

Offsetting the Demographic Drag? Endogenous Female Labor Supply in Latin American Economies*

Camila Casas[†] Nicolas Gomez-Parra[‡] Flavien Moreau[§]

September 14, 2026

Abstract

The share of Latin America’s working-age population has begun to fall, on the back of increased longevity and falling birth rates. The female participation response to falling child dependency is the largest domestic margin available to offset it. We quantify that margin, combining harmonized microdata from seven large economies with a tractable overlapping generations model in which female labor supply responds endogenously to household composition. We first decompose the region’s gender participation gap into several margins, with about one third arising from the differential labor supply response to household-composition between men and women (a ”motherhood penalty”). Embedding these micro-estimates in a general equilibrium framework, we find that the female participation response to falling child dependency can partially offset the demographic drag, raising regional GDP by between 3.1 and 11.3 percent by 2050 relative to a path with participation held fixed. The range reflects both fundamental uncertainty around future paths of comoving national trends and substantial heterogeneous labor supply responses across countries.

KEYWORDS: Demography, Labor Force Participation, Gender Gaps, Aging, Growth.

JEL CODES: J11, J21, J22, O4, O15

*Authors would like to thank Raquel Fernández, Inés Berniell, Gabriel Ulyssea, and Santiago Levy for enlightening conversations. We also thank Gustavo Adler, Lusine Lusinyan, Jiajia Gu, Ana Paula Nothen Ruhe, Cassie Chen Xiang, Manuk Ghazanchyan, Duncan MacDonald, and Júlia Cots-Capell, as well as seminar participants at IMF Western Hemisphere Department, Banxico, Banco de la República, LACEA, and Barcelona Summer Forum. We acknowledge support from Genevieve Lindow, and Kenji Moreno. We thank Petia Topalova for extensive and insightful comments, as well as Adrien Auclert and Hannes Malmberg for help with the sequence-space Jacobian. The work in this paper is partly supported by the Macroeconomic Research on Climate Change and Emerging Risks in Asia program of the Ministry of Economy and Finance of the Government of Korea. The views expressed in this article are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management. We also thank, without implicating, national statistical agencies for disseminating granular series to researchers.

[†]International Monetary Fund. E-mail: ccasas@imf.org.

[‡]International Monetary Fund. E-mail: ngomezparra@imf.org.

[§]International Monetary Fund. E-mail: fmoreau@imf.org

1 Introduction

The share of Latin America’s working-age population has already started declining, while the growth of female labor market participation has markedly slowed down. Women joining the labor force added more than ten percentage points to the region’s participation rate between the 1950s and 2010. Since then the rate has plateaued, with a gender gap still roughly twice the advanced-economy average. As fertility falls and populations age across middle-income countries, the standard demographic-dividend logic—in which a falling youth dependency ratio mechanically raises labor input per capita—reverses, and the working-age share begins to contract.¹ Whether this reversal translates into a sustained drag on growth depends critically on how household labor supply responds to the demographic transition. If female participation rises endogenously with declining child dependency ratios, the demographic drag may be partially offset; if it does not, the drag is amplified. Quantifying this endogenous response, and disentangling it from the policy levers that have historically been credited with rising female participation is crucially important to gauge future economic prospects and the extent to which they can be shaped by policies. This requires combining micro evidence on the household drivers of female labor force participation (FLFP) with a macro framework that maps those drivers into aggregate labor supply.

This paper provides such a quantification for Latin America, a region that is well advanced in the demographic transition. We combine harmonized microdata from seven of the region’s largest economies—Argentina, Brazil, Chile, Colombia, Mexico, Peru, and Uruguay (LA7)—with a tractable overlapping generations model in which female labor supply depends on household composition through a closed-form expression for the female-to-male labor ratio, and solve the model’s full perfect-foresight transition.

The participation response to falling child dependency raises LA7 output in 2050 by between 3.1 percent and 11.3 percent across specifications, relative to a path with participation held at its initial profile and fertility rates decline following UN projections. What separates the two estimates is whether the national secular trend in female participation—which can be attributed to factors such as rising education, falling discrimination, and changing norms, and which moved together with falling child dependency over our estimation window—is properly accounted for. This distinction is one of the paper’s results. The lower figure is obtained when the elasticity of home production to child dependency is estimated *net* of these national trends and therefore represents the proper demographic channel. Cross-sectional variation in dependency moves together with the secular trends over our estimation window and they cannot be separated empirically. These two forces are controlled for in our structural model by two parameters: the elasticity to child dependency and the pace of gender norms evolution. We find that in most cases the labor supply response to declining child dependency is substantial but temporary, a feature of the second half of the demographic transition rather than a permanent dividend. On the baseline trajectory, the demographic drag is only partially self-limiting.

¹See [Bloom et al. \(2010\)](#) on the dividend phase and [Maestas et al. \(2023\)](#) on the drag.

The main contributions of our paper are threefold. First, we develop a theoretical framework in which the female labor force participation is a function of the child dependence ratio, the old-age dependency ratio and a wage penalty that captures statistical discrimination, monoposony and the like, reflecting micro patterns well documented in the literature on the region (e.g. marriage premium, motherhood penalty, returns to education as [Berniell et al. \(2024\)](#); [Kleven et al. \(2023\)](#)). Second, we develop a methodology that allows us to disentangle the impact of child-dependency on the level vs change of FLFP, a technically difficult task. This allows us to use our model to forecast FLFP assuming perfect foresight across our sample of LA7 using UN’s Population Prospects fertility projects over the next 3 decades. Third, we recover an estimate of the social norm parameter from the model’s calibration, which we compare to the ordering observed in advanced economies. The close correlation between our estimates and the data suggests our model is able to correctly pin down this key parameter.

Identification and methodology. Our identification strategy combines within-country variation in household composition, which identifies the structural parameters of the labor supply equation, with cross-country and cross-time variation in fertility and aging, which identifies the demographic channels in the macro framework. The wage markdown is calibrated to match the Mincerian wage gap by country. The elasticity of substitution between female and male labor in the Constant-Elasticity-of-Substitution (CES) production function is estimated from sub-national variation in the female-to-male wage ratio, an approach we discuss in detail in [Section 5](#), together with its limitations.

Related literature Our work relates to four main strands of the literature. First, we contribute to the extensive literature documenting the secular relationship between economic development and female labor force participation, originating with [Lewis \(1954\)](#)’s seminal work in development economics. When unpaid work is properly accounted for, the development-FLFP relationship exhibits a U-shaped pattern ([Goldin, 2014](#)): FLFP initially declines as agriculture’s share of GDP falls, then rises alongside the expansion of services. [Ngai et al. \(2024\)](#) examine the interaction between structural transformation and FLFP using long-run U.S. time series, finding that sectoral shifts play a crucial role—particularly given evidence of women’s comparative advantage in service occupations ([Ngai and Petrongolo, 2017a](#)). Moreover, the omission of unpaid home labor, especially in small-scale farming, substantially understates women’s economic contributions in historical contexts such as the late nineteenth-century United States. [Acemoglu et al. \(2004\)](#) analyze how rising female labor force participation affects wages for both women and men. We extend these findings, which have focused primarily on the United States, by studying a panel of major Latin American economies within a unified analytical framework. In addition to the channels emphasized by previous studies, we focus on the endogenous response of FLFP to demographic changes that typically accompany economic development.

Second, we contribute to a growing macroeconomic literature examining the effects of population

aging on economic growth. This literature has focused predominantly on western economies, where demographic aging is most pronounced.² [Maestas et al. \(2023\)](#) find that aging significantly reduced U.S. GDP growth per capita by 0.3 percentage points annually between 1980 and 2010. [Fernández-Villaverde et al. \(2023\)](#), [Aksoy et al. \(2019\)](#), and [Jones \(2022\)](#) explore the implications of population slowdowns for long-run growth, while [Cravino et al. \(2022\)](#) quantify how aging increases the service share of consumption. [Jaimovich and Siu \(2009\)](#) demonstrate that demographic change accounts for one-fifth to one-third of the decline in U.S. output volatility. [Acemoglu and Restrepo \(2017\)](#) and [Acemoglu and Restrepo \(2022\)](#) link automation to population aging.³ [Kotschy and Bloom \(2023\)](#) emphasize that expanding labor supply can cushion the demographic drag on growth, but do not analyze how demographic structure interacts with endogenous labor supply decisions as we do here. Classic work on the demographic dividend ([Bloom and Williamson, 1998](#); [Bloom et al., 2010](#)) documents that falling dependency ratios create growth windfalls in developing economies; our paper analyses Latin America’s transition from this dividend phase into a structural drag. [Cavalcanti and Tavares \(2016\)](#) quantify the macro cost of gender wage gaps for long-run income—a closely related analysis that we complement by studying the endogenous labor supply channel directly.

Third, given that changes in FLFP unfold gradually across successive cohorts, we adopt an overlapping generations (OLG) framework in the tradition of [Allais \(1947\)](#), [Yaari \(1965\)](#), and [Blanchard \(1985\)](#). Specifically, we build on the class of OLG models developed to analyze the long-run fiscal implications of demographic change ([Auerbach and Kotlikoff, 1987](#)). Our main contribution is to endogenize female labor supply within the framework of [Auclert et al. \(2021\)](#), thereby connecting FLFP dynamics to demographic shifts. This extension highlights a labor supply channel that partially offsets, in the medium run, the adverse effects of declining fertility on long-run economic growth.⁴

Finally, our paper is also related to a fast-growing microeconomic and labor literature measuring the extent and drivers of gender participation gaps in Latin America. [Berniell et al. \(2024\)](#) provide a comprehensive survey of the recent literature.⁵ Full equal participation has often been prevented by traditional social norms ([Jayachandran, 2020](#)). In particular, causal evidence using event studies has highlighted the role of the child penalty ([Berniell et al., 2023](#)). We build on these findings and show how these gaps add up in a macroeconomic setting. The aggregate income cost of gender gaps in occupational allocation is quantified by [Cuberes and Teignier \(2016\)](#), whose GDP-level estimates complement our labor supply analysis. We generate scenarios consistent with the micro-

²See [Batog et al. \(2019\)](#) for an analysis of demographic headwinds in Eastern and Central Europe.

³Our analysis focuses on labor supply, abstracting from potentially evolving labor demand. For the impact of slower population growth on firm dynamics and entry, see [Karahana et al. \(2024\)](#) and [Hopenhayn et al. \(2022\)](#), who employ neoclassical models with endogenous entry. [Peters and Walsh \(2021\)](#) study a model where both entry and incumbent expansion are endogenous, analyzing the consequences of declining population growth.

⁴A similar offsetting mechanism appears in [Bloom et al. \(2009\)](#), though their static framework abstracts from intertemporal consumption-savings decisions. Unlike [Galor and Weil \(1996\)](#) and [Doepke and Tertilt \(2009\)](#), who endogenize fertility choices to explain the demographic transition in pre-modern economies, we incorporate household heterogeneity while taking aggregate child dependency ratios as exogenous, consistent with UN population forecasts.

⁵See [Edjigu et al. \(2026\)](#) and references therein for a survey of reforms aiming at boosting FLFP.

literature to study how policies and changes in the demographic structure can contribute to closing participation gaps.

Roadmap. Section 2 documents the long-run demographic and FLFP trends in the LAC sample and presents the growth-accounting decomposition that motivates the structural analysis. Section 3 presents the micro evidence on the household drivers of FLFP across the seven countries. Section 4 sets out the structural model and its closed-form benchmark. Section 5 brings the model to the data: it estimates the demographic elasticity, calculates the range of model-supported contribution of UN-projected demographic changes to GDP levels up until 2050, and validates the model’s norm parameter against evidence from advanced economies. Section 6 concludes.

2 Anatomy of a fall: Latin America’s declining demographic dynamism

Between the 1950s and 2010, labor force participation in Latin America rose by more than 10 percentage points—and nearly 25 percentage points among the prime working-age (25–54) population—driven by women joining the labor force (Figure 2a). However, participation rates have plateaued since then at uneven levels across gender and age groups, while the share of the working-age population has reached a turning point.

2.1 Demographic trends and stylized facts

Baseline scenario Latin America’s major economies have been undergoing a rapid demographic transformation since 1950, with population growth fast decreasing and projected to become negative by the end of the century. Figure 1a shows population growth rates exhibiting a consistent downward trend, declining from peaks of 2.5–3.5 percent in the mid-20th century to projected negative rates by the end of the 21st century. Figure A4 depicts the evolution of median age, revealing a dramatic aging process across the region. Starting from median ages of 15–20 years in 1950, all countries show substantial aging, with projected median ages reaching 40–55 years by 2100. Uruguay’s demographic pattern stands distinct from its regional peers, characterized by lower initial population growth rates and a consistently higher median age, indicating its earlier progression through the demographic transition relative to other Latin American nations.

These dynamics will translate into rising dependency ratios. Figure 1d presents the total dependency ratio, which peaked in most countries during the 1960s and 1970s at around 80–100 percent, primarily driven by high youth dependency (Figure 1b). This ratio subsequently declined to historical lows of approximately 45 percent by 2020, reflecting the demographic dividend period. However, projections indicate a significant reversal of this trend, with total dependency ratios expected to rise to 70–90 percent by 2100. Figure 1c shows the old age dependency ratio, revealing a dramatic transformation in the composition of demographic dependence.⁶ Starting from relatively low levels

⁶A more sanguine view of demographic changes focuses on prospective age and functional aging. Through that lens, the impact of population aging for the economy may be less dramatic (Sanderson and Scherbov, 2007; Kotschy




Figure 1: LAC population is rapidly aging

Note: UN Population Prospects. Authors' calculations for 2022-2100 based on median projections. The total dependency ratio is calculated as the ratio of population under 14 and over 65 to population 15-64, the young-age dependency ratio as the ratio of population under 15 to the working-age population (15-64), and the old-age dependency ratio as the ratio of population 65 and over to the working-age population (15-64).




Figure 2: Female labor force participation in Latin America

Sources: ILOSTAT and IMF staff calculations. Participation rates based on working-age population (15 and older). 1951-2023 series smoothed geometrically. AEs = advanced economies; Africa = Sub-Saharan Africa; EMDE = emerging market and developing economies; LAC = Latin America and the Caribbean; MECA = Middle East and Central Asia.

of 5-10 percent in 1950, all countries exhibit a sharp acceleration in old-age dependency, particularly post-2020, with ratios projected to reach 45-70 percent by 2100. Chile and Uruguay lead this aging process, with substantial heterogeneity in the timing and pace of demographic transition across the region based on UN's projections.

Downside and upside scenarios Perhaps more worrying, these baseline scenarios may paint an overly optimistic picture. Each scenario rests on a combination of assumptions on the evolution of fertility rates, mortality rates, and migration.⁷ The UN World Population Prospects provides a range of scenarios corresponding to a large range of assumptions, including several downside scenarios, with significantly lower fertility rates. In fact, recent demographics outturns have concentrated around the lower range of the alternative scenarios. We hold to the median variant throughout; the counterfactuals of Section 5.4 vary participation rather than fertility, leaving the downside fertility variants for future work.

2.2 Long term evolution of LFP

Female LFP ratios have risen enormously over the last decades (Figure 2a). Yet large disparities in participation rates persist across the region., mostly reflecting heterogeneity in FLFP. Men's prime working age (25–54 years old) participation rates are uniformly high across the region, in most cases

et al., 2024). See also Lee et al. (2014) for a fiscal argument in favor of lower fertility rates when social transfers are limited.

