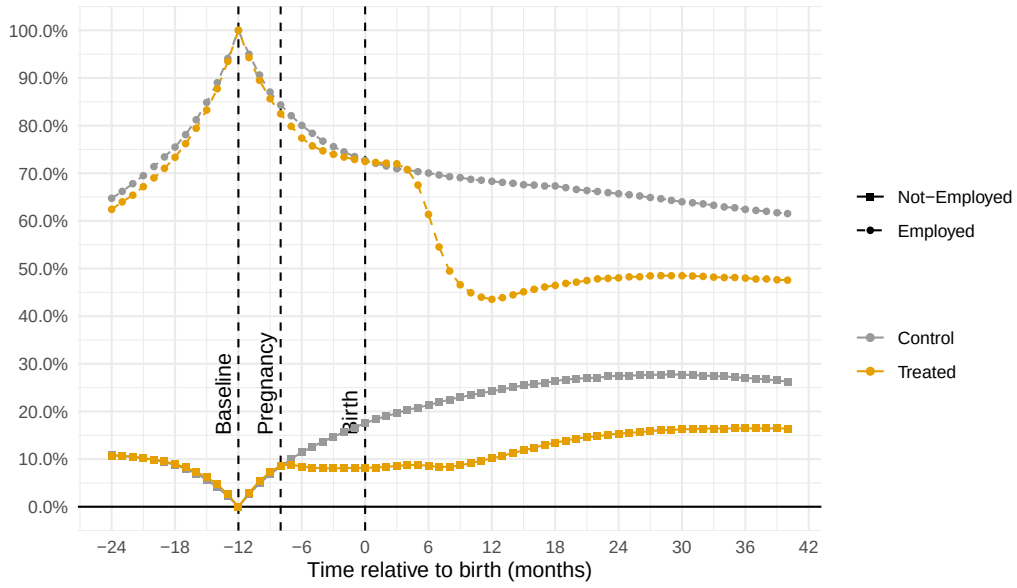
(a) Raw Data



(b) Treatment Effects

Notes:

Figure 2: Motherhood Penalty in Formal Employment, by baseline employment (Raw data)



Notes:

Figure 3: Within-period timing

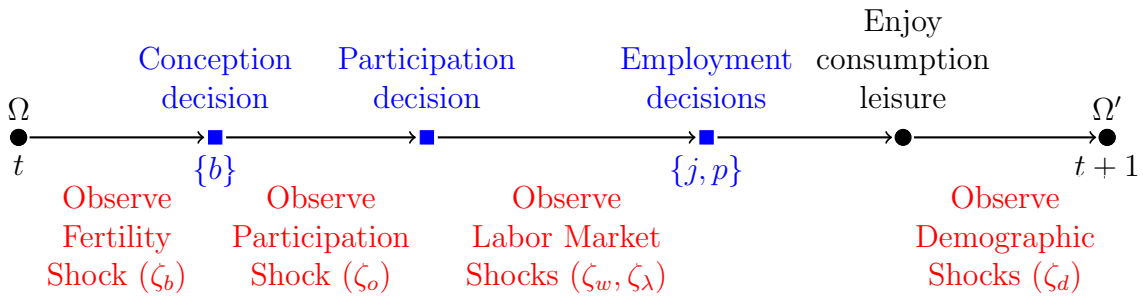
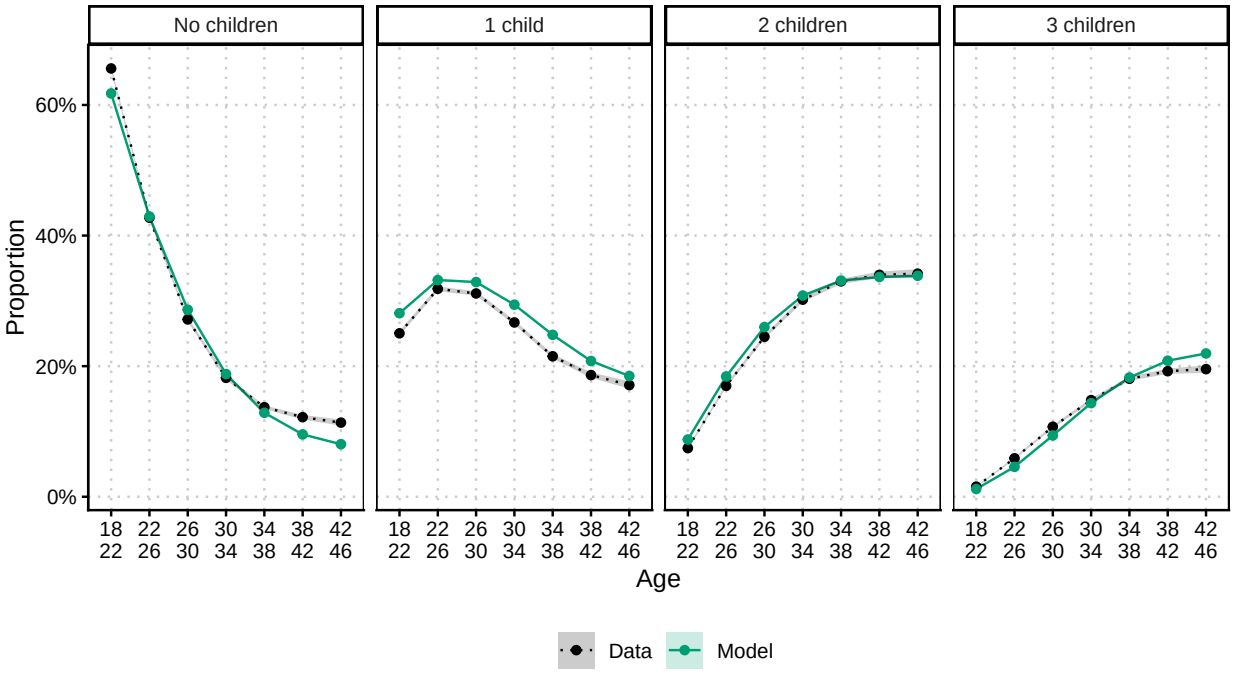
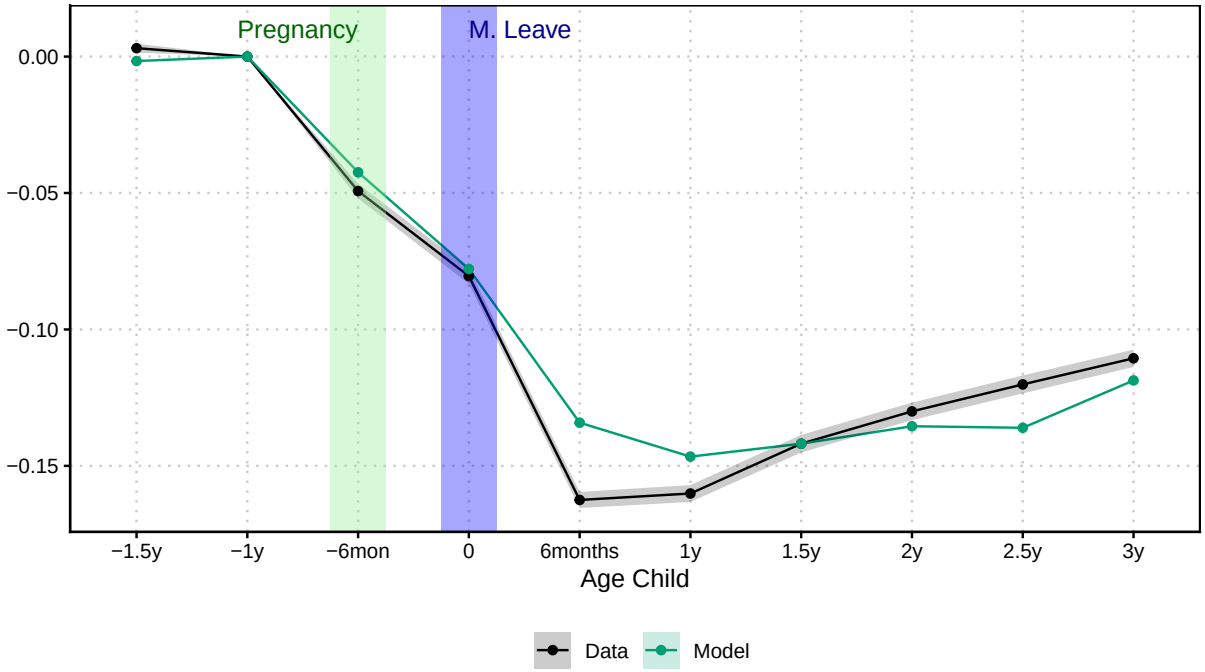