⁷See Onofri et al. (2026) for an analysis of the evolution of fertility rates in LAC




Figure 3: Decomposing the contribution of labor force participation to growth

Note: Contributions to real GDP growth are expressed in percentage points. LA7 = Latin America 7 (Argentina, Brazil, Chile, Colombia, Mexico, Peru, Uruguay). Sources: ILOSTAT; IMF Gender Data Hub; and IMF staff calculations. Notes: ILOSTAT; Penn World Table 10.0 database; 2024 UN Population Prospects; and IMF staff calculations. Working-age population is defined as 15 and older. Aggregates are simple averages based on modeled ILO estimates.

above the levels observed in advanced economies. By contrast, female participation rates exhibit greater heterogeneity both in levels and across age groups and are generally lower than those in advanced economies. Uruguay stands out as the country in the region with the highest female labor force participation rate, reaching over 85 percent among women 35–44.

Latin America’s average gender gap in participation is now close to 20 percentage points, twice as large as the average gap across advanced economies (Figure A7a). Within the region, however, there are significant cross-country differences: while Uruguay’s participation gap is 12 percentage points, just a few percentage points larger than those of advanced economies, Mexico’s is about 30 percentage points. And, while Mexico and Chile have reduced their gaps by 15 percentage points since 2000, Colombia’s gap has shrunk by only 5 percentage points since then.

2.3 Growth decomposition

We compute the contribution of capital, labor, and technological change to total output growth.⁸ Let Y_{ct} , K_{ct} , L_{ct} , and Z_{ct} denote, respectively, GDP, labor, capital, and labor-augmenting productivity of country c at time t . We assume aggregate output follows a Cobb-Douglas production function in which α_c denotes the capital share,⁹ therefore

$$Y_{ct} = K_{ct}^{\alpha_c} (Z_{ct} L_{ct})^{1-\alpha_c}. \quad (1)$$

⁸We abstract from international trade, as, contrary to many fast-growing Eastern Asian countries, Latin American countries are less open to trade and did not experience similar export-led growth Parente and Moreau (2024).

⁹See Appendix C for the general framework with CES production function.

Letting y , l , k , and z denote, respectively, the logarithm of GDP, labor, capital, and productivity of country c at time t , we therefore have the following growth identity:

$$\Delta y_{ct} = \alpha_c \Delta k_{ct} + (1 - \alpha_c) \Delta z_{ct} + (1 - \alpha_c) \Delta l_{ct} \quad (2)$$

The growth of the labor force can, in turn, be expressed as

$$\Delta l_{ct} = \Delta n_{ct} + \Delta wap_{ct} + \Delta lfpr_{ct} + \Delta e_{ct} \quad (3)$$

where N is the total population, WAP is the share of working-age population (those 15 and older), $LFPR$ is the labor force participation rate, and E is the employment rate (the share of the labor force currently employed). We measure K_{ct} and α_c using the Penn World Table (PWT); α_c is one minus the PWT labor share.

Figure 3b shows the decomposition of labor force growth across the 7 major Latin American countries since 2000. A large portion of labor growth has come from demographic factors: population growth has been the largest contributor, followed by a demographic shift that has led to a sustained increase in the share of the working-age population, referred to as the “demographic dividend.” In contrast, labor market outcomes, namely the participation and employment rates, have played a more limited role. Modest increases in the overall participation rate during 2000–10 followed by minor decreases thereafter resulted in a negligible contribution of participation rates over the past two decades in most countries, while increases in employment rates can explain roughly 0.3 percentage points of the labor force growth on average. Overall, these positive developments in demographic composition and labor outcomes have ensured that GDP has grown faster than population, hence GDP per capita has risen too.

Labor force growth across Latin American economies still exhibits some notable variations in its underlying drivers (Figure 3b). Colombia exhibits the highest overall labor growth rate at approximately 2.3 percent, driven primarily by demographic factors and increased employment rates, while Uruguay shows the lowest at around 0.8 percent. Population growth contributes significantly to labor force expansion in Mexico and Peru, accounting for roughly half of their total labor growth. Argentina and Chile demonstrate moderate labor growth rates of approximately 2 percent, though through different channels - Argentina’s growth is predominantly driven by employment rate increases, while Chile’s expansion is more evenly distributed across demographic factors and labor force participation. Brazil presents a balanced profile with relatively equal contributions from population growth, demographics, and employment rates, though its overall labor growth rate remains somewhat lower than its regional peers.

2.4 Labor growth contributions by gender

The contributions to labor growth can be further decomposed by gender. For each gender $g = m, f$ we have the following decomposition

$$\Delta l_{ct}^g = \Delta n_{ct}^g + \Delta wap_{ct}^g + \Delta lfp_{ct}^g + \Delta e_{ct}^g \quad (4)$$

For simplicity, assuming that female and male labor are substitute, we can describe the differences between female and male contributions.¹⁰ We start in 1980 due to data limitations.¹¹ Figure A1 breaks down the sources of labor growth by country and gender. Positive contributions from population and demographic changes were evident in most periods. For example, in Argentina and Colombia, large positive demographic shifts—exceeding 10 percentage points—boosted labor growth significantly in earlier decades. However, these favorable effects were increasingly offset by negative contributions from labor force participation and, to a lesser extent, employment rates in the more recent period (2011–2020).

The decomposition reveals distinct gender patterns. In several countries, such as Brazil, Mexico and Chile, female labor growth initially outpaced male labor growth, with females benefiting from relatively stronger positive contributions in certain periods. In fact, during the 1980–2000 period, the differential was largely driven by a marked increase in female labor force participation and demographics, even though employment rates did not exhibit significant gains for either gender. Yet, by 2011–2020, both genders experienced a sharp slowdown or even negative growth as mentioned previously. While past demographic dividends played a key role in expanding the labor force, recent declines in participation and employment dynamics may dampen future labor supply.

2.5 Household composition

The historical patterns suggest that household composition plays an important role in the rise of LFP: most of the participation gains came from married women, regardless of the number of children. The data reveal broad regularities across Latin American countries, similar to patterns described by Goldin (2006) and Olsson (2025). The ranking order of participation rates across groups remained consistent throughout the period. Married men maintained high labor force participation rates close to unity, while single men exhibited high but slightly eroding participation rates.¹² Single women started with participation rates around 40 percent in the 1960s and crossed over 50 percent in most countries by the end of the period. Married women’s labor force participation rates, typically around 25 percent in the 1960s, doubled over the subsequent decades.

Using harmonized IPUMS series, female participation rose by about 27 percentage points for married women, against 9 percentage points for single women, averaged over the five countries with

¹⁰Section 4 present a richer model allowing for imperfect substitution.

¹¹See de la Escosura (2007) for evidence of rising LFP in earlier decades. Auclert et al. (2021) use a linear labor aggregator.

¹²Both ‘push’ and ‘pull’ factors have played a role in the United States, in particular skills mismatch, caretaking responsibilities, and prolonged continuing education (Bengali et al., 2023).

comparable data.¹³ The country-level changes are reported in Table 1. The widening gap between married and single women’s participation rates is most pronounced in Brazil and Uruguay, where married-women participation increased by more than 30 percentage points while single-women participation rose by less than 10. Mexico and Chile show large gains in both groups, with married women still leading; Colombia is the exception, with limited movement on either margin. These patterns underscore that the headline rise in female participation in Latin America has been driven primarily by married women joining the labor force—consistent with the household-composition mechanisms documented in our probit analysis below.

Table 1: Change in women’s labor force participation, 2020 vs. 1980

Country	Married (p.p.)	Single (p.p.)
Brazil	+33.2	+3.4
Chile	+31.1	+16.8
Colombia	+7.2	-0.3
Mexico	+27.2	+19.1
Uruguay	+34.4	+7.7
<i>Average</i> [*]	<i>+27</i>	<i>+9</i>

Sources: harmonized IPUMS series; LIS used as a bridge to project missing 2020 IPUMS cells by applying growth rates from the closest comparable LIS survey, followed by linear interpolation back to 1980. Argentina and Peru excluded due to incomplete comparable data. ^{*}Simple average across the five countries shown.

3 The micro origins of LFP gaps

We now study the details of participation gaps across groups. Participation rates are also uneven across age groups (Figure 2b). Although a decline with age is expected, female participation rates notably start declining from a relatively early age as women in Latin America exit the labor force sooner than men, and sooner than women in advanced economies—in some cases, even while still in their prime working age. Moreover, participation rates among late working-age women lag those of Asian and European peer cohorts, with the differences becoming more notable in recent years when participation increased elsewhere while it remained flat in Latin America. In contrast, participation among men during late working-age years is in line with, and even slightly above, that in Asian and European peers. While there is a sharp decline for the eldest group of workers (65 and older) across all countries, the participation rate of men in Latin America remains relatively high—at around 40 percent—indicating that a significant fraction of this demographic does not retire when reaching official retirement age, a feature also observed in Asia.

¹³Missing 2020 observations are filled by applying growth rates from the closest comparable LIS survey to each country-cell. Argentina and Peru are dropped for want of close comparators.

3.1 A probit decomposition

Despite progress over the past decades, women in Latin America continue to face numerous obstacles to labor force participation, with unpaid work, including household work, explaining most of the remaining participation gap. Participation gaps widen after parenthood, as women often permanently exit the labor force while men’s participation remains largely unaffected. There is, however, a substantial heterogeneity both across and within countries, in particular along education levels.

Which agents’ characteristics can explain labor force participation decisions? We use simple probit regressions to explain the role of features such as household composition (whether a person is the head household, married, or has young children or older relatives living at home), education, and geographical domain (rural/urban) as potential drivers of participation for men and women separately.¹⁴ Suppose that a labor market outcome can be characterized by the following data generating process:

$$LFPR_i = \alpha + \beta_1 Head_i + \beta_2 Married_i + \beta_3 Young_children_i + \beta_4 Children_i + \beta_5 Elders_i + \beta_6 Education_i + \beta_7 Age_i + \beta_8 Age_i^2 + \beta_9 Rural_i + \varepsilon_i. \quad (5)$$

The mean outcome for a given group of workers at time t can be predicted using the group-specific estimates of equation 5 (obtained using each group’s subsample) and averages.

Results are presented in Figure 4, with full probit coefficients for all seven countries reported in Appendix Table A3. We find that some features have similar effects across all agents—for instance, the probability of being active in the labor market increases for both men and women with education, age, or when they are the head of their household. However, family characteristics play different roles for each group. While being partnered/married or having young children lowers the probability of women entering the labor force, it increases that of men. These patterns are broadly consistent across Argentina, Brazil, Chile, Colombia, Mexico, Peru, and Uruguay, though the magnitude of the motherhood penalty varies considerably across countries.

3.2 Institutional and social factors

The role of household responsibilities Household responsibilities play a key role in keeping women out of the labor force, especially among women with lower educational attainment. Women are systematically more likely to report in household surveys that household responsibilities, including childcare and elderly care are the main reason why they do not look for a job (Figure A6b), with some variation across countries. For instance, in Mexico, household work is the reported reason for inactivity for about two thirds of women who are out of the labor force as opposed to only about a tenth of men. This cross-country variation is consistent with differences in attitudes towards gender roles as measured, for instance, by Latinobarómetro. In every country in the sample, women are

¹⁴See Grigoli et al. (2020) for an analysis of the drivers of labor force participation in European economies.




Figure 4: Micro-drivers of LFP

Sources: IPUMS (Ruggles et al., 2025); national household surveys; and authors' calculations. Notes: average marginal effects from the probit model of equation (5), estimated separately for men and women on the working-age population (15 and older) in each of the seven countries. Each facet reports one regressor; markers are country point estimates and the dashed line is the unweighted seven-country average. Facets are ordered by increasing average marginal effect on female participation. Full coefficients and standard errors are in Appendix Table A3.

at least five times more likely than men to report household work as the main reason for inactivity. Within country, female labor force participation increases significantly with educational attainment (Figure A7b). For instance, in Brazil, women with educational attainment lower than high school are more than five times more likely to stay inactive because of household work compared to women with higher education who face a higher opportunity cost of staying home.

Unpaid work Women are the largest contributors to (unmeasured) household output.¹⁵ Once domestic responsibilities are taken into consideration, female participation rates are similar to (or higher than) men’s (Figure A2). However, from a national accounting perspective, unpaid household services produced and consumed by households are generally omitted in the scope of production while similar tasks performed by paid domestic workers are.¹⁶ Unpaid care and domestic work are estimated to range between 15 percent (Ecuador) and 24 percent of GDP (Mexico), when hours of domestic work are valued at their corresponding market price. Advancements in accounting unpaid household services (as part of the updates to the international statistical standards) could shed light on a more comprehensive measure of labor input and economic growth.

As a result of women’s larger share of household duties, female labor supply is more income-elastic than male labor supply. This is consistent with evidence that exogenous shifts in household non-labor income—for instance via remittances—generate larger participation responses for women than for men.¹⁷ The same higher elasticity also shows up in differences in cross-country labor force participation responses to tax-and-benefit design: female participation is more sensitive than male participation to the degree of individualization and progressivity of personal income tax systems, with notable cross-country heterogeneity in LAC consistent with the variation in joint-filing and family-based provisions documented in Leyva and Urrutia (2020).

Causal Effects of Parenthood Female labor force participation drops sharply around motherhood, while male labor force participation is largely unaffected or slightly increases (Figure A8). The pattern is broadly similar across the region although varies with income levels. The drop in participation rates in the year after childbirth is smaller and less pronounced (at least in the initial year) in higher-income countries (Uruguay, Chile), more in line with advanced economies with more extensive social protection systems. In middle-income countries, especially in Brazil and Colombia, the drop is sharper and plateaus around 20 percent, even comparing with peer countries in emerging Asia, where the motherhood effect is twice smaller on average at a five-year horizon (Kleven et al., 2023). In the region, motherhood effects are estimated to account for about a third of the overall gender participation gap (Kleven et al., 2023).

¹⁵See Casas and Moreau (2024) for an in-depth analysis of patterns of unpaid work in Latin America based on household surveys.

¹⁶In many contexts, the boundaries between paid and unpaid labor categories can be porous, for instance when various forms of non-wage labor – subsistence farming, informality, etc. – are prevalent. This calls for caution in the use of binary variables like female LFP as a single measure of labor (Burnette, 2025).

¹⁷Amuedo-Dorantes and Pozo (2006) and Hanson (2007) find that remittances raise the reservation wage of non-migrating women and reduce domestic female labor force participation. The asymmetric responses by gender are consistent with the household-bargaining structure of our framework in Section 4.

The motherhood effect is also heterogeneous within countries and across occupations. It reflects a host of factors, including cultural norms, availability of alternatives and personal preferences, but also legal/tax incentives such as tax codes that penalize dual earner households and asymmetries in parental roles can play an important role. Job flexibility is another important dimension for parents, and especially for mothers (Berniell et al., 2023), and this is closely connected to the issue of informality as formal and informal jobs differ in the protection, parental leave benefits, and the flexibility of hours. This can lead to changing informality rates around motherhood. In many countries, for instance in Peru (Figure A8), informality among women drops around motherhood, as mothers in a formal job can use maternal leaves to stay in the labor force, while women in an informal job are typically not granted the same protection and need to exit of the labor force. However, when re-entering the job market, mothers are more likely to take an informal job which tends to allow for more flexible hours (Berniell et al., 2023) or due to reduced job opportunities and higher search frictions. Higher informality rates for women after childbirth are also observed in countries with lower overall formality rates such as Chile (Berniell et al., 2021).¹⁸

Childcare availability and female labor supply A large quasi-experimental literature has estimated the female labor supply response to exogenous expansions in childcare availability and quality. Estimates from preschool expansions in Argentina, Chile, and Uruguay (Berlinski et al., 2011; Martínez and Perticará, 2017; Padilla-Romo and Cabrera-Hernández, 2019; International Monetary Fund, 2023) and from the extension of the school day in Mexico place the elasticity of eligible women’s participation with respect to formal childcare access in the 5–15 percentage point range. These estimates discipline the calibration of the home-production cost parameter χ in the structural model of Section 4, where childcare availability is captured as a shifter of the time cost of children in the household problem.

Age-specific LFP gaps At the opposite end of the age spectrum, a significant share of youth, ranging from 17 to 27 percent in the region, is neither employed, nor enrolled in education or training (the so-called NEETs). Although a high portion of the youngest remains in school, the share of NEETs increases around the age when most people finish high school, and this increase reflects the incipient gender participation gaps that appear around this time (Figure A2). Moreover, for those who do enter the labor force, unemployment is an important concern, as the youth unemployment rate is markedly higher than the overall unemployment rate in all the countries in the sample (Figure A5).

4 A model of demographic shifts and endogenous labor force participation

Given the microeconomic drivers of FLFP highlighted in the previous section, the large shifts in demography expected in current population forecasts are likely to be accompanied by potentially

¹⁸Many informal jobs and self-employment can be considered close substitutes in Latin America (Narita, 2020; Herreño and Ocampo, 2023)

large endogenous LFP responses. To assess the magnitude of this channel we build a simple dynamic model embedding endogenous female labor supply decisions into an overlapping generation framework with a rich demographic structure (Auerbach and Kotlikoff, 1987; Gagnon et al., 2021; Auclert et al., 2021). We then simulate the effect of shifting demographics on output growth under counterfactual scenario. Shifting demographics could induce an increasing female labor force participation which can partially offset the impact of aging on the workforce.

4.1 Environment

We consider a simple environment consisting of a small open economy with overlapping generations of households. Time is discrete and runs from $t = 0$ to ∞ , agents have perfect foresight. Households can save by investing in capital or in government bonds, which share the same global return; other variables and parameters can vary across countries c . We drop indices c unless there is risk of ambiguity.

Households At each time t , a country has a population $N_t = \sum_j N_{jt}$ growing at rate $1 + n_t \equiv N_t / N_{t-1}$, with N_{jt} being the number of individuals of age j . To reflect evolving age structures over time, we depart from the usual perpetual youth model (Yaari, 1965; Blanchard, 1985) with age-dependent mortality rates as in Auclert et al. (2021). Each individual faces an exogenous probability ϕ_j of surviving from age j to age $j+1$, so the probability of surviving from birth to age j is $\Phi_j \equiv \prod_{j'=0}^{j-1} \phi_{j'}$. The maximal lifespan is J , so that $\phi_J = 0$. We assume that this survival profile is constant over time, and that there is no migration.¹⁹ Hence, the age distribution, $\mu_{jt} \equiv \frac{N_{jt}}{N_t}$, only varies over time due to changes in fertility and the ensuing convergence dynamics as the equilibrium age distribution converges to its stationary limit.

Individuals face idiosyncratic income risk, and can partially self-insure and smooth income over their life cycle by saving in an actuarially fair annuity, which has a purchase price equal to the survival probability ϕ_j , and which pays out $(1 + r_t)$ in case of survival. One implication of this setup is that mortality does not appear in the Euler equation: the utility shifter coming from survival is exactly offset by the price of the annuity.

Households labor supply is semi-endogenous. Denote ℓ_j^g , $g \in \{f, m\}$ the effective labor supply for each working-age member of the household. We assume that each partner in the household is endowed with one unit of time. As in Olivetti et al. (2024) we assume for simplicity that the husband works full-time in the market and supply a fixed amount of time $\bar{\ell}_{jt}^m$ at wage w_m , taking as given for now. This assumption is consistent with prime-age male LFP rate being remarkably stable over the past decades and during the life-cycle (Figure 2b), and close to one. This allows us to focus on the labor supply decision of women, jointly decided by the household. We abstract from education decisions.²⁰ Men supply labor exogenously, $\bar{\ell}_j^m$, while female labor supply is endogenous.

¹⁹In practice survival rates have evolved over time generating a 'rectangularization' of the survival curves (Lleras-Muney and Moreau, 2022).

²⁰See Malta et al. (2019) for a model taking account endogenous education decisions.

This captures the larger variations of FLFP over time and during the life-cycle, while MLFP is remarkably flat

Households with birth year k choose sequences of consumption c_{jt} , labor supply ℓ_{jt}^f and annuities $a_{j+1,t+1}$ priced in an actuarially fair manner with a purchase price equal to the survival probability ϕ_{jt} , and a payout of $1 + r_t$ in case of survival, for all ages $j = j_0, \dots, J$ (with $t = j + k$) to solve the utility maximization problem

$$\begin{aligned} \max_{\{c_{jt}, \ell_{jt}^f, a_{j+1,t+1}\}} \quad & \sum_{j=j_0}^J \beta \Phi_{jt} u(c_{jt}, \ell_{jt}^f; \kappa_{jt}) \\ \text{s.t.} \quad & c_{jt} + \phi_j a_{j+1,t+1} \leq w_t(1 - \tau^l) \left[(1 - \tau^f) \ell_j^f + \bar{\ell}_j^m + T_j \right] + (1 + r_t) a_{jt} \\ & a_{j+1,t+1} \geq -\bar{a} Z_t, \end{aligned} \tag{1}$$

where $\bar{a}$ is a borrowing constraint, w_t is the real wage per efficiency unit of labor at time t , r_t is the return on assets, τ^l is the labor tax rate, τ^f a wedge on female labor (the wage markdown, reported as π in the calibration of Section 5), and T_j denotes transfers from the government, including wage-indexed social insurance and retirement transfers, for agents of age j . This stylized annuity market keeps the model tractable by simplifying the Euler equation.

The households' per-period utility depends on the consumption of market goods, c_{jt} , and a home-produced good

$$u(c_{jt}, \ell_{jt}^f; \kappa_{jt}) = \ln c_{jt} + h(\kappa_{jt}) \ln H \tag{6}$$

where κ_{jt} indexes the household's dependent-care needs at age j and $h(\kappa_{jt})$ is the weight on home production, which covers tasks such as house chores, child rearing.²¹ We assume that $h' \geq 0$ and $h'' \leq 0$, consistent with the fact that having more children require more child-rearing time but with possible returns to scale (Jones et al., 2025). Home production, H , combines unpaid household services of both partners measured in units of time

$$H = \left(1 - \ell_{jt}^f\right)^\rho \left(1 - \bar{\ell}_{jt}^m\right)^{1-\rho} \tag{7}$$

where ρ captures social norms about gender roles. A higher value of ρ would correspond to context where social norms assign a larger role to women in home production. For now, we take ρ as exogenous but see Olivetti and Petrongolo (2016) or Fernández (2013) for how to endogenize cultural norms about gender roles.

Production There is a single good used for private consumption, government consumption, and investment. Final output Y_t of this good is produced competitively from physical capital K_t and

²¹While Olsson (2025) allows for inter-temporal substitution of consumption with elasticity $\sigma \in (0, \infty)$, her calibration leads to $\sigma = 1$, that is, log utility. We follow this specification for simplicity.

effective labor input L_t according to a Cobb-Douglas aggregate production function

$$Y_t = K_t^\alpha (Z_t L_t)^{1-\alpha} \quad (8)$$

with labor-augmenting technological progress $Z_t \equiv Z_0(1 + g_Z)^t$. Capital has a law of motion $K_{t+1} = (1 - \delta)K_t + I_t$ where I_t is aggregate investment. Effective labor input is built from age-specific labor supplies: $\bar{\ell}_j^m$ and ℓ_{jt}^f denote average effective labor input per person of age j , capturing variation in experience and hours of work over the life cycle. Aggregating within gender gives $L_{mt} = \sum_j \frac{N_{jt}}{2} \bar{\ell}_j^m$ and $L_{ft} = \sum_j \frac{N_{jt}}{2} \ell_{jt}^f$, and the two combine into aggregate effective labor according to

$$L_t = L_{ft}^\xi L_{mt}^{1-\xi}, \quad \xi \in (0, 1). \quad (9)$$

The Cobb-Douglas case is the $\eta \rightarrow 1$ member of the constant-elasticity family, which we set out in full in Appendix C. We adopt it in the body for two reasons. It is the only value of the elasticity for which the participation decision admits the closed form of Section 4.2, and that closed form is what allows the female labor supply response to be decomposed analytically into a children-dependency channel, an old-age-dependency channel and a residual. And, once ξ is calibrated to the base-year female share of the labor bill, the restriction is quantitatively close to free: Appendix C shows that a twenty percent rise in female labor input raises aggregate labor by 7.50 percent under Cobb-Douglas against 7.59 percent at $\eta = 1.3$ and 7.71 percent at $\eta = 2.0$. Our subnational estimate of $\hat{\eta} = 1.3$ is discussed in Section 5; we cannot instrument it, and we cannot rule out that it varies across countries or over time, so a fully general quantification with an identified and possibly heterogeneous η is left to future work.

For simplicity we assume that the labor market for men is perfectly competitive, so that men are paid their marginal product. There is ample evidence that the market for female labor is not, and that women are paid less than men for the same observable characteristics. We therefore introduce a wage penalty π , so that women receive a fraction $1 - \pi$ of their marginal product; π is the same object as the wedge τ^f in the household budget constraint.²² Factor prices are then

$$w_{mt} = (1 - \alpha)(1 - \xi) \frac{Y_t}{L_{mt}}, \quad (10)$$

$$w_{ft} = (1 - \pi)(1 - \alpha) \xi \frac{Y_t}{L_{ft}}. \quad (11)$$

Two features of (10)–(11) matter for what follows. The female wage falls as female labor input rises, at an elasticity of one, so the participation response to falling child dependency is self-limiting: as more women enter, the wage that draws them in declines. And the share parameter

²²For a microfoundation of a wage markdown for women, [Ransom and Oaxaca \(2010\)](#) considers a model where both men and women face oligopolistic labor markets and differ in their supply elasticities. Oligopsonistic competition arises when labor demand is concentrated in a smaller share of sectors and firms ([Bhaskar et al., 2002](#)), as is the case for female workers in Latin America, and if women face higher switching costs, informational frictions, or search costs. Discrimination by firms can also reinforce gender roles in the household through self-fulfilling expectations ([Dolado et al., 2013](#)).

ξ is not free. Dividing (11) by the sum of the two wage bills, the female share of the labor bill is $(1 - \pi)\xi / [(1 - \pi)\xi + 1 - \xi]$, an observable moment. We calibrate ξ to it country by country in Section 5, rather than recovering it from the intercept of the relative-demand regression—Appendix C explains why the latter does not identify it.

Government The government purchases G_t goods, maintains a constant tax rate on labor income τ , gives individuals state-contingent transfers T_j indexed to current wages w_t , and finances itself using a risk-free bond with real interest rate r_t . It faces the flow budget constraint

$$G_t + w_t \sum_{j=0}^J \frac{N_{jt}}{2} T_{jt} + (1 + r_t)B_t = \tau^l w_t L_t + B_{t+1} \quad (12)$$

where a positive B_t denotes government borrowing. In practice, the government chooses a retirement age j_r and a replacement rate b such that $T_{jt} = b w_t$ if $j \geq j_r$ and 0 otherwise. When demographic change disturbs the balance of aggregate tax receipts and expenditures, the government adjusts G_t to ensure that the debt-to-output ratio $\frac{B}{Y}$ remains fixed.

Equilibrium With labor-augmenting growth at a constant rate g_Z and no adjustment costs in capital, let us define $k_t = \frac{K_t}{Z_t L_t}$ the amount of capital per unit of labor. The net rental rate of capital and the wage per efficiency unit of labor are

$$r_t = \alpha \left(\frac{1}{k_t} \right)^{1-\alpha} - \delta \quad (13)$$

$$w_t = Z_t (1 - \alpha) k_t^\alpha \quad (14)$$

For women to enter the labor market, their realized market wage must be higher than the reservation wage $\underline{w}^f$ corresponding to the marginal utility of time of women when all their time is allocated to home production, $\ell^f = 0$. The reservation wage is therefore

$$\underline{w}_{jt}^f = \underline{w}^f(c_{jt}, \kappa_{jt}) = \frac{\rho h(\kappa_{jt})}{\lambda_{jt}(1 - \tau^l)(1 - \tau^f)} \quad (15)$$

where $\lambda_{jt} = \beta^j \Phi_j \frac{1}{c_{jt}}$ is the Lagrange multiplier characterizing the date-zero marginal utility of one unit of revenue for households of age j and capturing the income effect. In this baseline framework, we abstracted from within-cohort heterogeneity. Different assets and income level would be reflected in different values of the Lagrange multiplier, leading some women with a higher reservation wage. More broadly, at the cohort-level, the reservation wage is higher, i) the higher the number of children κ_{jt} , ii) the stronger the gender role norm ρ for women to stay home, iii) the higher the wage penalty. When the equilibrium wage rate of women is higher than $\underline{w}_{jt}^f$, then the maximization

problem has an interior solution and the equilibrium labor supply of women of age j is

$$\ell_{jt}^f(w_t) = \left(1 - \frac{w_{jt}^f}{w_t}\right), \quad w_t \geq \underline{w}_{jt}^f \quad (16)$$

The Euler equation is

$$\frac{c_{j+1,t+1}}{c_{jt}} = \beta(1 + r_{t+1}) \quad (17)$$

where the mortality rates cancel out due to the formulation of the actuarially fair annuities.

Definition 1 (Equilibrium). *Given demographics, government policy, an initial distribution of assets, and initial levels of bonds and capital across countries such that $F_K - \delta$ is equal to r_0 in each country, an equilibrium is i) a sequence of returns $\{r_t\}$ and country-level allocations such that, in each country, ii) individuals chose labor and consumption optimally (16, 17), iii) firms optimize and input market clear iv) asset demand from individuals equals asset supply from firms and governments.*

The total wealth held by individuals of age j is the product of N_{jt} and the average wealth at age j , a_{jt} . Aggregate (private) wealth W_t is the sum across age groups:

$$W_t \equiv \sum_{j=0}^J N_{jt} a_{jt}. \quad (18)$$

4.2 A closed-form benchmark

Before turning to the transition, it is worth isolating the forces that govern female participation in a stripped-down version of the environment. Consider a static household with a working-age couple, in which the husband supplies $\bar{L}_m$ inelastically and the wife chooses L_f to maximize

$$\ln(c - \underline{c}) + \chi [\rho \ln(1 - L_f) + (1 - \rho) \ln(1 - \bar{L}_m)], \quad \text{s.t.} \quad c = w_m \bar{L}_m + w_f L_f, \quad (19)$$

where χ is the weight on home production, ρ the share of home production borne by women, and $\underline{c}$ a subsistence requirement covering dependents. Write $\underline{c} = \phi w_m \bar{L}_m$, so that ϕ is the consumption claim of dependents expressed as a fraction of the husband's full-time earnings. The interior solution is

$$L_f(w_f) = \frac{1}{1 + \chi \rho} \left(1 - \frac{\chi \rho (w_m \bar{L}_m - \underline{c})}{w_f}\right), \quad w_f \geq \underline{w}_f \equiv \chi \rho (w_m \bar{L}_m - \underline{c}). \quad (20)$$

Combining (20) with the labor demand schedule implied by (10)–(11) and solving for the fixed point gives the equilibrium female labor supply in closed form.