Figure 4: Model fit: Proportion of women by number of children and age



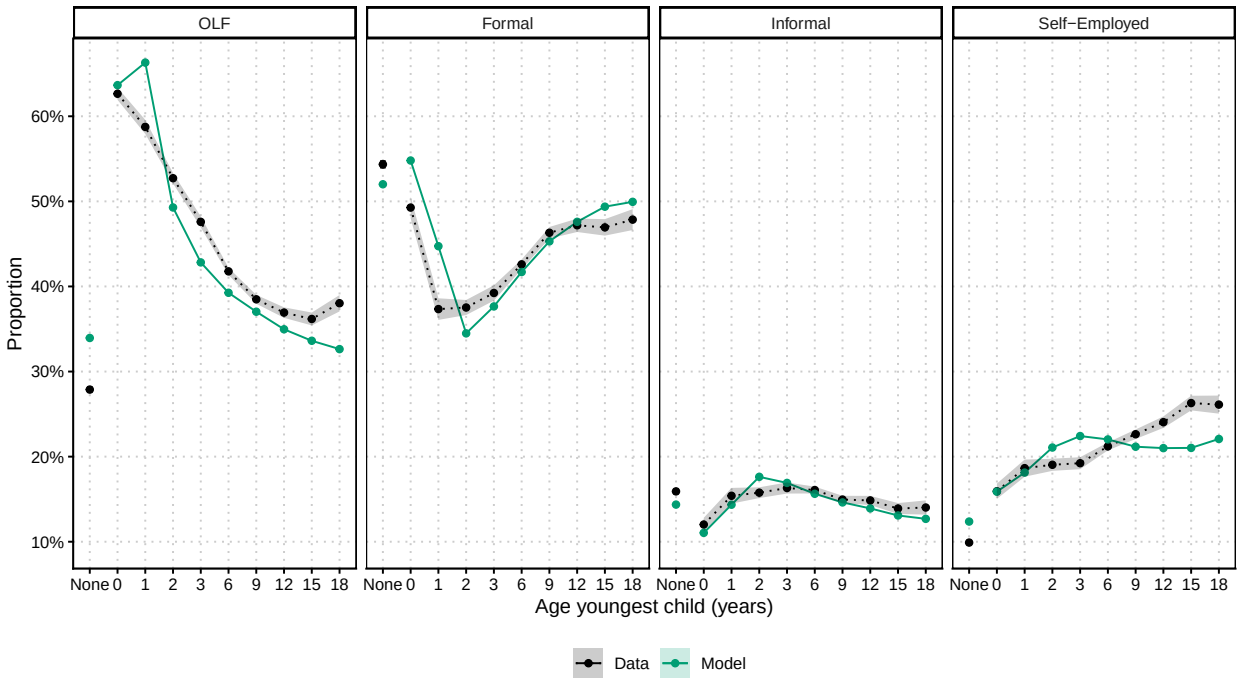
Notes:

Figure 5: Model fit: Reduced-form results of the motherhood penalty



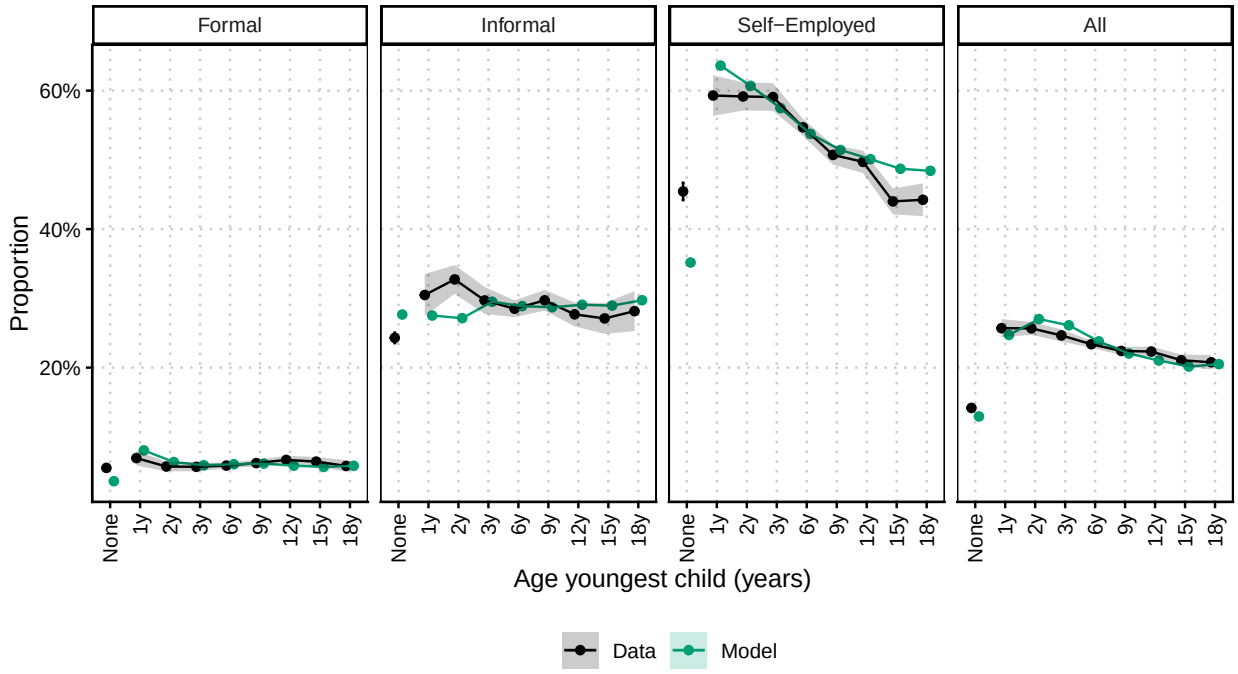
Notes:

Figure 6: Model fit: Reduced-form results of the motherhood penalty



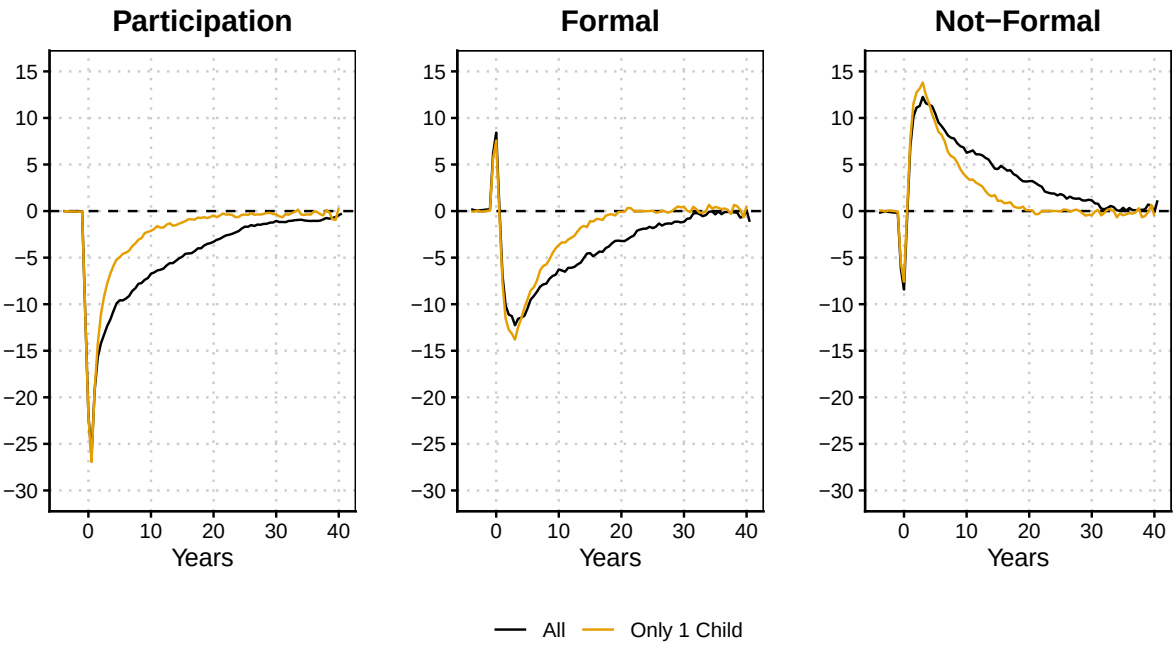
Notes:

Figure 7: Model fit: Reduced-form results of the motherhood penalty



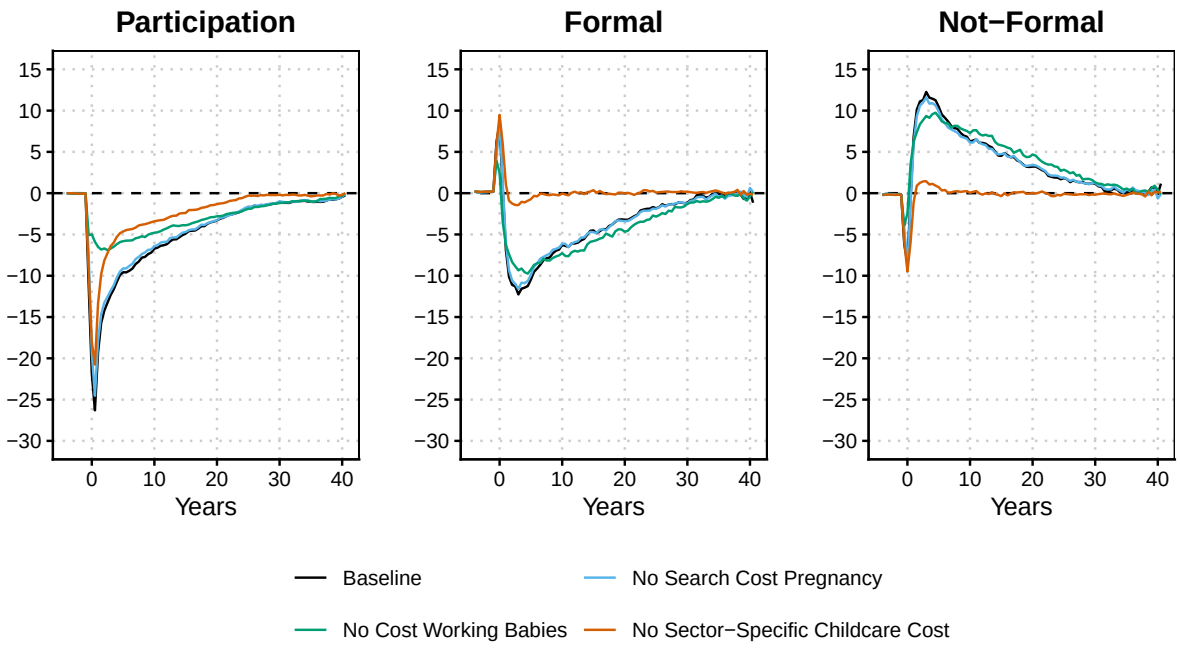
Notes:

Figure 8: Motherhood penalties

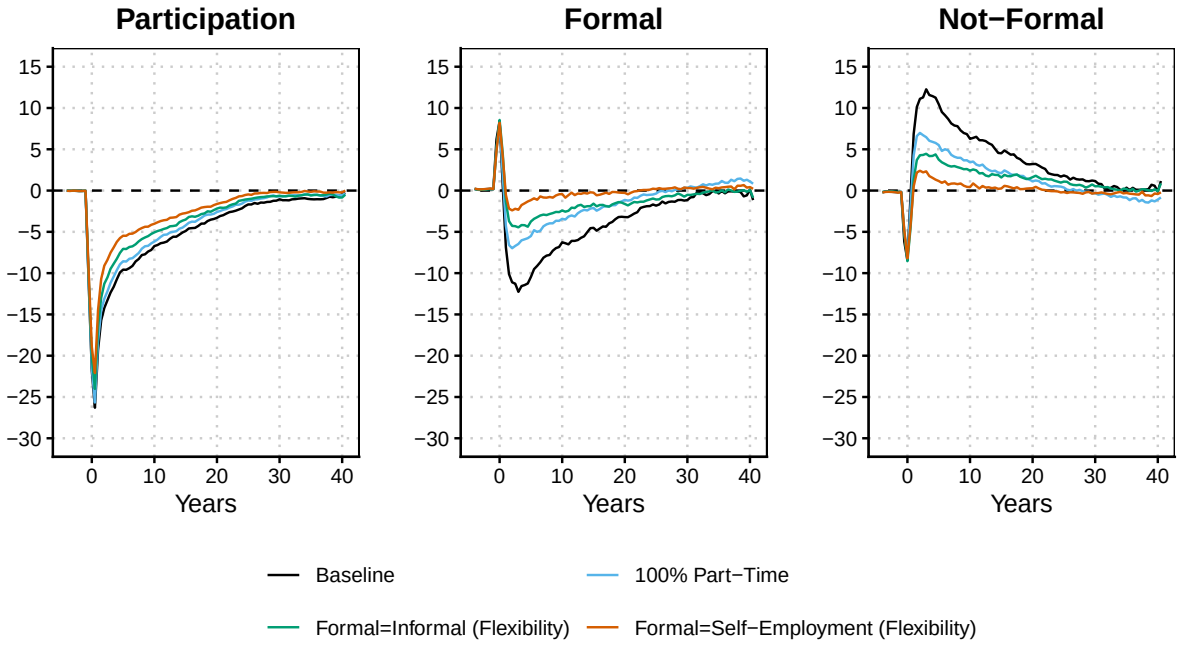


Notes: The figure reports event-study estimates of the motherhood penalty in labor market participation, formal and non-formal employment conditional on participation.

Figure 9: Counterfactual results - Short-run



(a) Counterfactuals (i)–(iii)



(b) Counterfactuals (iv)–(vi)

Notes:

Table 1: Descriptive Statistics

	Women	Men
Year of birth	1988	1986
	(5.30)	(6.02)
Age at first birth	25.7	28.3
	(5.65)	(6.32)
Number of children	1.21	1.23
	(0.48)	(0.50)
Years in sample after first birth	4.02	4.02
	(2.07)	(2.07)
High school graduates	0.61	0.56
	(0.49)	(0.50)
Observed in RAIS	0.47	0.70
	(0.50)	(0.46)
Observed in *CadUnico*	0.22	0.14
	(0.41)	(0.35)
Observations	273,560	286,308

Summary statistics for the final analysis sample after data linkage and sample restrictions. Means are reported with standard deviations in parentheses. The last two rows refer to presence in administrative records in the year prior to conception: formal employment in RAIS and registration in the Cadastro Único.

Table 2: Counterfactual exercises

	LF Participation (pp)	Total Formal (pp)	Intensive Formal (pp)	Formal Experience (years)	Total Experience (years)	Consumption (%)	Number of Children (avg)
Baseline	-4.01	-3.76	-2.61	-1.24	-1.38	-8.19	2.33
No Cost Working Babies	-2.62	-3.53	-3.37	-1.16	-1.24	-6.66	2.96
No Sector-Specific Childcare Cost	-2.32	-1.32	0.09	-0.43	-0.53	-3.74	2.51
No Search Cost Pregnancy	-3.80	-3.56	-2.58	-1.18	-1.31	-7.68	2.41
100% Formal Part-Time	-3.60	-3.08	-1.24	0.00	-1.27	-7.25	2.36
Formal=Informal (Flexibility)	-3.18	-2.62	-1.03	-0.46	-1.03	-5.52	2.40
Formal=Self-Empl. (Flexibility)	-2.57	-1.79	-0.17	0.00	-0.47	-3.53	2.47

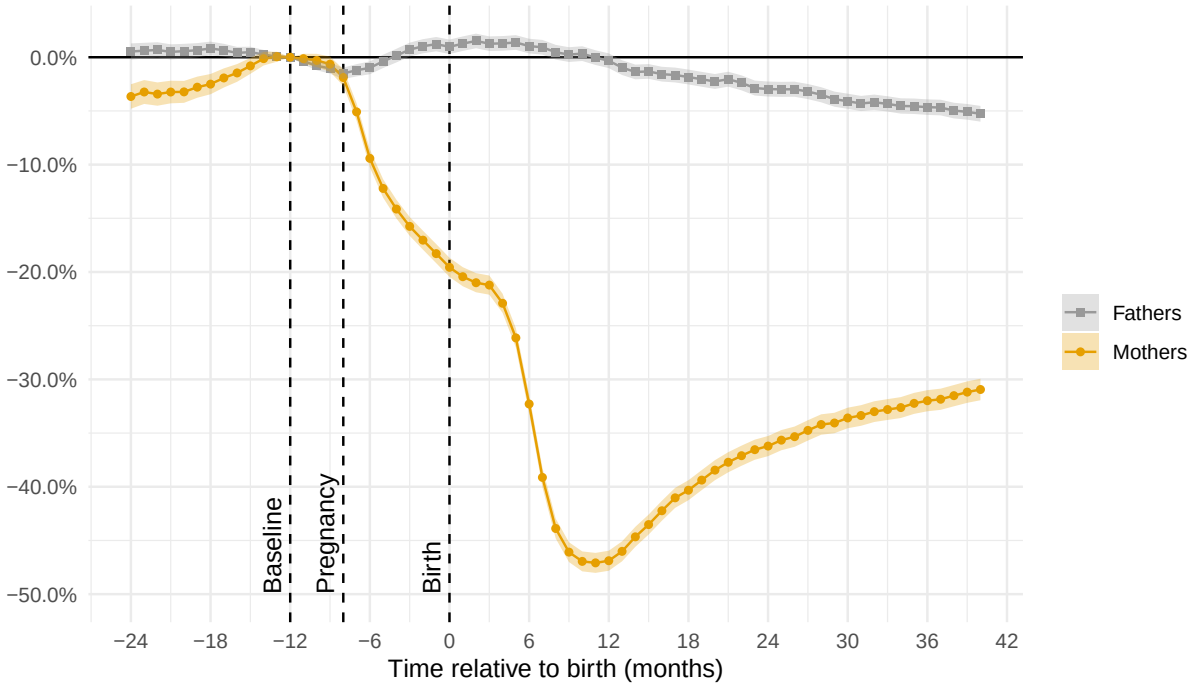
Notes: This table reports model-based counterfactual exercises for the motherhood penalty. The baseline row compares outcomes in the baseline economy to those in a counterfactual economy in which women do not have children. The remaining rows shut down or relax specific mechanisms: the cost of working with newborns, sector-specific childcare costs, and pregnancy-related search costs; part-time availability in formal jobs; and formal jobs with the same flexibility as informal work or self-employment, in terms of costs and part-time work.