Proposition 1 (Equilibrium female labor supply). *Under the Cobb-Douglas aggregator (9) and the*

static household problem (19), equilibrium female labor supply is

$$L_f^* = \frac{1}{1 + \chi\rho + (1 - \phi)\chi\rho \frac{1 - \xi}{(1 - \pi)\xi}}. \quad (21)$$

Expression (21) captures the determinants of aggregate female participation, and provides a direct mapping with observables. Female labor supply is higher when the weight on home production χ is lower, when gender norms ρ assign less of that work to women, when the wage penalty π is smaller, when dependents absorb a larger share ϕ of the husband's earnings, and when women's share ξ in market production is larger. The last two deserve comment. A higher ϕ raises participation because it lowers the residual consumption the household can finance out of the husband's earnings alone—this is the income effect that makes female labor supply more responsive than male labor supply to shifts in non-labor income. A higher ξ raises participation through demand: where women are more productive at the margin, the wage that clears the market is higher.

The demographic content enters through χ and ϕ . Child-rearing time is what makes home production costly, so χ rises with the children dependency ratio; elderly dependents consume without earning, so ϕ rises with the old-age dependency ratio. We write the first mapping as $\chi_t = \chi_0(CD_t/CD_0)^\gamma$ and estimate the exponent rather than imposing it. This matters more than it looks. With γ set to one, χ_0 alone governs both the base-year level of female participation and its elasticity to dependency, and the two moments cannot be matched at once; freeing γ separates them, and Section 5 pins it from the subnational evidence. Substituting these into (21) makes equilibrium participation a function of the two dependency ratios and the four structural parameters, which is the analytic decomposition the introduction promises: falling child dependency raises participation through χ , rising old-age dependency raises it through ϕ , and whatever is left is the residual we do not attribute to demography. Section 5 calibrates the parameters and Section 5.1 reports the decomposition.

Two extensions of (21) are worth recording. First, a single woman with no partner solves (19) with $\bar{L}_m = 0$, so her reservation wage is zero and she participates at any positive wage—consistent with the pattern in Figure A3. If a share ω_H of women live in dual-adult households, aggregating over the two types replaces $(1 - \phi)$ in (21) with $(\omega_H - \phi)$, so that the secular shift in household composition is itself a source of rising participation. Second, allowing χ_i and c_i to differ across households by the number and type of dependents and by the partner's education preserves the structure of (20) household by household but no longer aggregates in closed form; the quantitative model of the next section handles this numerically.

The benchmark serves to illustrate the mechanism. In contrast, the quantitative results in Section 5 come from the overlapping-generations transition, in which the same production block (9)–(11) is embedded in a life-cycle household problem with savings, annuities and an endogenous capital stock. The two agree on the direction of every comparative static in (21); they differ in that the transition also prices the intertemporal margin, which the static benchmark cannot.

Two features of (21) matter for what follows. First, the demographic term enters through $\chi_t = \chi_0 (CD_t / CD_0)^\gamma$, so the exponent γ governs the *responsiveness* of participation to dependency independently of the *level* at which participation starts; imposing $\gamma = 1$, as a linear shifter does, forces one parameter to do both jobs. Second, normalizing χ_0 to one makes ρ interpretable as the share of child-related home time borne by women, so the expression can be inverted to recover a norm state from observed participation and demography. We do not use attitudes data anywhere in this inversion, which makes it a genuine out-of-sample check; we carry it out in Section 5.2.

4.3 The impact of shifting demographics

We first define an asymptotic Balanced Growth Equilibrium and study the long-run demographic impacts.

Proposition 2 (Asymptotic Balanced Growth Path). *An asymptotic balanced growth path is a dynamic equilibrium path where the global interest rate is constant equal to r , the capital intensity, the normalized wage levels and age-level consumption are constant and the equilibrium labor supply converges towards a constant value*

$$L_t^* \rightarrow \bar{L}^* = \sum_{j=j_0}^J \frac{N_j}{2} (\bar{\ell}_j^m + \bar{\ell}_j^f) \quad (22)$$

with steady-state female labor force participation

$$\bar{\ell}_j^f = 1 - \frac{\rho h(\bar{\kappa}_j) \tilde{c}_j}{\beta \Phi_j \tilde{w}^* (1 - \tau^l) (1 - \tau^f)} \quad (23)$$

The existence of an asymptotic balanced growth path follows from the fact that if the child dependency ratio is fixed, the model is similar to the standard OLG growth models and all the FOC conditions are satisfied, in particular:

$$k^* = \left(\frac{\alpha}{r + \delta} \right)^{\frac{1}{1-\alpha}} \quad (24)$$

$$\tilde{w}^* \equiv \frac{w_t}{Z_t} = (1 - \alpha) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}} \quad (25)$$

$$\frac{\bar{c}_{j+1}}{\bar{c}_j} = \beta(1 + r) \quad (26)$$

where $\tilde{c}_j = \frac{c_{jt}}{Z_t}$ is normalized consumption for age j on the balanced growth path. Note that working-age female labor supply is declining with age as $\frac{1-l_{jt+1}}{1-l_j} = \frac{1+r}{\phi_j} > 1$ in line with empirical evidence on life-cycle FLFP (Figure 2b). The behavior of FLFP over time is not necessarily monotonic, *even if* child dependency ratios fall monotonically after a given date. This is because the demographic structure carries some momentum.²³

²³“Echoes” can occur when the age structure is not single-peaked as is the case in China for instance

Proposition 3 (Permanent Effects of Endogenous Female LFP). *Consider a change in children dependency ratios following a decline in fertility rates. The permanent effect on labor supply is*

$$\Delta L = -\frac{L^m - L^f}{1 + \frac{L^m - L^f}{L^*}} \frac{h'(\kappa)}{h(\kappa)} \Delta \kappa \quad (27)$$

and both capital and GDP levels are higher by a similar amount.

The result follows from taking the total differential of the FOC for female labor supply (16) and simplifying

$$d\ell_j^f = (1 - \ell_j^f) \left[\frac{dw}{w} - \frac{dc_j}{c_j} - \frac{h'(\kappa_j)d\kappa_j}{h(\kappa_j)} \right] \quad (28)$$

On the balanced growth path and under the small open-economy assumption, the capital-labor ratio is fixed and $k = \left(\frac{\alpha}{r+\delta}\right)^{\frac{1}{1-\alpha}}$. As a result, the wage per a unit of efficient labor is constant equal to $\tilde{w} = (1 - \alpha)k^\alpha$ hence $d\tilde{w} = 0$.

Transitional dynamics While the mortality rates are fixed over time, the age structure continue to evolve for some time after a change in fertility. This 'momentum' effect (Keyfitz, 1971; Kim and Schoen, 1997) can generate additional growth during the transition, as the share of working age population continues to evolve.

Proposition 4 (Transitory Dividends). *During the transition, as the child dependency ratio $CD_t \rightarrow \overline{CD}$ declines, growth rates can increase driven by labor increases.*

While the level effects are permanent, the additional growth generated by declining fertility is transitory. Whether the transition is smooth depends on the initial age structure. Starting from a stable age distribution, a monotone decline in fertility produces a monotone decline in the child dependency ratio, and the labor supply response inherits that monotonicity. Oscillations—demographic echoes—require the age structure to depart from stability, as it does after a war, a baby boom, or an abrupt policy-induced fertility break (Keyfitz, 1971; Kim and Schoen, 1997). The Latin American economies in our sample are close to the stable case: fertility has fallen steadily since the 1960s without a sharp discontinuity, which is why the participation paths in Section 5.1 rise and then plateau rather than cycling.

A numerical illustration To illustrate one component of the labor increase underlying Proposition 3, consider a local partial-equilibrium experiment. Holding consumption, wages and demographic composition fixed, the female labor-supply condition implies

$$d\ell_j^f = -\left(1 - \ell_j^f\right) d\ln h(\kappa_j). \quad (29)$$

This is illustrative rather than a calibration of the full transition. In particular $d\ln h(\kappa_j) = \varepsilon_{h,\kappa} d\ln \kappa_j$, so translating a change in the child dependency ratio into a change in $h(\kappa_j)$ requires an assumption

about the elasticity $\varepsilon_{h,\kappa}$. That elasticity is the object we estimate in Section 5 and denote γ ; the range of values it can take is what generates the range of results reported there.

Suppose the decline in child dependency lowers $h(\kappa_j)$ by ten percent. Starting from $\ell_j^f = 0.6$, female labor supply rises by

$$d\ell_j^f = -(1 - 0.6)(-0.10) = 0.04,$$

or four percentage points. For a transparent aggregation, suppose working-age individuals are forty percent of the population, half of them women, male labor supply is one, and workers have equal efficiency. Effective labor per capita rises by $0.4 \times 0.5 \times 0.04 = 0.008$, against an initial level of $0.4 \times 0.5 \times (1 + 0.6) = 0.32$, so aggregate effective labor rises by

$$\frac{dL}{L} = \frac{0.008}{0.32} = 0.025,$$

or 2.5 percent. Under the balanced-growth small-open-economy assumptions used above capital adjusts with labor, so the long-run output level rises by the same proportion; spread over T years its average contribution to annual growth is $(1.025)^{1/T} - 1$.

Household heterogeneity We can introduce household heterogeneity. Their effective labor supply is $\ell(z_j)$, where z_j is a stochastic process, and unless stated otherwise, all individual variables at age j are a function of the whole history of the idiosyncratic shocks z_j , which we denote z^j . The stochastic process z_j can be different across countries, and is arbitrary subject to being the same across cohorts, meaning that over time, individuals of the same age have identical distributions across histories regardless of their birth year.

5 Quantification

This section explains how we take our baseline model to the data. We rely both on micro and macro moments to set the parameters $\{\alpha, \chi, \gamma, \eta, \pi, \rho, \phi, \bar{L}_m\}$ that characterize our baseline model.

The elasticity of substitution between men and women in production is identified from the slope of the log ratio of female to male labor force participation against the log inverse wage ratio. More precisely, we run the following regression

$$\ln\left(\frac{LFPR_{f,i}}{LFPR_{m,i}}\right) = a_0 + a_1 \ln\left(\frac{w_{m,i}}{w_{f,i}}\right) \quad (30)$$

and we therefore have $\eta = \hat{a}_1$. The regression is shown in Figure A9 and gives $\hat{\eta} = 1.28$, with a standard error of 0.14 under the classical formula and 0.10 clustering on administrative division, on 196 observations across 57 divisions; we adopt $\eta = 1.3$. We use the regression for its slope alone. Its intercept does not identify the share parameter ξ , because the left-hand side is the ratio of participation *rates* rather than of effective labor inputs, so the female-to-male ratio of working-age population and any systematic difference in hours per participant load onto it; Appendix C sets this

out and reports the intercept-based value, $\xi = 0.12$, as a rejected diagnostic. We calibrate ξ_c country by country to the female share of the labor bill through equation (C.4). The point estimate lies above one, so female and male labor are substitutable rather than complementary, but it lies close to one: the clustered interval excludes unity while the classical interval, $[0.99, 1.56]$, does not. We read this as licensing the Cobb-Douglas aggregator of the baseline rather than as rejecting it, and Appendix C carries the general CES case. The elasticity we find for the region is in any case smaller than the one estimated on US data.²⁴ We set the capital share α_c country by country to match each economy’s capital-output ratio in the Penn World Table, which yields values between 0.28 and 0.42 rather than a single common value. The parameters capturing the interaction between the demographic structure and the female LFP deserve more attention. An increasing old age dependency ratio raises ϕ whereas the weight on home production increases with the children dependency ratio. We adopt $\phi = \underline{\phi} \times \text{Old-age dependency ratio}$ and $\chi_t = \chi_0 (CD_t / CD_0)^\gamma$, and we calibrate γ rather than setting it to one. In the runs reported below the old-age loading is switched off ($\underline{\phi} = 0$); the significant old-age coefficient in column (2) of Appendix Table A2 suggests that quantifying this channel is a natural next step. Differentiating (21) at the base year gives $\left| d\ln L_f^* / d\ln CD \right| = \gamma B_c / (1 + B_c)$, where B_c is the composite in the denominator of (21), so γ is identified by the elasticity of participation to child dependency once B_c is anchored on the base-year participation level. That elasticity comes from the subnational regression of the female-to-male labor force ratio on child dependency shown in Figure 5, estimated with and without census-year fixed effects and converted to an elasticity at the sample-mean dependency ratio: 0.26 with year effects and 1.06 without. Appendix Table A2 reports the pooled level specification; the census-year-effects estimates and their standard errors are available from the authors, pending the country-level re-estimation in logs that reports the elasticity directly.

Parameter	Value	Data and target
α_c	0.28–0.42	Capital share, calibrated to K/Y (PWT)
χ_0	1	Normalization; gender norms absorbs the level
γ	0.45 / 1.85	Elasticity of home-production weight to child dependency (Sec. 5)
ξ_c	0.37–0.48	Female share of the labor bill, via eq. (C.4)
π	0.25	Female wage penalty, $w_f = (1 - \pi)MPL_f$; Mincerian regression
ρ_c	0.40–0.48	Share of child-related home time borne by women, eq. (5.2)
$\bar{L}_m$	0.92	Male LFPR

Table 2: Summary of model parameters

Note: γ is reported at the two estimated values, with and without census-year fixed effects; the solves of Table 3 hold the composite $\chi\rho = 0.28$ fixed while varying γ . Only the product $\chi_0\rho_c$ is identified. We report it on the normalization of Section 5.2, $\chi_t = \chi_0 CD_t^\gamma$ with $\chi_0 = 1$, so that ρ_c here is the same object recovered by (5.2) and plotted in Figure 7. The solves of Table 3 impose $\rho_c CD_c(2016)^\gamma = 0.28$ uniformly across countries; the country range above is what that single restriction implies given each country’s base-year dependency, not calibrated heterogeneity.

Several parameter choices warrant additional discussion. First, ξ_c is calibrated country by

²⁴Ngai and Petrongolo (2017b) finds an elasticity of 2.27 and Acemoglu et al. (2004) more than 3.




Figure 5: Empirical relationship between dependency ratios and female labor force participation

Sources: IPUMS International ([Ruggles et al., 2024](#)); and IMF staff calculations. Notes: the figure presents a binned scatter plot of the female-to-male labor force ratio against the child dependency ratio (population aged 0–15 relative to the working-age population, 15–64). The specification includes fixed effects for first-level administrative divisions, a first-order polynomial fit, and 95 percent confidence intervals. Further detail is in Appendix Table [A2](#).

country to the female share of the labor bill through equation (C.4), which yields values between 0.37 and 0.48. For Brazil, Mexico and Colombia, which together account for roughly two-thirds of the LA7 population, we can check this against the share computed directly from household-survey microdata rather than from published aggregates. The two routes agree to within 0.02 in every case, a difference of well under half a percentage point in the 2050 results. For the remaining four countries the share is taken from published aggregates and an assumed ratio of male to female hours, and we flag it as the calibration input we would most like to replace with a direct measurement. Because women are paid less than their marginal product, ξ_c exceeds the observed earnings share: the penalty drives a wedge between what female labor is paid and what it contributes at the margin, which is the sense in which the gender wage gap causes the national accounts to understate female labor’s marginal contribution. Recovering ξ instead from the intercept of the relative-demand regression delivers 0.12, implying a female share of the labor bill near 9 percent against roughly a third in the data; Appendix C explains why that route is unavailable, and we report the value only as a diagnostic. Second, π is the female wage penalty, defined so that women are paid $(1 - \pi)$ times their marginal product; it is the same object as the wedge τ^f in the household budget constraint. The Mincerian coefficient of -0.29 log points (Appendix Table A1) implies a retention factor of $e^{-0.29} = 0.75$, so we set $\pi = 0.25$. We treat π as a reduced-form catch-all for the several sources of the gender wage gap—human capital differences, occupational sorting, discrimination, and monopsony power—rather than attributing it to any one mechanism. Third, the elasticity of substitution $\eta = 1.3$ is lower than estimates for the United States (Ngai and Petrongolo, 2017b; Acemoglu et al., 2004). This likely reflects greater labor market segmentation and higher informality in Latin America, which compresses the responsiveness of female relative labor supply to wage differentials. A lower η implies a larger wage response to gap-closing policies, which we discuss in the sensitivity analysis.

Remember that the growth of the labor force can be expressed as:

$$\Delta l_{ct} = \Delta N_{ct} + \Delta WAP_{ct} + \Delta LFPR_{ct} + \Delta ER_{ct} \quad (31)$$

where N_{ct} is total population, WAP_{ct} is the share of working-age population (those 15 and older), $LFPR_{ct}$ is the labor force participation rate, and ER is the employment rate. In turn, the change in $LFPR_{ct}^g$ for both genders $g \in \{f, m\}$ is linked to the micro parameters that we have estimated in the previous section.

$$LFPR_{ct}^g = \Phi(X_{ct}^g, \beta_c^g) + \varepsilon_{ct} \quad (32)$$

where X_{ct} is a set of population characteristics (age, education, and so on), and β_c^g connects these characteristics with the propensity to join the labor force. Note that we assume the LFPR is separable in gender, that is, there is no crowding out of male LFP when more women join the labor force. This assumption of gender separability is broadly consistent with the data for Latin America, where we find little evidence of substitution between male and female participation across

subnational units. We then construct scenarios where we use population projections for X_{ct} and examine the impacts of closing or shrinking various gaps $\hat{\beta}_c^m - \hat{\beta}_c^f$ in the propensity to work.

We bring the model to the data in two complementary ways, and organize the results around the model itself. We first ask how much of the gender participation gap closes *endogenously* as demographics evolve (Section 5.1), then feed that response through the model’s full general-equilibrium transition and decompose its impact on output in a manner consistent with the production side of the model (Section 5.3). A set of purely mechanical growth-accounting benchmarks—which hold capital and productivity growth at their 2010–19 averages and close the gender gap by assumption—serves as an external check and is reported in Section 5.4. Throughout, the economies are treated as small open economies with frictionless access to world capital; introducing frictions in capital flows is the natural next refinement.

5.1 The demographic response of female participation

Female labor supply is now a model object. At each date the closed form of Section 4.2 returns equilibrium female participation given the child dependency ratio along the UN median-fertility path, the production share ξ_c , the wage penalty, and gender norms; effective labor and factor prices follow; and the transition is solved as in Section 5. Nothing is interpolated toward an assumed benchmark, and the participation path is no longer an input to the exercise.

What the exercise delivers is a range, and the width of the range is the finding.

Table 3 reports the 2050 GDP level effect relative to a path holding female participation at its 2016 profile. Under the specification that includes census-year fixed effects, falling child dependency raises regional GDP by 3.1 percent by 2050. Under the specification that omits them, the same mechanism delivers 11.3 percent. Between those two figures lies a factor of nearly four, and it is generated entirely by a single parameter: γ , the elasticity of home-production time to the child dependency ratio. Figure 6a plots the mapping. The relationship is close to linear over the relevant range, so the headline is not fragile in any subtle way. It is simply proportional to an elasticity we cannot pin down. One number is worth pausing on: at the elasticity implied without year effects the model returns 11.3 percent, against the 11.6 percent reported by the interpolation-based exercise this section replaces. The earlier headline is, almost exactly, what one obtains by attributing the entire national trend to demography.




Figure 6: Impact of alternative dependency elasticities

Sources: 2024 UN Population Prospects; Penn World Table 10.0; ILOSTAT; and authors' calculations (Table 3). Notes: Panel (a): the solid line traces the population-weighted LA7 GDP level effect in 2050, relative to a path with female participation held at its 2016 profile, as the elasticity γ varies. Markers indicate the three values discussed in the text: the elasticity implied by Appendix Table A2 estimated with census-year fixed effects, the proportional mapping $\gamma = 1$, and the elasticity implied by the same regression without year effects. The shaded region spans the two estimated values. The dashed horizontal line is the figure reported by the interpolation-based exercise this specification replaces; note that it is almost exactly what the model delivers at the elasticity estimated without year effects. Panel (b): the solid line is the preferred specification, $\gamma = 0.45$; the dashed and dotted lines are $\gamma = 1.85$ and $\gamma = 1.00$; the effect accumulates fastest between 2025 and 2040, when the child dependency ratio falls most steeply, and flattens thereafter as dependency approaches its floor.

The interpretation of the two ends is not symmetric, and the paper should not pretend otherwise. The upper figure comes from a regression in which child dependency is measured against a national mean that is itself falling. Over the estimation window, female education rose, employment shifted out of agriculture, contraception spread, and household norms changed, all of them moving in step with the fertility decline. A specification without year effects assigns that entire common movement to demography. Adding year effects identifies the parameter from cross-state deviations within a census round, which removes the national trend and roughly three-quarters of the identifying variation with it. The lower figure is what remains once the confound is purged.

Table 3: GDP level effect of the demographic participation response in 2050, by the assumed elasticity of home-production time to child dependency

Country	GDP level effect in 2050 (percent)		
	$\gamma = 0.45$	$\gamma = 1.00$	$\gamma = 1.85$
	census-year FE	imposed	no year FE
Argentina	3.9	8.1	13.6
Brazil	2.6	5.5	9.6
Chile	3.5	7.4	12.5
Colombia	3.6	7.7	12.9
Mexico	3.5	7.3	12.4
Peru	3.5	7.3	12.4
Uruguay	3.4	7.3	12.3
LA7 (pop. weighted)	3.1	6.6	11.3

Sources: 2024 UN Population Prospects (median fertility scenario); Penn World Table 10.0; ILOSTAT; and authors' calculations. Notes: the level effect is the percent difference in GDP between the transition with the model's endogenous female participation response and one holding female participation at its 2016 profile, cumulated from 2016. The parameter γ is the elasticity of the home-production weight to the child dependency ratio, $\chi_t = \chi_0(CD_t/CD_0)^\gamma$. The outer columns are pinned by the subnational regression of Appendix Table A2, estimated with and without census-year fixed effects. The composite $\chi\rho_g$ is held at 0.28 in all three columns, so the only quantity varying across them is γ . Figures are from the general-equilibrium solves; the closed form of Section 4.2 reproduces them to within 0.05 percentage points. The share parameter ξ_c is provisional; see Section 5.

We therefore take 3.1 percent as our estimate of the demographic channel and report 11.3 percent as the upper bound that a cross-sectional design would deliver. The gap between them is not model uncertainty in the usual sense. It is a quantity with an economic name: the contribution of the secular non-demographic trend in female participation, which no design available to us can separate from the demographic channel using the variation we have.

Two features of the country pattern survive across the range. The effect is smallest in Brazil, at 2.6 percent in the preferred specification, because it begins the projection period with child dependency already low and therefore has the least demographic distance left to travel. And the ordering across countries is stable in γ , so the ranking is a robust statement even where the level is not.

The timing is the same across the range even though the level is not (Figure 6b). The effect accumulates fastest between the mid-2020s and 2040, which is when child dependency falls most steeply, and flattens after that as dependency approaches its floor. By 2040 the preferred specification has delivered 2.8 of its eventual 3.1 percent, so nearly the whole of the demographic contribution arrives within two decades. This matters for how the result should be read against the aging literature. The participation response is front-loaded and the working-age drag is not, so the two do not offset each other evenly through time even when their cumulative magnitudes are comparable.

5.2 The role of gender norms

The exercise above holds fixed gender norms, a variable notoriously difficult to pin down. Below we propose an approach to recover estimates from observables instead of attitudes data. Normalizing χ_0 to one makes ρ the share of child-related home time borne by women, and setting the old-age loading to zero, (21) then solves for gender norms directly:

$$\rho_c(t) = \frac{1/L_c^f(t) - 1}{S_c CD_c(t)^\gamma}, \quad S_c \equiv 1 + \frac{1 - \xi_c}{(1 - \pi)\xi_c}, \quad (33)$$

where $L_c^f(t)$ is the observed female participation rate, $CD_c(t)$ the child dependency ratio, and S_c the composite of the production share and the wage penalty. Conditional on the calibrated parameters, equation is an accounting identity: it reports the value of the home-care allocation wedge required for the model to reproduce observed female participation. Our gender-norm parameter is a residual and, therefore, may also absorb other persistent determinants of female participation not represented explicitly in the model. We evaluate the inversion on the population aged 15–64. The denominator of $CD_c(t)$ is the working-age population, so a participation rate defined instead on the population aged 15 and over would place two different populations on the two sides of one expression, which turns out to matter in practice. Our inversion differs from occupation-based measures of gender barriers. Li et al. (2025) use a Roy model to infer occupation-specific wedges from gender wage gaps adjusted for the endogenous selection of women and men across occupations. Their identification is relative to a reference economy and relies on cross-country invariance in relative gender differences in innate occupational aptitude. Our exercise instead recovers a single model-implied parameter from aggregate female participation after accounting for child dependency and calibrated production conditions. It therefore captures the participation margin associated with the gendered allocation of home time, but does not identify barriers to occupational entry or the productivity gains from reallocating talent across market occupations.

Applied to advanced economies at 2019, the inversion places Nordic countries and the Netherlands with the more egalitarian norms, followed by Canada, the United Kingdom, Germany and Japan, and then Southern Europe (Figure 7). The ordering is well-aligned with the comparative literature on gender norms drawing from survey and time-use evidence (Fernández, 2013; Olivetti and Petrongolo, 2016). Two cases sit awkwardly with that reading: the United States ranks thirteenth of sixteen and Japan only ninth. We did not target this moment and we do not claim it is reproduced exactly. That most of it follows from participation and demography alone is our main evidence that ρ captures social norms, at least to some extent.




Figure 7: Implied gender norm parameter, Advanced Economies and the LA7 sample

Sources: ILOSTAT; World Development Indicators; 2024 UN Population Prospects; and authors' calculations. Notes: gender norms ρ is recovered by inverting equation (5.2) at each country's observed female participation rate and child dependency ratio, with χ_0 normalized to one and $\gamma = 0.45$. Participation is measured on the population aged 15–64, the same population that forms the denominator of the child dependency ratio; the ILO modelled series for that age band is used throughout. Advanced economies are evaluated at 2019; the LA7 countries, in bold, at the 2020–24 average. A higher value denotes a larger share of home production borne by women. The advanced-economy ordering is an untargted moment. Measuring participation on the population aged 15 and over instead raises every value, by more where participation at ages 65 and over is high, and is reported in Section 5.2.

Our own sample can then be read against that benchmark. Peru returns 0.24, Uruguay 0.30, Argentina 0.40, Brazil 0.43, Chile and Mexico 0.48, and Colombia 0.48. Six of the seven lie above the advanced-economy median of 0.29; the regional median of 0.43 is about one and a half times it; and of the sixteen advanced economies only Italy lies above the highest value in our sample. The seven span $[0.24, 0.48]$, an interval contained within the advanced-economy $[0.17, 0.64]$. Latin America is more norm-constrained than the comparison group by a moderate margin, measured on a common population and without attitudes data anywhere in the calculation. Both the direction and the order of magnitude are what one would expect, which is a further check on the inversion and not a result we impose.

Peru's estimate present an interesting anomaly. Having the highest female participation rate among LA7, it naturally returns the lowest implied norm state ($\rho = 0.24$). The discrepancy in FLFP between Peru and its neighbors is concentrated among women above 50 and rural areas. In particular, female participation at ages 65 and over reaches 35.7 percent in Peru, compared with a sample average of 15.2 percent and 8.8 percent in Brazil. This is consistent with Peru's weaker

pension coverage. while Brazil’s rural pension is close to universal and its rate at those ages is 8.8 percent.

Gender norms are held fixed in the counterfactuals, not because the two cross-sections cannot be placed on one scale, but because we have no evidence on the speed at which such a state moves. Appendix F supplies the only direct evidence available: across Brazilian states over 2012–2025 gender norms do not converge, and their dispersion widens. Projecting convergence would be an assumption, and we instead choose to provide a plausible range.

Two consequences follow for what the paper claims. Our range is a range over the demographic channel with norms held constant, so it is a statement about demography and not a forecast of participation. And the historical decomposition that would separate the demographic contribution from the norm contribution requires participation data reaching back before 1990, since the region’s norm transition was already well advanced by the time our series begins. We return to this in the conclusion.

One channel these counterfactuals hold fixed deserves comment, because its sign is not obvious. Higher female participation is generally expected to reduce fertility (Galor and Weil, 1996; Doepke and Tertilt, 2009), which would lower child dependency further and reinforce the labor supply effect. Within this framework the opposite is also possible: with male and female labor imperfect substitutes, a rise in female labor supply lowers the female wage relative to the male wage, and it is the relative wage that governs the opportunity cost of children. We therefore hold the UN median projections fixed in every scenario and treat the direction of the fertility feedback as an open question. Quantifying it requires closing the loop between participation and fertility, which we leave to future work.

5.3 A model-consistent decomposition

We solve the general-equilibrium transition to gauge the extent to which the rise in FLFP due to falling child dependence ratios can offset the declining share of working-age adults in LA7. We solve the model twice: once along a demographic-only baseline with female participation frozen at its 2016 profile, and once allowing the response of Section 5.1. Because capital and the interest rate adjust each time we solve the model, differencing the two output paths captures both the direct contribution of the additional labor and the capital deepening it induces. The decomposition uses each country’s own calibrated capital share α_c , pinned by its capital-output ratio in the Penn World Table.

Figure 8 shows the anatomy of growth country by country. The demographic dividend that powered past growth is fading on two margins at once: population growth slows, and the working-age share turns from tailwind to headwind. Against that, two forces push the other way, the participation response and the capital it pulls in. Under the preferred elasticity the participation channel offsets roughly a third of the working-age drag through the late 2030s and less thereafter. Under the upper elasticity it covers the drag entirely until the early 2040s. Across the whole range the offset is real and temporary; only its size is in question.



Figure 8: Decomposition of projected GDP growth under the endogenous participation path

Sources: 2024 UN Population Prospects (median fertility scenario); Penn World Table 10.0; and authors' calculations. Notes: bars decompose projected annual GDP growth, in percentage points, into the contribution of capital ($\alpha_c \Delta \ln K$) and four components of effective labor, each scaled by $1 - \alpha_c$: population growth, the changing working-age share, labor quality, and the female participation response. Markers trace total GDP growth. Figures are five-year averages; the economies are treated as small open economies. The participation channel is identified as $(1 - \alpha_c) [\Delta \ln(L/WAP)_{FLFP} - \Delta \ln(L/WAP)_{base}]$: because the demographic-only baseline holds female participation at its 2016 profile, movement in labor per working-age person along that path is pure age composition, and the difference between the two solves is the participation response. The five contributions sum to $\Delta \ln Y$ to within 10^{-12} percentage points. Labor-augmenting productivity is exogenous and common across scenarios so its contribution is identically zero and is omitted. The LA7 aggregate is population-weighted. The panel is drawn at $\gamma = 1.00$.

5.4 Exogenous gap-closing benchmarks

As an external check on the model’s response, we also compute purely mechanical growth-accounting counterfactuals that hold capital and productivity growth at their 2010–19 averages and vary only the labor force. Holding the participation and employment terms of the labor identity at their 2010–19 levels and updating population and the working-age share for the UN median projections, the demographic-only path lowers average annual GDP growth to about 2.5 percent from the roughly 2.8 percent observed before the pandemic. We then raise female participation, holding male participation fixed, to levels consistent with gender gaps of 15, 10, and 0 percentage points.