Table 3: Counterfactual exercises: fertility

Counterfactual	Participation rate	Formal Employment (Total)	Formal Employment (intensive)
Baseline	0.610	0.287	0.574
No fertility (not-anticipated)	0.644	0.314	0.588
No fertility (anticipated)	0.650	0.325	0.600

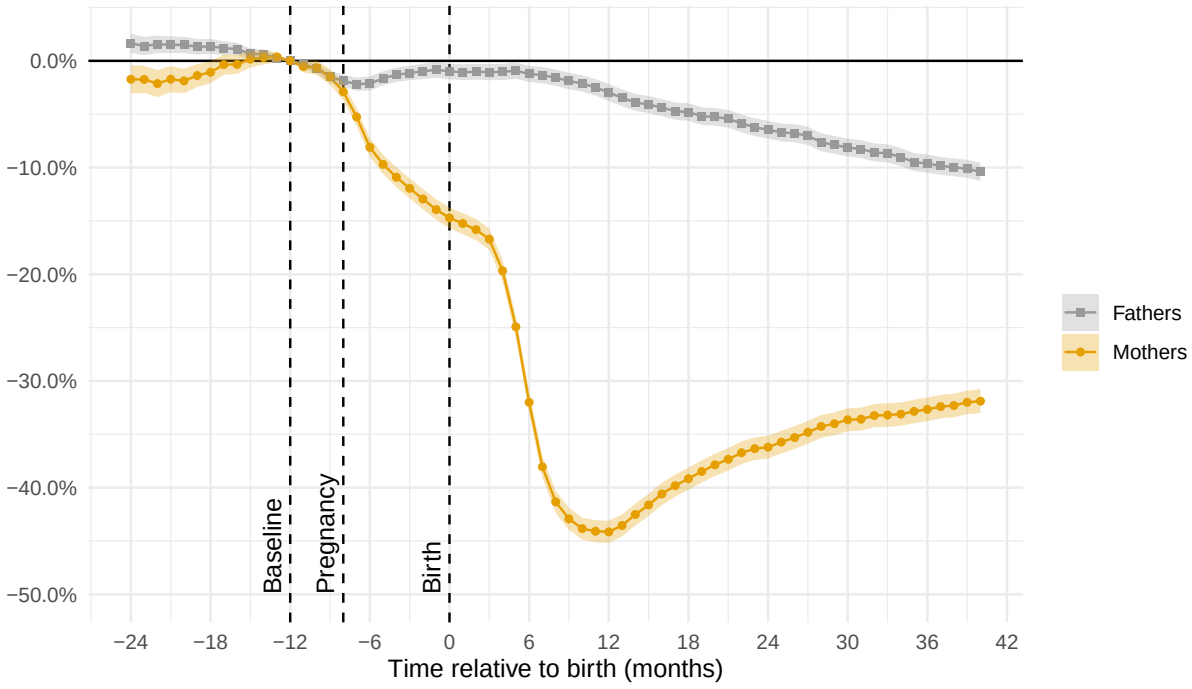
A Additional Figures and Tables

Figure A.1: Effect of child birth on labor earnings in the formal sector



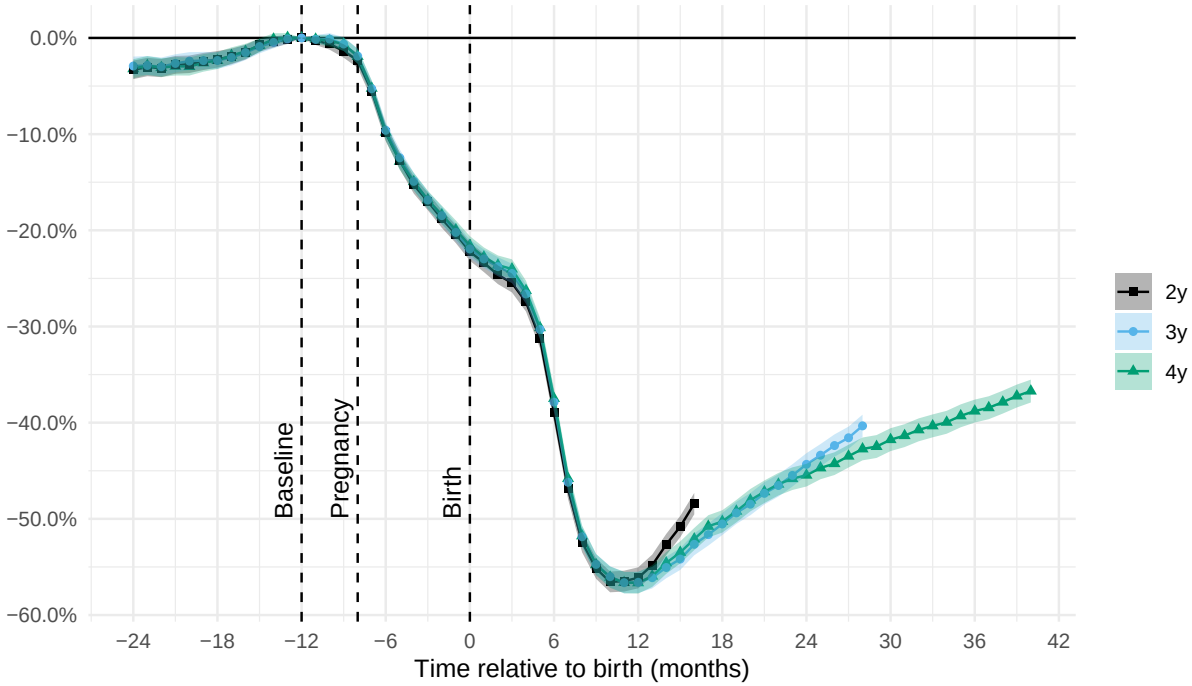
Notes:

Figure A.2: Effect of child birth on months worked in the formal sector



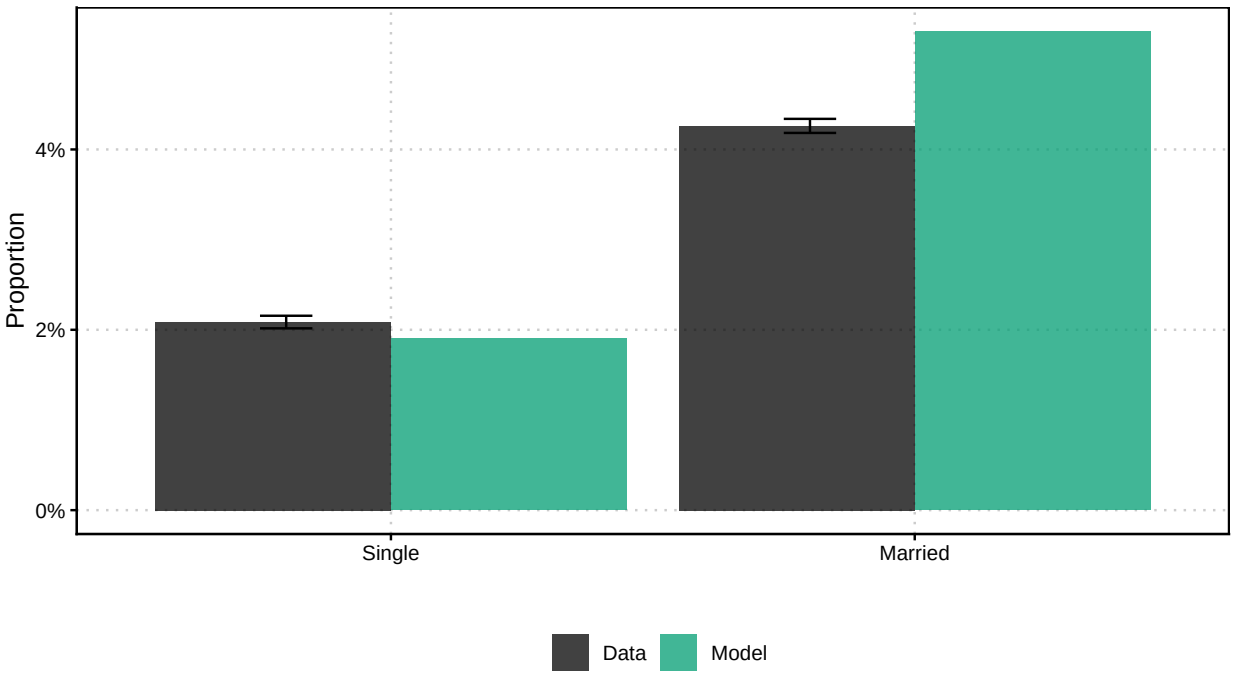
Notes:

Figure A.3: Effect of child birth on formal employment - different control groups



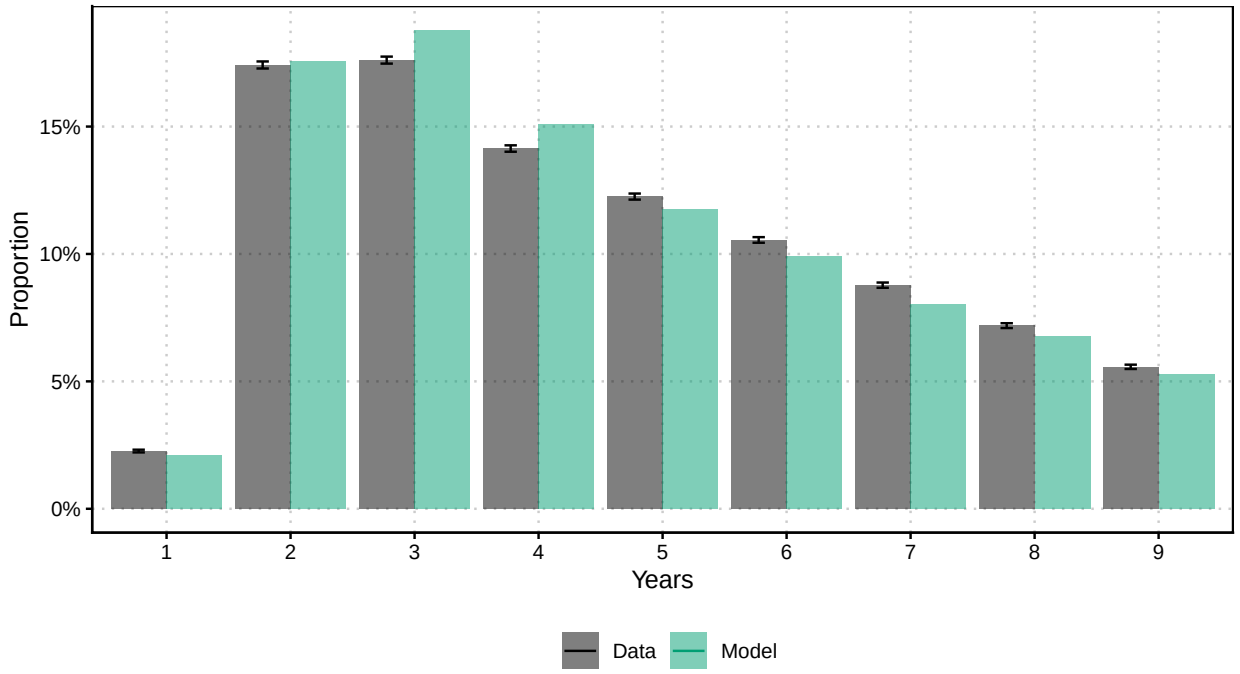
Notes:

Figure A.4: Model fit: Proportion of women with a newborn by marital status



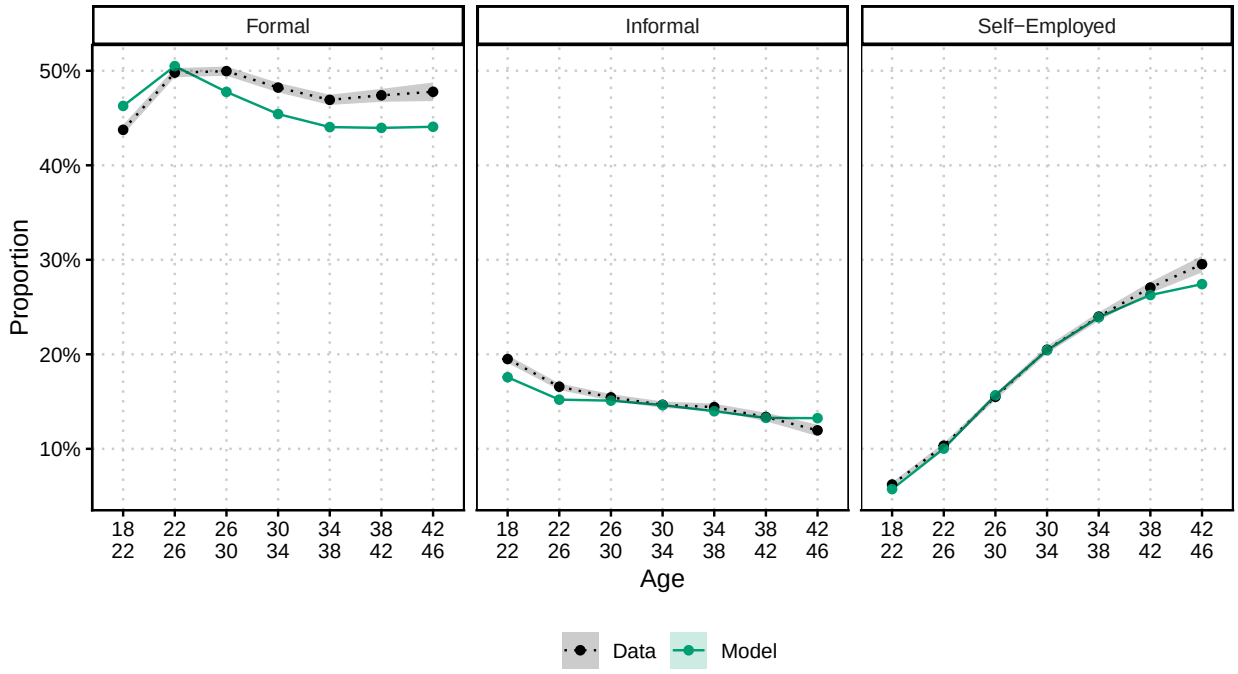
Notes:

Figure A.5: Model fit: Distribution of spacing between first and second child in years



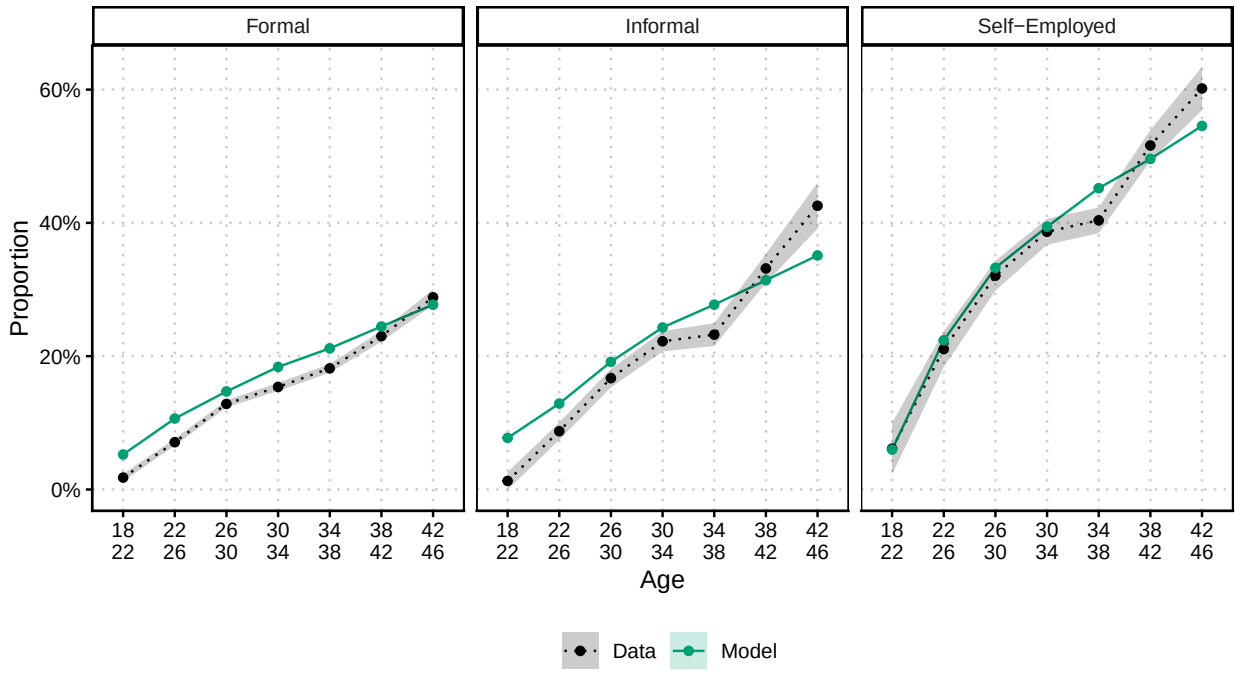
Notes: The figure plots the distribution of spacing between first and second child in years. This is computed as the different in years between the birth of the first and second child for all women that have at least two children in at most ten years. The gray bars plot the distribution in the data, while the green bars plots the distribution in the model. For the data we also present the 95% confidence interval.