Closing the participation gap to advanced-economy levels would offset the demographic drag in most of the sample. The model’s demographic response under the preferred elasticity reaches roughly a quarter of that benchmark by 2050 and, under the upper elasticity, roughly three-quarters of it. Demographic change on its own therefore does let Latin American economies’ FLFP reach the levels currently enjoyed by advanced-economy.

6 Conclusion

This paper develops a unified framework for studying the interaction between demographic change and female labor supply in Latin America. We combined harmonized microdata from seven of the region’s largest economies with a tractable overlapping generations model in which female participation responds endogenously to household composition through a closed-form expression for the female-to-male labor ratio, and solved the model’s perfect-foresight transition under UN median demographic projections. The framework yields a mapping from observed cross-country differences in household composition, fertility, and aging to aggregate female labor force participation.

Three findings emerge. First, empirical analysis which uses harmonized cross-country micro data uncovers substantial heterogeneity across the the 7 Latin American countries in the marriage premium, the motherhood penalty, and the education gradient, which we link to differences in fertility and household structure. Second, the structural model maps these micro estimates into an aggregate labor supply equation. Model simulations suggest that the FLFP response to projected falling child dependency ratio could raise LA7 output in 20250 by between 3.1 and 11.3 percent relative to a path with participation held fixed. The two estimates reflect differences in assumptions about the relative contribution of social norms (or secular trends) and demographics . Third, the width of that interval is informative of a fundamental identification challenge arising from the co-movement between falling child dependency and the secular rise in female participation, which a design resting on cross-sectional variation in dependency cannot decompose without a structural model.

Our findings show that the participation response is a substantial but transitory force during the later part of the demographic transition. Under our baseline elasticity calibration, this response offsets roughly a third of the drag from the contracting working-age share through the late 2030s and less thereafter; under the upper elasticity it covers that drag entirely into the early 2040s. Both

readings are consistent with the findings of the long-run growth literature (Galor and Weil, 1996; Doepke and Tertilt, 2009; Bloom et al., 2009) which has emphasized the trade-offs between female labor supply and fertility. Our contribution is quantifying the role of the demographic transition in a region that has received less attention in the existing literature and where the demographic transition is set to occur while participation gender gaps are significantly higher than in advanced economies

Four extensions could help narrow the range of the predicted trajectories. The first is data rather than method: our estimation window begins in 1990, by which point Latin America’s norm transition was already well advanced, so the secular trend we cannot separate from demography is observed only in its tail. Earlier census rounds would place the trend against a period in which participation and dependency did not move together, and would in addition permit the demographic elasticity to be estimated country by country rather than pooled.

The second is the role of informality. We hold the norm parameter fixed in every counterfactual, and we do so because it cannot presently be compared across our sample and the advanced-economy benchmark: measured participation in the region includes substantial informal, unpaid family and subsistence work in a proportion that varies across countries. An informality correction disciplined by measured informality rates would let gender norms carry part of the projection instead of being held constant.

The third is the direction of the fertility feedback. Higher participation is conventionally expected to lower fertility and so to reinforce the labor supply effect. Within a framework in which male and female labor are imperfect substitutes the opposite is also available, because a rise in female labor supply lowers the female wage relative to the male wage and it is the relative wage that governs the opportunity cost of children (Galor and Weil, 1996).

The fourth is the capital side. With a small open economy facing a fixed world interest rate and no adjustment costs, the split of the effect between labor and the capital it induces is an identity rather than a finding, which is why Section 5.3 declines to report it. Adjustment costs, or an upward-sloping supply of external funds, would make that division informative. Three further refinements are worth naming without developing: the old-age dependency channel, which the model carries through ϕ and the reported runs hold at zero; the labor aggregate, which conflates formal and informal employment and which a three-way Roy model over home production, informal and formal work would separate; and the elasticity of substitution between male and female labor is estimated from subnational wage variation, which a shift-share instrument built from national sectoral shocks and baseline subnational composition could sharpen.

References

- ACEMOGLU, D., D. H. AUTOR, AND D. LYLE (2004): “Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury,” *Journal of Political Economy*, 112, 497–551.
- ACEMOGLU, D. AND P. RESTREPO (2017): “Secular stagnation? The effect of aging on economic growth in the age of automation,” *American Economic Review*, 107, 174–179.
- (2022): “Demographics and automation,” *The Review of Economic Studies*, 89, 1–44.
- AKSOY, Y., H. S. BASSO, R. P. SMITH, AND T. GRASL (2019): “Demographic structure and macroeconomic trends,” *American Economic Journal: Macroeconomics*, 11, 193–222.
- ALLAIS, M. (1947): *Économie et Intérêt*, vol. 2, Paris: Imprimerie Nationale et Librairie des Publications Officielles, two volumes.
- AMUEDO-DORANTES, C. AND S. POZO (2006): “Migration, Remittances, and Male and Female Employment Patterns,” *American Economic Review*, 96, 222–226.
- AUCLERT, A., H. MALMBERG, F. MARTENET, AND M. ROGNLIE (2021): “Demographics, Wealth, and Global Imbalances in the Twenty-First Century,” Working Paper 29161, National Bureau of Economic Research.
- AUERBACH, A. J. AND L. J. KOTLIKOFF (1987): *Dynamic Fiscal Policy*, Cambridge: Cambridge University Press, april 1987.
- BATOG, C., E. CRIVELLI, A. ILYINA, Z. JAKAB, J. LEE, A. MUSAYEV, I. PETROVA, A. SCOTT, A. SHABUNINA, A. TUDYKA, X. C. XU, AND R. ZHANG (2019): *Demographic Headwinds in Central and Eastern Europe*, no. 19/12 in IMF Departmental Paper, Washington, DC: International Monetary Fund.
- BENGALI, L., E. A. DUZHAK, AND C. ZHAO (2023): “Men’s Falling Labor Force Participation across Generations,” *FRBSF Economic Letter*, 2023, 1–6.
- BERLINSKI, S., S. GALIANI, AND P. J. MC EWAN (2011): “Preschool and maternal labor market outcomes: Evidence from a regression discontinuity design,” *Economic Development and Cultural Change*, 59, 313–344.
- BERNIELL, I., L. BERNIELL, D. DE LA MATA, M. EDO, AND M. MARCHIONNI (2021): “Gender gaps in labor informality: The motherhood effect,” *Journal of Development Economics*, 150, 102599.
- (2023): “Motherhood and flexible jobs: Evidence from Latin American countries,” *World Development*, 167, 106225.

-
- BERNIELL, I., R. FERNÁNDEZ, AND S. KRUTIKOVA (2024): “Gender inequality in Latin America and the Caribbean,” Tech. Rep. Working Paper 32104, National Bureau of Economic Research, Cambridge, MA.
- BHASKAR, V., A. MANNING, AND T. TO (2002): “Oligopsony and monopsonistic competition in labor markets,” *Journal of Economic Perspectives*, 16, 155–174.
- BLANCHARD, O. J. (1985): “Debt, deficits, and finite horizons,” *Journal of political economy*, 93, 223–247.
- BLOOM, D. E., D. CANNING, AND G. FINK (2010): “Implications of population aging for economic growth,” *Oxford Review of Economic Policy*, 26, 583–612.
- BLOOM, D. E., D. CANNING, G. FINK, AND ET AL. (2009): “Fertility, female labor force participation, and the demographic dividend,” *Journal of Economic Growth*, 14, 79–101.
- BLOOM, D. E. AND J. G. WILLIAMSON (1998): “Demographic transitions and economic miracles in emerging Asia,” *The World Bank Economic Review*, 12, 419–455.
- BURNETTE, J. (2025): “The Problems with Labor Force Participation as a Measure of Women’s Work,” *Capitalism A Journal of History and Economics*, 5, 234–265.
- CASAS, C. AND F. MOREAU (2024): “Labor Market Participation in Latin America,” in *Western Hemisphere Regional Economic Outlook*, International Monetary Fund.
- CAVALCANTI, T. V. D. V. AND J. TAVARES (2016): “The influence of gender roles on economic development: Evidence from wages,” *Journal of Economic Growth*, 21, 1–34.
- CRAVINO, J., A. LEVCHENKO, AND M. ROJAS (2022): “Population aging and structural transformation,” *American Economic Journal: Macroeconomics*, 14, 479–498.
- CUBERES, D. AND M. TEIGNIER (2016): “Aggregate effects of gender gaps in the labor market: A quantitative estimate,” *Journal of Human Capital*, 10, 1–32.
- DE LA ESCOSURA, L. P. (2007): “When did Latin America fall behind?” in *The decline of Latin American economies: Growth, institutions, and crises*, University of Chicago Press, 15–58.
- DOEPKE, M. AND M. TERTILT (2009): “Women’s liberation: What’s in it for men?” *The Quarterly Journal of Economics*, 124, 1541–1591.
- DOLADO, J. J., C. GARCÍA-PENALOSA, AND S. DE LA RICA (2013): “On Gender Gaps and Self-Fulfilling Expectations: Alternative Implications of Paid-For Training,” *Economic Inquiry*, 51, 1829–1848.
- EDJIGU, H., L. KOLOVICH, AND M. NEWIAK (2026): “Which Reforms Travel? Identifying Reform Packages to Boost Female Labor Force Participation,” IMF Working Paper 26/XX, International Monetary Fund, Washington, DC, forthcoming.
-

-
- FERNÁNDEZ, R. (2013): “Cultural change as learning: The evolution of female labor force participation over a century,” *American Economic Review*, 103, 472–500.
- FERNÁNDEZ-VILLAYERDE, J., G. VENTURA, AND W. YAO (2023): “The Wealth of Working Nations,” Tech. rep., National Bureau of Economic Research.
- GAGNON, E., B. K. JOHANNSSEN, AND D. LÓPEZ-SALIDO (2021): “Understanding the New Normal The Role of Demographics,” *IMF Economic Review*, march 2021.
- GALOR, O. AND D. N. WEIL (1996): “The Gender Gap, Fertility, and Growth,” *The American Economic Review*, 86, 374–387.
- GOLDIN, C. (2006): “The quiet revolution that transformed women’s employment, education, and family,” *American economic review*, 96, 1–21.
- (2014): “A grand gender convergence: Its last chapter,” *American economic review*, 104, 1091–1119.
- GRIGOLI, F., Z. KOCZAN, AND P. TOPALOVA (2020): “Automation and labor force participation in advanced economies: Macro and micro evidence,” *European Economic Review*, 126, 103443.
- HANSON, G. H. (2007): “Emigration, Remittances and Labor Force Participation in Mexico,” INTAL Working Paper 1456, Inter-American Development Bank, INTAL.
- HERREÑO, J. AND S. OCAMPO (2023): “The macroeconomic consequences of subsistence self-employment,” *Journal of Monetary Economics*, 136, 91–106.
- HOPENHAYN, H., J. NEIRA, AND R. SINGHANIA (2022): “From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share,” *Econometrica*, 90, 1879–1914.
- INTERNATIONAL MONETARY FUND (2021): “Recessions and Recoveries in Labor Markets: Patterns, Policies, and Responses to the COVID-19 Shock,” in *World Economic Outlook: Managing Divergent Recoveries*, Washington, DC.
- (2023): “Mexico: Selected Issues,” Tech. Rep. IMF Country Report No. 23/357, Washington, DC.
- JAIMOVICH, N. AND H. E. SIU (2009): “The young, the old, and the restless: Demographics and business cycle volatility,” *American Economic Review*, 99, 804–826.
- JAYACHANDRAN, S. (2020): “Social norms as a barrier to women’s employment in developing countries,” Tech. rep., National Bureau of Economic Research.
- JONES, C. I. (2022): “The end of economic growth? Unintended consequences of a declining population,” *American Economic Review*, 112, 3489–3527.
-

-
- JONES, T. R., E. KARGER, AND V. MICHELMAN (2025): “Cheaper by the Dozen Parental Time Use and Returns to Scale in Raising Children,” Draft first posted on March 20, 2025.
- KARAHAN, F., B. PUGSLEY, AND A. ŞAHİN (2024): “Demographic Origins of the start-up Deficit,” *American Economic Review*, 114, 1986–2023.
- KEYFITZ, N. (1971): “On the Momentum of Population Growth,” *Demography*, 8, 71–80, 10 pages.
- KIM, Y. J. AND R. SCHOEN (1997): “Population Momentum Expresses Population Aging,” *Demography*, 34, 421–427.
- KLEVEN, H., C. LANDAIS, AND G. LEITE-MARIANTE (2023): “The child penalty atlas,” Tech. Rep. Working Paper 31649, National Bureau of Economic Research, Cambridge, MA.
- KOTSCHY, R. AND D. E. BLOOM (2023): “Population aging and economic growth: From demographic dividend to demographic drag?” .
- KOTSCHY, R., D. E. BLOOM, AND A. J. SCOTT (2024): “On the Limits of Chronological Age,” .
- LEE, R., A. MASON, N. NETWORK, R. LEE, A. MASON, E. AMPORFU, C.-B. AN, L. R. BIXBY, J. BRAVO, M. BUCHELI, ET AL. (2014): “Is low fertility really a problem? Population aging, dependency, and consumption,” *Science*, 346, 229–234.
- LEWIS, W. A. (1954): “Economic Development with Unlimited Supplies of Labour,” *The Manchester School*, 22, 139–191.
- LEYVA, G. AND C. URRUTIA (2020): “Informality, labor regulation, and the business cycle,” *Journal of International Economics*, 126, 103340.
- LI, N., L. ZHANG, AND R. ZYMEK (2025): “The Allocation of Talent and Income Differences Across Countries,” Tech. Rep. IMF Working Paper 2025, International Monetary Fund, Washington, DC.
- LLERAS-MUNEY, A. AND F. MOREAU (2022): “A unified model of cohort mortality,” *Demography*, 59, 2109–2134.
- MAESTAS, N., K. J. MULLEN, AND D. POWELL (2023): “The effect of population aging on economic growth, the labor force, and productivity,” *American Economic Journal: Macroeconomics*, 15, 306–332.
- MALTA, V., A. MARTINEZ, AND M. M. M. TAVARES (2019): *A Quantitative Analysis of Female Employment in Senegal*, International Monetary Fund.
- MARTÍNEZ, C. AND M. PERTICARÁ (2017): “Childcare effects on maternal employment: Evidence from Chile,” *Journal of Development Economics*, 126, 127–137.
- NARITA, R. (2020): “Self-employment in developing countries: A search-equilibrium approach,” *Review of Economic Dynamics*, 35, 1–34.
-

-
- NGAI, L. R., C. OLIVETTI, AND B. PETRONGOLO (2024): “Gendered Change: 150 Years of Transformation in US Hours,” NBER Working Paper 32475, National Bureau of Economic Research.
- NGAI, L. R. AND B. PETRONGOLO (2017a): “Gender Gaps and the Rise of the Service Economy,” *American Economic Journal: Macroeconomics*, 9, 1–44.
- (2017b): “Gender gaps and the rise of the service economy,” *American Economic Journal: Macroeconomics*, 9, 1–44.
- OLIVETTI, C., J. PAN, AND B. PETRONGOLO (2024): “The Evolution of Gender in the Labor Market,” Working Paper 33153, National Bureau of Economic Research.
- OLIVETTI, C. AND B. PETRONGOLO (2016): “The evolution of gender gaps in industrialized countries,” *Annual review of Economics*, 8, 405–434.
- OLSSON, J. (2025): “Singles, couples, and their labor supply: Long-run trends and short-run fluctuations,” *American Economic Journal: Macroeconomics*, 17, 1–34.
- ONOFRI, M., I. BERNIELL, R. FERNÁNDEZ, AND A. MENDUIÑA (2026): “Understanding Latin America’s Fertility Decline: Age, Education, and Cohort Dynamics,” Tech. rep., National Bureau of Economic Research.
- PADILLA-ROMO, M. AND F. CABRERA-HERNÁNDEZ (2019): “Easing the Constraints of Motherhood: The Effects of All-Day Schools on Mothers’ Labor Supply,” *Economic Inquiry*, 57, 890–909.
- PARENTE, R. M. AND F. MOREAU (2024): *The dynamics of trade integration and fragmentation in lac*, International Monetary Fund.
- PETERS, M. AND C. WALSH (2021): “Population growth and firm dynamics,” Tech. rep., National Bureau of Economic Research.
- RANSOM, M. R. AND R. L. OAXACA (2010): “New Market Power Models and Sex Differences in Pay,” *Journal of Labor Economics*, 28, 267–289.
- RUGGLES, S., L. CLEVELAND, R. LOVATON, S. SARKAR, M. SOBEK, D. BURK, D. EHRLICH, Q. HEIMANN, J. LEE, AND I. MN (2024): “Integrated public use microdata series, international: Version 7.5 [dataset],” *Minneapolis, MN: IPUMS*.
- RUGGLES, S., L. L. CLEVELAND, R. L. DÁVILA, S. SARKAR, M. SOBEK, D. BURK, D. E. EHRLICH, J. LEE, AND N. MERRILL (2025): “IPUMS International: Version 7.7 [dataset],” <https://doi.org/10.18128/D020.V7.7>, accessed: September 14, 2026.
- SANDERSON, W. C. AND S. SCHERBOV (2007): “A new perspective on population aging,” *Demographic research*, 16, 27–58.
- YAARI, M. E. (1965): “Uncertain lifetime, life insurance, and the theory of the consumer,” *The Review of Economic Studies*, 32, 137–150.
-

Appendix A: Additional tables and figures




Figure A1: Labor growth decomposition in major LAC economies

Sources: ILOSTAT; 2024 UN Populations Prospects; and IMF staff calculations. Working-age population is defined as 15 and older.

	(1)
	Wage (log)
Female (dummy)	−0.2923*** (0.0078)
Constant	9.0412*** (0.0029)
R ² -squared	0.8031
N	39,816,273

Table A1: Female wage penalty from Mincerian regression

Sources: IPUMS International ([Ruggles et al., 2024](#)); and IMF staff calculations. Notes: The regression estimates wage (*log*) as a function of sex and a rich set of controls (e.g., urban–rural status, geographic region, household head’s location, age, education, and employment characteristics). The model is estimated with high-dimensional fixed effects, robust standard errors clustered at the first administrative division level, and sampling weights, yielding an estimate of the wage penalty for women relative to men with similar observable characteristics. Samples from Brazil (1980, 1991, 2000, and 2010) and Mexico (2010, 2015, and 2020) are included in the analysis.






Figure A2: Household work and employment over the lifecycle
(Percent of working—age population)



(a) By gender and marital status



(b) Married women with and without children

Figure A3: Rising labor force participation among married women in major LAC economies (*Share*)

Sources: IPUMS International ([Ruggles et al., 2025](#)); Luxembourg Income Study; and authors' calculations. Notes: Missing 2020 IPUMS cells use the LIS bridge described in Appendix E.

Dependent var.	Labor force ratio (Female:Male)		
Model	OLS	OLS	GLM
	(1)	(2)	(3)
Children dependency ratio	-0.956*** (0.130)		
Children dependency ratio ²	-0.050 (0.126)		
Old age dependency ratio		14.666** (6.072)	31.615*** (1.872)
Old age dependency ratio ²		-37.322 (31.431)	
Constant	1.020*** (0.033)	-0.529* (0.291)	
R ² -squared	0.927	0.569	
AIC			1.481
BIC			-723.952
N	196	196	196

Table A2: Empirical relationship between dependency ratios and female labor force participation

Sources: IPUMS International ([Ruggles et al., 2024](#)); and IMF staff calculations. Notes: Column (1) presents the results of a regression analysis for female-to-male labor force ratio against the children dependency ratio (population aged 0–14 relative to the working-age population, 15–64), while Column (2) shows the corresponding relationship with the old-age dependency ratio (population aged 65+ relative to the working-age population). Column (3) estimates a binomial general linearized model (GLM) with a logistic link function to truncate the predicted values between zero and one. All columns include fixed effects for the first geographical administrative divisions. Standard errors are clustered using the same geographical units. Samples from Brazil (1980, 1991, 2000, and 2010) and Mexico (2010, 2015, and 2020) are included in the analysis.

Table A3: Probit Regression Results by Country and Gender

	Argentina		Brazil		Chile		Colombia	
	Women	Men	Women	Men	Women	Men	Women	Men
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Gender head	0.350*** (0.0131)	0.312*** (0.0182)	-0.0764*** (0.00460)	0.00980 (0.00597)	0.824*** (0.0102)	0.574*** (0.0120)	0.402*** (0.00571)	0.301*** (0.00840)
Education	0.226*** (0.00429)	0.0602*** (0.00583)	0.243*** (0.00142)	0.142*** (0.00211)	0.255*** (0.00300)	0.167*** (0.00380)	0.209*** (0.00187)	0.0372*** (0.00266)
Age	0.274*** (0.00281)	0.319*** (0.00313)	0.119*** (0.00119)	0.111*** (0.00166)	0.185*** (0.00196)	0.236*** (0.00211)	0.216*** (0.00121)	0.283*** (0.00148)
Married	-0.204*** (0.0132)	0.567*** (0.0184)	-0.252*** (0.00523)	0.229*** (0.00712)	-0.00834 (0.00944)	0.463*** (0.0120)	-0.304*** (0.00566)	0.499*** (0.00864)
Young children	-0.156*** (0.0191)	0.162*** (0.0316)	-0.339*** (0.00653)	0.0269*** (0.0101)	-0.0589*** (0.0141)	0.121*** (0.0219)	-0.0954*** (0.00859)	0.181*** (0.0145)
Children	-0.147*** (0.0148)	-0.000299 (0.0201)	-0.138*** (0.00568)	0.0895*** (0.00802)	-0.187*** (0.0106)	-0.0481*** (0.0142)	-0.106*** (0.00686)	-0.0776*** (0.00977)
Elders	-0.0270 (0.0299)	0.00423 (0.0370)	-0.178*** (0.00862)	-0.216*** (0.0107)	0.0234* (0.0140)	0.00283 (0.0173)	-0.0497*** (0.00941)	-0.0403*** (0.0132)
Age ²	-0.00334*** (3.47e-05)	-0.00388*** (3.92e-05)	-0.00166*** (1.39e-05)	-0.00160*** (1.86e-05)	-0.00225*** (2.39e-05)	-0.00276*** (2.62e-05)	-0.00267*** (1.52e-05)	-0.00343*** (1.86e-05)
Rural	-0.213*** (0.0110)	-0.181*** (0.0139)	-0.348*** (0.00462)	-0.216*** (0.00577)	-0.127*** (0.00935)	-0.0102 (0.0110)	-0.207*** (0.00784)	0.241*** (0.0104)
Constant	-5.192*** (0.0494)	-5.121*** (0.0546)	-1.755*** (0.0252)	-0.926*** (0.0363)	-4.125*** (0.0350)	-4.551*** (0.0371)	-4.140*** (0.0209)	-4.492*** (0.0248)
Observations	66,595	61,879	479,824	396,936	134,582	121,124	311,079	265,489

	Mexico		Peru		Uruguay	
	Women	Men	Women	Men	Women	Men
	(9)	(10)	(11)	(12)	(13)	(14)
gender head	0.319*** (0.00499)	0.245*** (0.00709)	0.397*** (0.0217)	0.488*** (0.0326)	0.259*** (0.0110)	0.262*** (0.0152)
education	0.166*** (0.00138)	0.0190*** (0.00189)	0.0781*** (0.00578)	0.0609*** (0.00883)	0.215*** (0.00397)	0.122*** (0.00564)
age	0.190*** (0.000864)	0.259*** (0.00103)	0.150*** (0.00344)	0.192*** (0.00442)	0.282*** (0.00240)	0.330*** (0.00271)
married	-0.448*** (0.00439)	0.484*** (0.00657)	-0.195*** (0.0198)	0.420*** (0.0300)	-0.166*** (0.0115)	0.490*** (0.0157)
young children	-0.0772*** (0.00607)	0.150*** (0.0102)	-0.129*** (0.0223)	0.104*** (0.0346)	-0.0490*** (0.0185)	0.149*** (0.0302)
children	-0.103*** (0.00464)	0.0296*** (0.00647)	-0.0437*** (0.0189)	0.0341 (0.0245)	-0.143*** (0.0141)	-0.000352 (0.0185)
elders	-0.0323*** (0.00735)	-0.0763*** (0.0101)	-0.0568*** (0.0290)	-0.0744* (0.0387)	-0.0843*** (0.0198)	-0.0160 (0.0250)
age sq	-0.00234*** (1.09e-05)	-0.00327*** (1.31e-05)	-0.00170*** (4.35e-05)	-0.00227*** (5.68e-05)	-0.00347*** (2.92e-05)	-0.00409*** (3.34e-05)
rural	-0.0660*** (0.00368)	0.139*** (0.00497)	0.349*** (0.0182)	0.365*** (0.0244)	-0.119*** (0.0224)	0.406*** (0.0292)
Constant	-3.488*** (0.0147)	-3.783*** (0.0170)	-2.553*** (0.0574)	-2.910*** (0.0706)	-4.947*** (0.0419)	-5.276*** (0.0476)
Observations	582,438	535,415	36,739	34,220	92,472	84,053

Note: The coefficients are the marginal effects. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.




Figure A4: Median age
(Years)

Sources: UN Population Prospects and IMF staff calculations. Note: Calculations for 2022-2100 based on median population projections.




Figure A5: Youth unemployment by age group
(Percentage)

Sources: Sources: EPH (INDEC); PNADC (IBGE); ENE(INE); GEIH (DANE); ENOE (INEGI); ENAHO (INEI); EPHC (INE); ECH (INE); and IMF staff calculations. Note: Unemployment rates based on active population. For youth (15–24), students are excluded.




Figure A6: Obstacles to FLFP

Sources: EPH (INDEC); PNADC (IBGE); ENE(INE); GEIH (DANE); ENOE (INEGI); ENAHO (INEI); EPHC (INE); ECH (INE); and IMF staff calculations. Note: Results from Oaxaca-Blinder decomposition of a probit regression. Observable worker features include education, household structure (such as marital status and the presence of children and/or older adults), age group, and location (urban or rural). The errors bars denote the 95 percent confidence interval for the unexplained gap. Self-reported by working-age population not in the labor force. “Household responsibilities” include taking care of children or other relatives and other chores, and cannot be separated due to data limitations. “Other reasons” vary across countries.





Figure A7: The role of individual characteristics – country, gender, age, and education – in LFP gaps

Sources: ILOSTAT; IMF Gender Data Hub; and IMF staff calculations. Notes. Aggregates are simple averages based on modeled ILO estimates. AEs = advanced economies; Africa = Sub-Saharan Africa; EMDE = emerging market and developing economies; LAC = Latin America and the Caribbean; MECA = Middle East and Central Asia.



(a) LFPR
(Fraction of pre-child rate)

(b) Informality rate
(Pre-child rate=100)

Figure A8: Parenthood effects on labor force

Sources: Berniell and others (2024); ENAHO 2007–22; and authors' calculations.




Figure A9: Elasticity of substitution in the production function.
(Ratio in logs)

Sources: IPUMS International (Ruggles et al., 2024);. Notes: Fixed effects at the first geographical administrative division level are included. The scatter plot displays binned estimates of the log female-to-male labor force ratio against the log median wage ratio (female-to-male), with a first-order polynomial fit and 95% confidence intervals.

Appendix B: Did the pandemic affect LFPR Trends?

The severe disruptions caused by the pandemic and the accompanying containment measures may have set in motion structural changes to participation trends. Although employment has returned to its pre-pandemic trend in most Latin American labor markets in 2023, the aggregate recovery may envelope different patterns across demographic groups, given the unequal effects of the COVID-19 pandemic [International Monetary Fund \(2021\)](#). We provide a more detailed analysis of the observed changes in participation between 2013–19 and between 2019–23, including a shift-share analysis to disentangle the effect of population shifts from changes in the underlying participation rates of different demographic groups.

The results indicate that the aggregate recovery conceals two important changes—after a small but generalized increase in female participation up to 2019 (except for Colombia), the recovery of female participation rates has been heterogeneous across the region, and the participation of older adults, that was increasing before 2019, was still subdued by the end of 2023. Importantly, the latter can be explained by behavioral changes that resulted in less participation within this group; as the share of population 55 and older is increasing, their lower participation rates weigh more on the overall participation rate (Figure B1; Table B1).




Figure B1: LFP before and after the pandemic by age and gender

Labor markets were severely disrupted by the COVID-19 pandemic. As containment measures were imposed, employment and labor force participation rates contracted rapidly in Latin America and across the world [International Monetary Fund \(2023\)](#), [International Monetary Fund \(2021\)](#). This subsection breaks down the changes in participation rates before and after the pandemic across different demographic groups to assess whether the shock changed any underlying pre-pandemic trends.

The overall changes in participation rates are analyzed for two periods: 2013–19 and 2019–23. The pre-pandemic period was characterized by diverging male and female trends—in line

with historical patterns, female labor force participation continued to increase, except in Colombia, while male labor force participation declined across the region, except in Paraguay, especially among young men (Box Figure 2.1).

	LFPR Change	Population Shift	Participation Shift
2013–19			
Total population	-0.1	-0.4	0.3
Female			
15–24		-0.9	-0.2
25–54		-0.5	1.2
55–64		0.6	0.4
65+		0.3	0.1
<i>Total</i>	1.1	-0.6	1.6
Male			
15–24		-1.2	-0.9
25–54		-0.5	-0.4
55–64		0.9	0.0
65+		0.5	0.1
<i>Total</i>	-1.4	-0.3	-1.2
2019–23			
Total population	-1.0	-0.2	-0.8
Female			
15–24		-0.6	-0.2
25–54		-0.3	0.3
55–64		0.3	-0.1
65+		0.2	-0.2
<i>Total</i>	-0.6	-0.3	-0.2
Male			
15–24		-0.7	-0.7
25–54		-0.1	-0.3
55–64		0.4	-0.2
65+		0.4	-0.3
<i>Total</i>	-1.5	-0.1	-1.4

Table B1: Shift-share decomposition of changes in LFP before and after the pandemic

Sources: 2024 UN Population Prospects; ILOSTAT; and IMF staff calculations. Note: The total change in the labor force participation is the sum of the population shift, the participation shift, and a second-order interaction term (omitted). LFPR = labor force participation rate.

Heterogeneity across population subgroups: a shift-share decomposition

The overall change in participation, however, is determined by the interaction between demographic changes and changes in participation rates. Using a shift-share analysis, this box then decomposes the total change in participation with respect to a base year into a population share shift, that

captures the demographic effect, and a participation rate shift, that captures changes in the participation rate of each group, holding the population composition constant. This decomposition can be expressed as:

$$p_t - p_0 = \sum_g [p_0^g (s_t^g - s_0^g) + s_0^g (p_t^g - p_0^g) + (p_t^g - p_0^g)(s_t^g - s_0^g)] \quad (\text{B.1})$$

where p_t denotes the aggregate participation rate in year t , and p_t^g and s_t^g denote the participation rate and the population share of group g in year t , respectively. The first term captures the demographic effect, by weighting changes in the population shares of each demographic group by its initial participation rate, while the second term measures the change in the overall participation attributable to changes in participation rates of the different groups, assuming constant population shares. The interaction term is typically negligible for relatively short periods.