Figure A.6: Model fit: Proportion of women across sectors



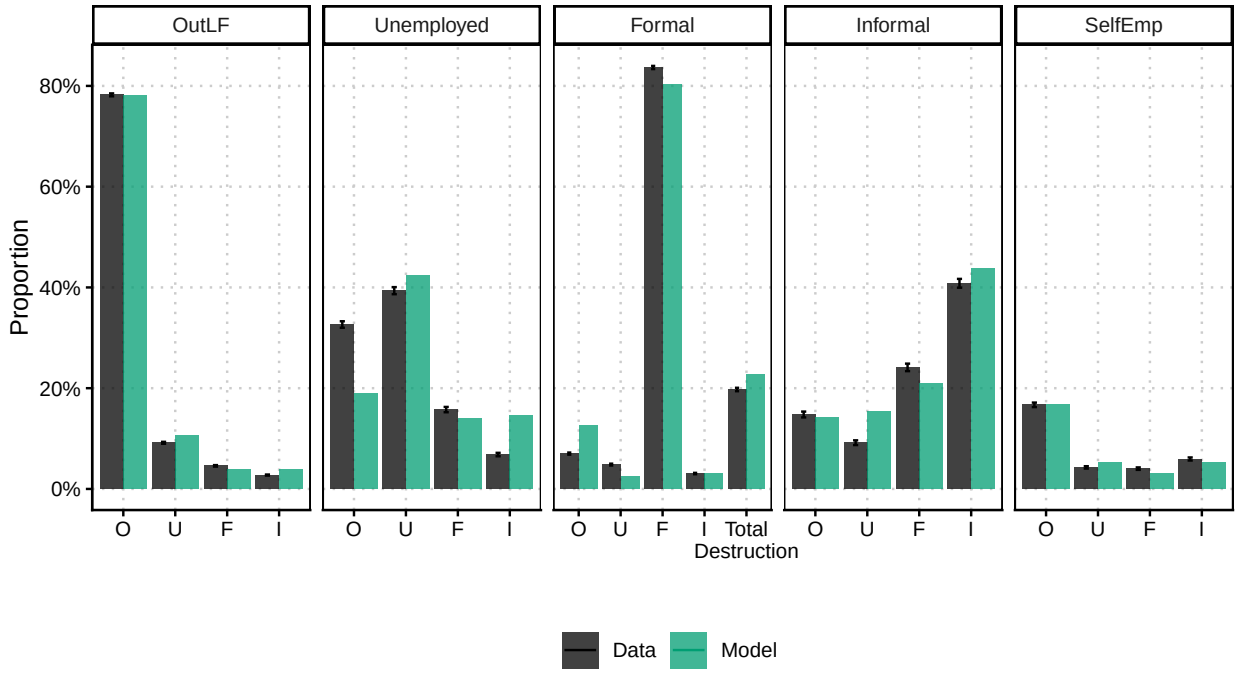
Notes:

Figure A.7: Model fit: Income profile (growth)



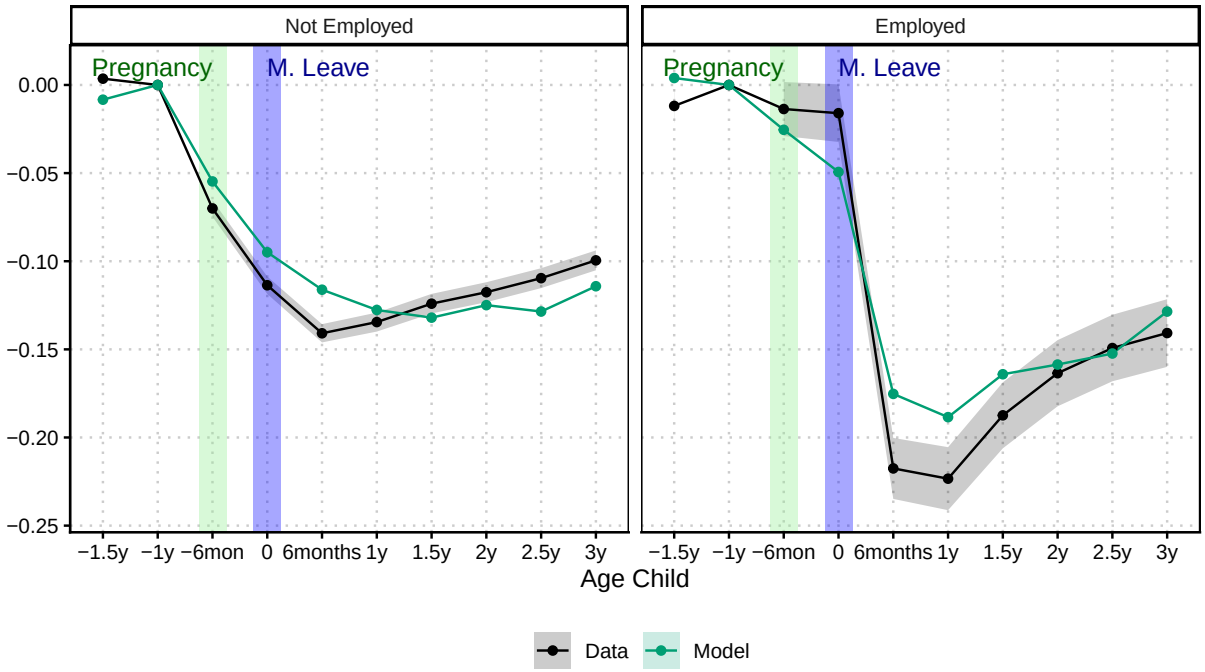
Notes: This figures show the average log-income for individuals working in each sector by age. The figure plots the difference in logs of the current average income and the intercept, being therefore, interpreted as the income growth profile. The black dots plot the data, while the green line plots the model fit. For the data we also present the 95% confidence interval.

Figure A.8: Model fit: Transitions



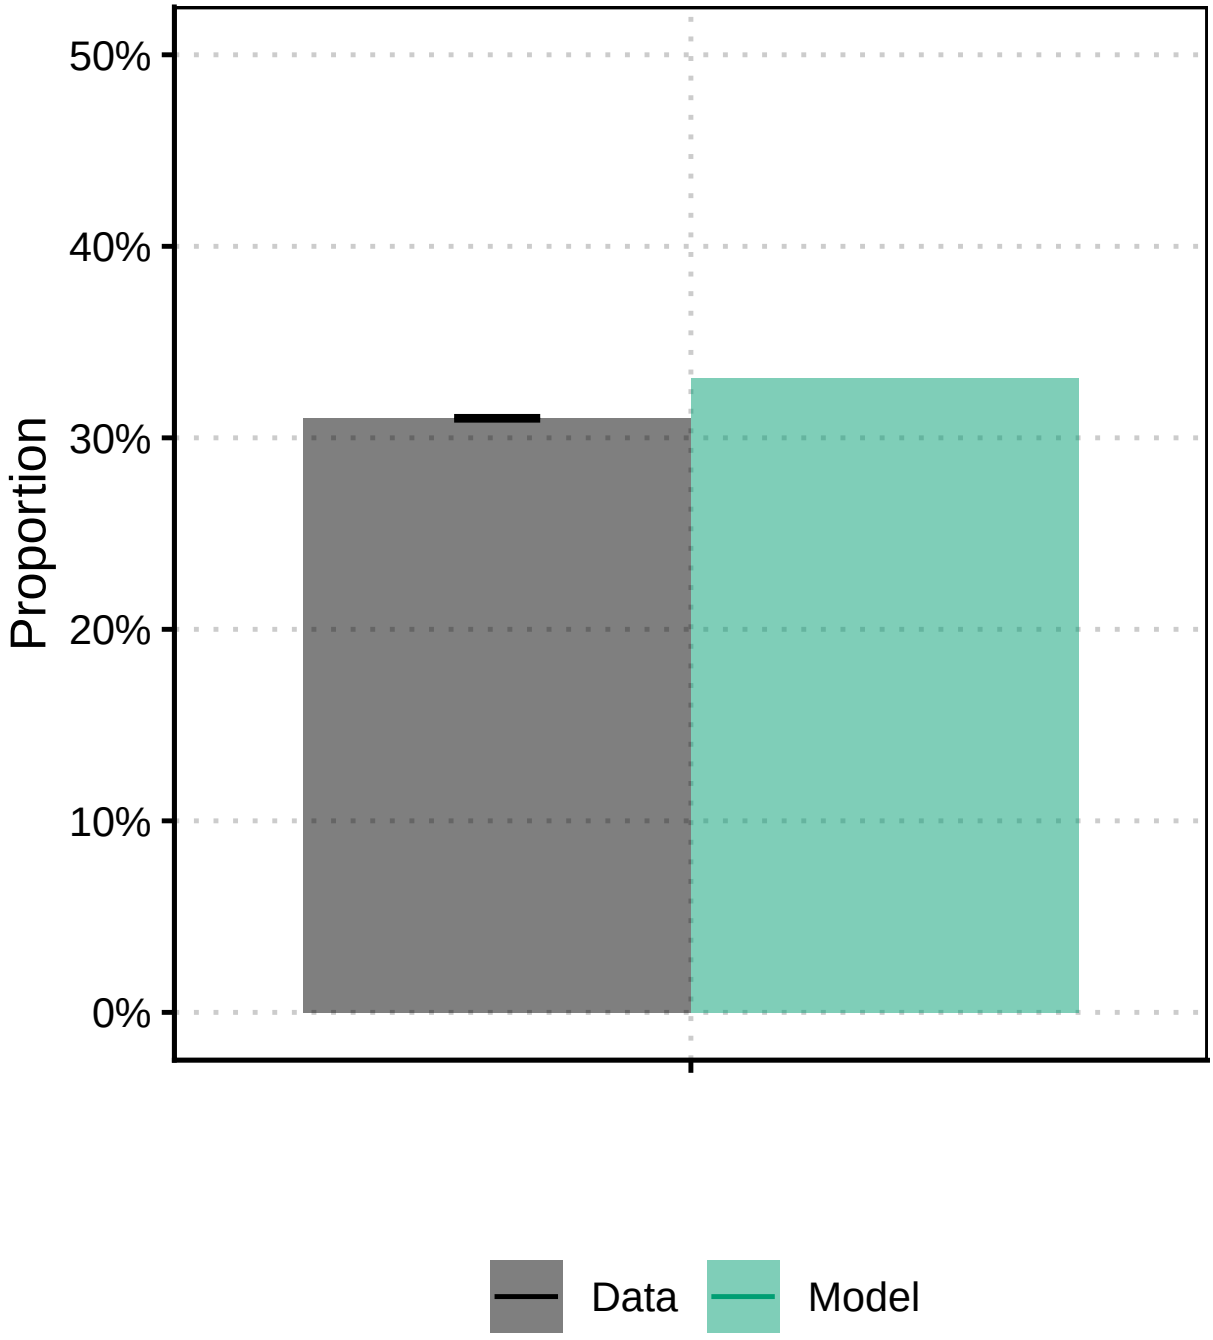
Notes: The figure plots the transitions across all labor market states. Each panel represents the origin status. Therefore, the first panel shows the transitions rates for women that were out of the labor force, towards all the sectors (being out of labor force: O, formal: F, and informal: I). For all panels the omitted category is the probability of moving to self-employment. For the formal sector, we also plot the total destruction probability, that is, the probability of either moving from formal to any other sector, plus the probability of moving from formal to formal but with a different job. The black bars plot the data, while the green bars plots the model fit. For the data we also present the 95% confidence interval.

Figure A.9: Model fit: Reduced-form results of the motherhood penalty, by employment status



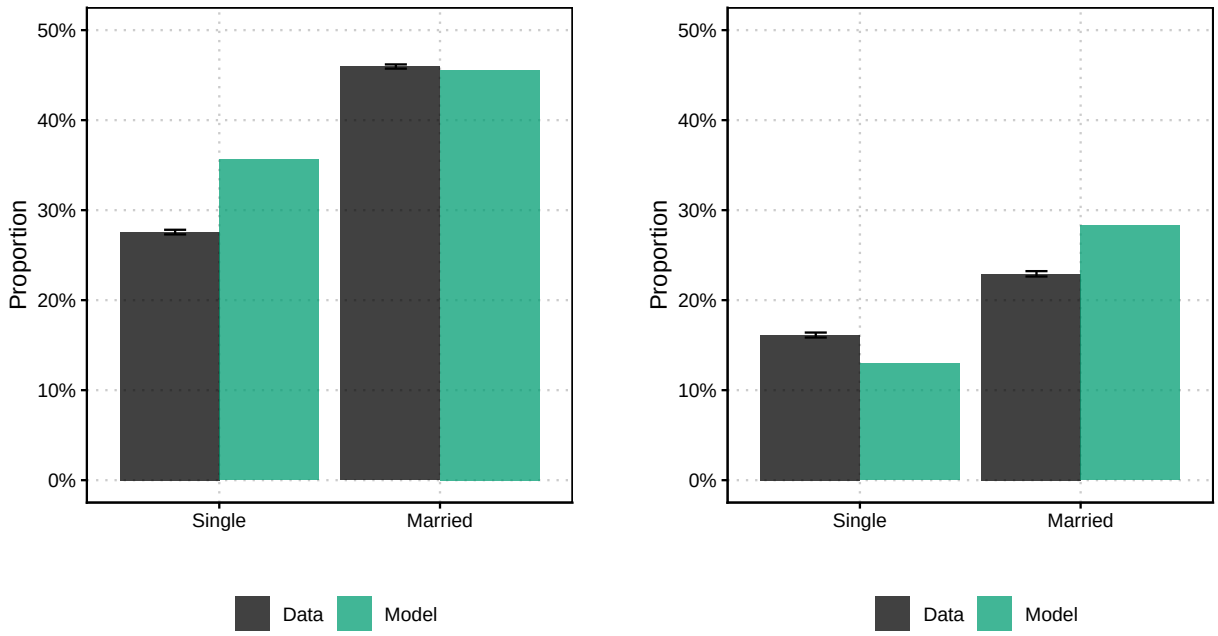
Notes:

Figure A.10: Model fit: Among pregnant women, proportion formally employed at the baseline period



Notes:

Figure A.11: Model fit: Employment patterns by marital status

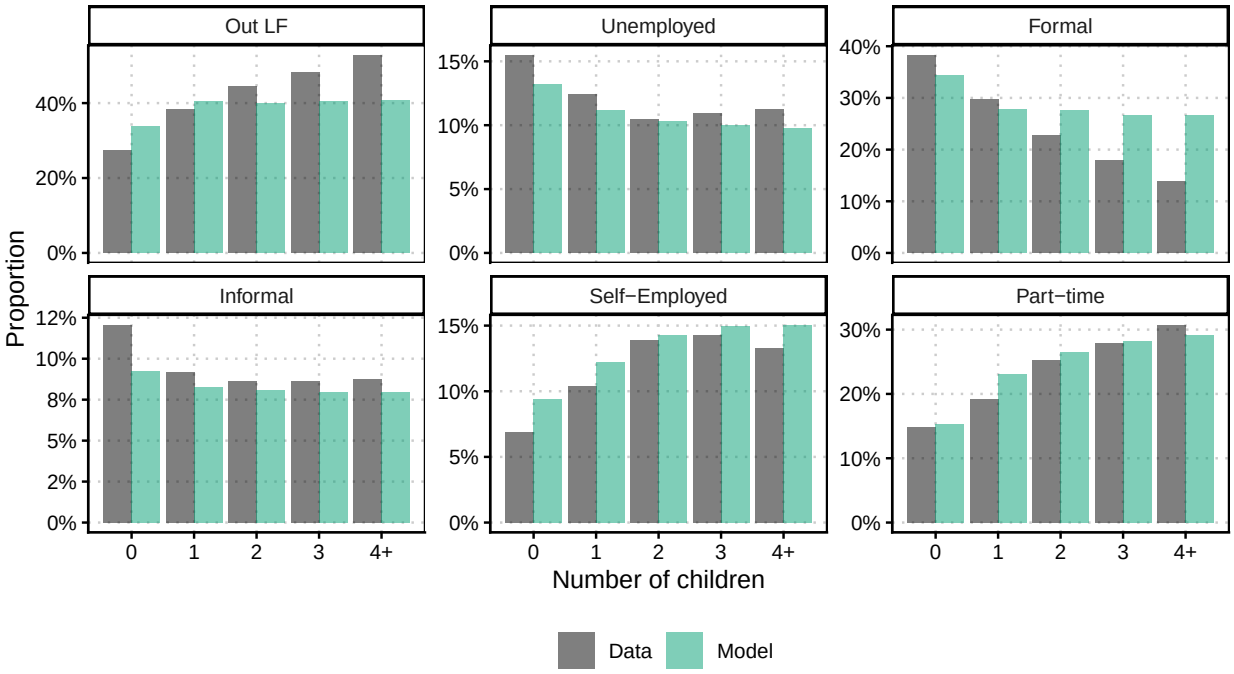


(a) Proportion out of the labor force

(b) Proportion working part-time

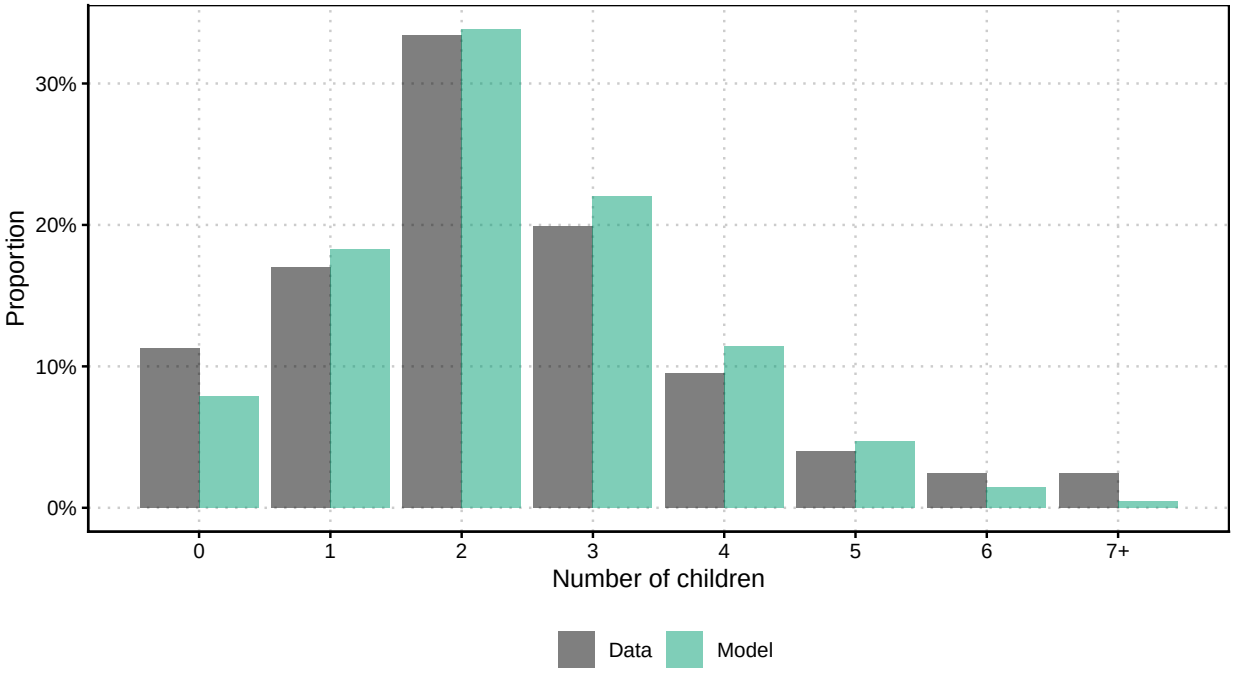
Notes:

Figure A.12: Untargeted moments: Labor market status by number of children



Notes:

Figure A.13: Untargeted moments: complete fertility



Notes:

Table A.1: Moments description

Description	By	Count
Fertility:		
Share with zero children	Age	7
Share with one child	Age	7
Share with two children	Age	7
Share with three children	Age	7
Share with newborns	Marital status	2
Share with second child born x years apart	Distance (in years)	9
Reduced form results: (Motherhood penalty)		
Overall penalty	Age of youngest child	9
Overall penalty - baseline formally employed	Age of youngest child	9
Overall penalty - baseline not formally employed	Age of youngest child	9
Proportion baseline employed	-	1
Labor market:		
Share formally employed	Age	7
Share informally employed	Age	7
Share self-employed employed	Age	7
Share out of the labor force	Marital status	2
Share working part-time	Marital status	2
Formal Earnings	Age	7
Informal Earnings	Age	7
Self-employment Earnings	Age	7
Transition rates from j_{t-1} to j_t	Occupation	5×4+1=21
Labor market & fertility		
Share out of the labor force	Age youngest child	10
Share formally employed	Age youngest child	10
Share informally employed	Age youngest child	10
Share self-employed employed	Age youngest child	10
Share working part-time (all)	Age youngest child	9
Share working part-time (formal)	Age youngest child	9
Share working part-time (informal)	Age youngest child	9
Share working part-time (self-employed)	Age youngest child	9

Notes: Age groups are: [18–21], [22–25], [26–29], [30–33], [34–37], [38–41], [42–45]
Age of youngest child groups are: None, [0–0.5], (0.5–1], (1–2], (2–3], (3–6], (6–9], (9–12], (12–15], (15–18].

Table A.2: Parameter estimates

Parameter	Description	Estimate	Std Error
ν	Consumption weight in the utility function	0.2657	(0.0025)
μ_{ζ_b}	Mean fertility shock	-1.1987	(0.0463)
$\mu_{\zeta_{b,m}}$	Mean fertility shock (married, additional)	0.9289	(0.0757)
$\mu_{\zeta_{b,n}}$	Mean fertility shock ($n \geq 2$, additional)	-0.4693	(0.0418)
γ	Utility cost of not spacing children	0.6265	(0.0274)
μ_{ζ_o}	Mean participation shock	-0.0004	(0.0108)
σ_{ζ_o}	Std dev participation shock	0.9902	(0.0676)
χ	Participation entry cost	2.6178	(0.2013)
δ_f	Exogenous destruction formal	0.0624	(0.0060)
δ_i	Exogenous destruction informal	0.3612	(0.0192)
δ_s	Exogenous destruction self-employed	0.1521	(0.0192)
$\lambda_{u,f}$	Arrival rate from unemployment to formal	0.2840	(0.0142)
$\lambda_{i,f}$	Arrival rate from informal to formal	0.5289	(0.0479)
$\lambda_{s,f}$	Arrival rate from self-employment to formal	0.0200	(0.0039)
$\lambda_{u,i}$	Arrival rate from unemployment to informal	0.2303	(0.0084)
$\lambda_{f,i}$	Arrival rate from formal to informal	0.0703	(0.0111)
$\lambda_{s,i}$	Arrival rate from self-employment to informal	0.0728	(0.0123)
$\lambda_{u,s}$	Arrival rate from unemployment to self-employment	0.3794	(0.0248)
$\lambda_{f,s}$	Arrival rate from formal to self-employment	0.0351	(0.0053)
$\lambda_{i,s}$	Arrival rate from informal to self-employment	0.1180	(0.0259)
λ_f^{pt}	Arrival rate part-time opportunities (formal)	0.0139	(0.0012)
λ_i^{pt}	Arrival rate part-time opportunities (self-employment)	0.1543	(0.0081)
κ	Depressed arrival rates for pregnant women	0.6050	(0.0796)
ω_f	Experience-adjustment to formal arrival rates	0.7317	(0.0715)
ω_s	Experience-adjustment to self-employment arrival rates	8.1185	(1.5256)
ϕ	Utility cost when working with babies	0.3825	(0.0278)
θ	Child sector-specific cost (level)	0.1142	(0.0101)
θ_I	Child sector-specific cost (informal adjustment)	0.5601	(0.0760)
θ_S	Child sector-specific cost (self-employment adjustment)	0.2431	(0.0510)
π	Child sector-specific cost (curvature)	0.4670	(0.0831)
r_1^f	Income experience returns (formal)	0.0490	(0.0015)
r_1^i	Income experience returns (informal)	0.0379	(0.0038)
r_1^s	Income experience returns (self-employment)	0.0849	(0.0032)
Y^m	Marital transfers	1.4989	(0.0821)

Table A.3: Counterfactual exercises

	LF Participation (pp)	Total Formal (pp)	Intensive Formal (pp)	Formal Experience (years)	Total Experience (years)	Consumption (%)	Number of Children (avg)
Baseline	-4.01	-3.76	-2.61	-1.24	-1.38	-8.19	2.33
Formal=Self-Employment							
Fixed cost of work	-4.0	-3.7	-2.5	-1.2	-1.4	-8.0	2.3
Childcare cost	-2.8	-1.8	-0.2	-0.6	-0.7	-4.9	2.5
Part-time option	-3.6	-3.1	-1.2	0.0	-1.3	-7.3	2.4
All	-2.6	-1.8	-0.2	0.0	-0.5	-3.5	2.5

Notes: This table decomposes the *Formal=Self-Employment (Flexibility)* counterfactual into its main components. The first row reproduces the baseline motherhood penalty. The remaining rows sequentially replace specific features of formal employment with those of self-employment: fixed costs of working, sector-specific childcare costs, and part-time option availability. The last row combines all three changes.