Appendix C: The general CES labor aggregator

The body of the paper aggregates male and female labor with a Cobb-Douglas function. This appendix sets out the general constant-elasticity case, shows that the body's specification is its $\eta \rightarrow 1$ limit, and reports what is lost by imposing that limit. The short answer is: very little, provided the share parameter is calibrated to the base-year female revenue share rather than backed out of the relative-demand regression.

The production block. Let effective male and female labor inputs be $L_{mt} = \sum_j \frac{N_{jt}}{2} \bar{\ell}_j^m$ and $L_{ft} = \sum_j \frac{N_{jt}}{2} \bar{\ell}_j^f$. Suppose the two combine into aggregate labor according to

$$L_t = \left[\xi L_{ft}^{\frac{\eta-1}{\eta}} + (1-\xi) L_{mt}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (\text{C.1})$$

where $\eta > 0$ is the elasticity of substitution between female and male labor and $\xi \in (0, 1)$ is a share parameter. Output is $Y_t = K_t^\alpha (Z_t L_t)^{1-\alpha}$ as in the body, so the marginal products of the two labor inputs are

$$MPL_{ft} = (1-\alpha) \frac{Y_t}{L_t} \xi \left(\frac{L_t}{L_{ft}} \right)^{\frac{1}{\eta}}, \quad MPL_{mt} = (1-\alpha) \frac{Y_t}{L_t} (1-\xi) \left(\frac{L_t}{L_{mt}} \right)^{\frac{1}{\eta}}. \quad (\text{C.2})$$

Men are paid their marginal product; women are paid a fraction $1 - \pi$ of theirs, where π is the wage penalty of Section 5 and coincides with the wedge τ^f in the household budget constraint. Taking the ratio and rearranging gives the relative demand schedule

$$\ln \left(\frac{L_{ft}}{L_{mt}} \right) = \eta \ln \left[\frac{(1-\pi)\xi}{1-\xi} \right] + \eta \ln \left(\frac{w_{mt}}{w_{ft}} \right). \quad (\text{C.3})$$

Equation (C.3) is what the subnational regression of Section 5 estimates. Its slope identifies η ; its intercept identifies the composite $(1-\pi)\xi/(1-\xi)$, and hence ξ given π , *provided the left-hand side is the ratio of effective labor inputs*. We return to this proviso below, because it is the reason we do not calibrate ξ this way.

The Cobb-Douglas case. As $\eta \rightarrow 1$, (C.1) converges to $L_t = L_{ft}^\xi L_{mt}^{1-\xi}$, the specification used in Section 4. Two properties make this case attractive beyond its familiarity. First, it is the only value of η for which the household's participation decision admits the closed form of Section 4, which is what allows the FLFP response to be decomposed analytically into a children-dependency channel, an old-age-dependency channel and a residual. Second, the share parameter acquires a direct empirical counterpart: with $w_{ft}L_{ft} = (1-\pi)(1-\alpha)\xi Y_t$ and $w_{mt}L_{mt} = (1-\alpha)(1-\xi)Y_t$, the female share of the labor bill is

$$s \equiv \frac{w_f L_f}{w_f L_f + w_m L_m} = \frac{(1-\pi)\xi}{(1-\pi)\xi + 1 - \xi}, \quad \text{so} \quad \xi = \frac{s}{(1-\pi)(1-s) + s}. \quad (\text{C.4})$$

Because women are paid less than their marginal product, ξ exceeds s : the observed earnings share understates women's elasticity contribution to output, and the gap between the two is exactly the wage penalty. This is the sense in which the gender wage gap causes the national accounts to understate what female labor contributes at the margin. We use (C.4) to calibrate ξ country by country.

Cost of the unit-elasticity restriction. Imposing $\eta = 1$ when the true elasticity exceeds one makes the model conservative, for two reinforcing reasons. The revenue share of female labor in (C.1) is $\theta_f = \xi L_f^{\frac{\eta-1}{\eta}} / [\xi L_f^{\frac{\eta-1}{\eta}} + (1-\xi)L_m^{\frac{\eta-1}{\eta}}]$, which is constant at ξ under Cobb-Douglas but rises with L_f when $\eta > 1$; and the relative wage falls by $1/\eta < 1$ rather than one-for-one as female labor expands, so the wage compression that discourages further participation is weaker. Both channels imply that a given increase in female labor supply raises output by more under $\eta > 1$ than under $\eta = 1$.

The magnitude, however, is small once ξ is calibrated to match the base-year revenue share. Holding $s = 0.33$ and $L_f/L_m = 0.55$ at the base year, recalibrating ξ so that (C.1) reproduces that share under each value of η , and raising female labor input by 20 percent, aggregate labor rises by 7.50 percent under Cobb-Douglas, 7.59 percent at $\eta = 1.3$, and 7.71 percent at $\eta = 2.0$: differences of 1.3 and 2.9 percent of the effect, well inside the uncertainty attaching to the demographic projections themselves. The restriction binds on the level of ξ , not on the elasticity, which is why the calibration target matters and the choice of η does not.

The intercept and the share parameter. The regression underlying (C.3) uses $\ln(LFPR_f/LFPR_m)$ on the left-hand side—the ratio of participation *rates*, not of effective labor inputs. The two differ by the female-to-male ratio of working-age population and by any systematic difference in hours and efficiency units per participant, both of which load onto the intercept. The composite

recovered from the intercept is therefore not $(1 - \pi)\xi / (1 - \xi)$, and the implied share parameter is correspondingly biased: taken at face value it delivers $\xi = 0.12$, which by (C.4) corresponds to a female share of the labor bill of about 9 percent, against roughly a third in the data. We therefore use the regression for the slope alone and calibrate ξ from the earnings share directly, reporting the intercept-based value only as a diagnostic.

Implications for the counterfactuals. Nothing in the growth decomposition of Section 2 or the general-equilibrium differencing of Section 5.3 depends on η . Under any member of the family (C.1), $\Delta \ln L$ is a share-weighted sum of $\Delta \ln L_f$ and $\Delta \ln L_m$; the Cobb-Douglas case simply fixes those weights at ξ and $1 - \xi$ rather than letting them drift with the input ratio. What changes across specifications is the weight, not the structure of the decomposition—which is why calibrating ξ correctly matters considerably more than choosing η .

Appendix D: Computing the perfect-foresight transition

We solve the model as a deterministic perfect-foresight transition using the sequence-space Jacobian method of Auclert et al. (2021), whose code we gratefully acknowledge.

Horizon and terminal condition. The transition runs from 2016 ($t = 0$) to 2300 ($t = T = 284$). The initial condition is each country’s observed capital stock and age distribution in 2016; the terminal condition is the asymptotic balanced growth path of Proposition 1, along which the age distribution, the capital–labor ratio and normalized age-specific consumption are constant. The horizon is long enough that the terminal condition does not bind over the 2024–2060 window we report, since the demographic forcing has converged well before 2300.

Demographic forcing. Age-specific survival probabilities ϕ_j and the sequence of birth cohorts are taken from the 2024 UN World Population Prospects under the median-fertility variant, and are held fixed across scenarios. Because survival profiles are constant over time in the model, the age distribution μ_{jt} evolves only through fertility and the resulting convergence dynamics.

Computing the transition. We solve for the perfect-foresight transition path between 2016 ($t = 0$) and 2300 ($t = 284 \equiv T$) using sequence-space Jacobian methods, iterating on the path of the interest rate until goods and asset markets clear in every period. Aggregate labor supply along the path is shifted by the participation response described above, and capital accumulation and factor prices adjust to it.

Algorithm. We proceed in four steps. First, demographics are simulated forward from the initial population distribution and the forcing variables, which—together with the effective labor supply profile and the retirement policy—delivers a path for aggregate labor $\{L_t\}_{t=0}^T$. Second, given a

candidate path for the interest rate $\{r_t\}_{t=0}^T$ and the technology parameters, the optimal capital–labor ratio and hence the wage path follow from the firm’s first-order conditions (14). Third, given the government’s fiscal rule, we solve each cohort’s problem backwards from age J using the Euler equation (17) and the female labor supply rule (16), and aggregate across cohorts to obtain private wealth W_t . Fourth, we iterate on the 285×1 interest-rate path using a Newton-based method, with Jacobians computed by the algorithm of Auclert et al. (2021), until goods and asset markets clear in every period. The government adjusts G_t so that the debt-to-output ratio in (12) remains fixed.

Country coupling. The seven LAC economies are solved jointly with the United States, China and a low-income-country aggregate, so that the world interest rate is consistent across scenarios. Holding the country set fixed across runs is what makes the scenario differences of Section 5.3 interpretable: the counterfactuals differ only in the female participation path, not in the general-equilibrium environment.

Scenario differences. Each counterfactual is a complete general-equilibrium solve. The decomposition in Section 5.3 differences two such solves—the demographic-only baseline, in which female participation is frozen at its 2016 profile, and the endogenous-participation path—so that capital accumulation and the interest rate respond in both. Labor-augmenting productivity Z_t is exogenous and common across scenarios, so it cancels in the difference; only its common level remains an assumption. Endogenizing productivity would be a separate extension.

Appendix E: Data

E.1 Data families and uses

The paper relies on six types of data. National household surveys support the seven-country probit analysis. IPUMS International (Ruggles et al., 2024) supports the long-run household-composition series and the Brazil–Mexico subnational regressions. Luxembourg Income Study (LIS) files provide growth rates for missing 2020 descriptive cells. ILOSTAT provides aggregate participation series. The 2024 UN World Population Prospects provides demographic paths. Penn World Table 10.0 provides capital stocks and capital-output ratios.

The national surveys are EPH for Argentina, PNADC for Brazil, ENE for Chile, GEIH for Colombia, ENOE for Mexico, ENAHO for Peru, and ECH for Uruguay. Table E2 reports the estimation counts reproduced in Appendix Table A3. The supplied files do not identify the survey waves or the complete restriction sequence for these probits.

E.2 IPUMS samples

Long-run descriptive series. The IPUMS data contains 39 country-year samples for eight countries. Table E1 reports the exact IPUMS sample codes and weighted person totals after the restrictions

below. The person-level extract is not supplied, so unweighted counts by country-year cannot be recovered.

Table E1: IPUMS samples used in the long-run household-composition series

Country	IPUMS sample	Year	Weighted persons
Argentina	AR1970A	1970	16,065,400
	AR1980A	1980	19,449,884
	AR1991A	1991	22,556,013
	AR2001A	2001	26,015,180
Brazil	BR1960A	1960	42,924,735
	BR1970A	1970	54,382,482
	BR1980A	1980	72,448,107
	BR1991A	1991	95,176,892
	BR2000A	2000	119,564,413
	BR2010A	2010	144,821,878
Chile	CL1970A	1970	5,381,840
	CL1982A	1982	7,468,460
	CL1992A	1992	9,421,170
	CL2002A	2002	11,240,890
Colombia	CO1973A	1973	10,857,390
	CO1985A	1985	16,849,078
	CO1993A	1993	20,589,990
	CO2005A	2005	27,300,373
Mexico	MX1970A	1970	25,376,400
	MX1990A	1990	48,850,540
	MX1995A	1995	58,398,166
	MX2000A	2000	63,501,454
	MX2010A	2010	78,615,717
	MX2015A	2015	86,308,122
	MX2020A	2020	93,214,983
Peru	PE1993A	1993	13,734,760
	PE2007A	2007	19,091,830
	PE2017A	2017	20,505,423
Uruguay	UY1963A	1963	1,811,780
	UY1975A	1975	2,013,000
	UY1985A	1985	2,132,400
	UY1996A	1996	2,333,770
	UY2006A	2006	2,355,228
	UY2011A	2011	2,478,420

Source: IPUMS International ([Ruggles et al., 2024](#)); and authors' calculations. *Notes:* The unit is a person. Weighted persons sum the IPUMS person weights after the age, sex, marital-status, child, and labor-force restrictions.

National-survey probit sample. The national-survey counts reported are estimation observations, not full survey sample sizes.

Table E2: National-survey probit estimation samples

Country	Survey	Women	Men	Total
Argentina	EPH (INDEC)	66,595	61,879	128,474
Brazil	PNADC (IBGE)	479,824	396,936	876,760
Chile	ENE (INE)	134,582	121,124	255,706
Colombia	GEIH (DANE)	311,079	265,489	576,568
Mexico	ENOE (INEGI)	582,438	535,415	1,117,853
Peru	ENAHO (INEI)	36,739	34,220	70,959
Uruguay	ECH (INE)	92,472	84,053	176,525
Total		1,703,729	1,499,116	3,202,845

Sources: National statistical agencies and authors' calculations. Notes: Counts are estimation observations from Appendix Table A3. They are not full survey sample sizes. The supplied materials do not include the generating code or survey-wave identifiers.

Subnational sample. The subnational participation, dependency, and earnings files use only Brazil and Mexico. Each observation is one harmonized first-level administrative division in one sample year.

Table E3: Subnational IPUMS analytical cells

Country	IPUMS samples	GEOLEV1 units	Cells
Brazil	BR1980A, BR1991A, BR2000A, BR2010A	25	100
Mexico	MX2010A, MX2015A, MX2020A	32	96
Total	Seven samples	57	196

Source: IPUMS International (Ruggles et al., 2024); and authors' calculations. Notes: Every sample contains all listed GEOLEV1 units. The GEOLEV1 variable label states that boundaries are consistent over time.

Limitations Three limitations bear directly on the interpretation of the results and are stated here rather than left for a reader to discover. First, our microdata window opens in 1990. Latin America's transition in female participation was already well advanced by then, so the national secular trend that Section 5 cannot separate from demography is observed only in its later phase. The 1960, 1970 and 1980 census rounds would place that trend against a period in which participation and child dependency did not move together. Second, measured participation includes informal, unpaid family and subsistence work, and the share of activity so classified varies across the region. The part of this that the norm-state inversion would otherwise absorb runs through the age margin rather than through informality as such, which is why Section 5.2 evaluates (5.2) on the working-age population: the excess is concentrated at ages 65 and over, where pension coverage rather than gender norms governs the decision to work. Appendix F shows that across Brazilian states informality is associated with *lower* measured female participation, so the informality reading cannot be applied to the cross-section without a mechanism. What remains uncertain is not the level of gender norms but its speed, which is why we hold it fixed in the counterfactuals rather

than projecting it. Third, the subnational dependency–participation relationship is estimated on Brazil and Mexico, the two countries with enough census rounds to support it. The cross-country dispersion reported in Section 5 therefore comes from differences in demographic paths and in ξ_c , not from differences in estimated behavior. Appendix F takes the inversion itself to Brazilian states and finds a within-country dispersion in gender norms nearly as wide as the dispersion across the seven countries, which is a caution against reading the country cross-section as a ranking of national cultures.

Appendix F: Gender norms within a country

Section 5.2 recovers a norm state ρ for each country. However note that equation (5.2) can be applied to any population regardless of the aggregation level. Given female participation, the child dependency ratio and the production share for any population, it returns the ρ that rationalizes them. This appendix applies it to Brazilian states. The exercise is descriptive and supports the interpretation of the parameter in two ways: i) the ordering recovered is plausible and ii) it highlights that substantial gender norms dispersion is present within-states.

Construction. We use the Brazilian continuous household survey (PNAD Contínua) for 2012–2025, restricted to the sixteen states with continuous coverage, giving 224 state-years and 7.26 million person-records. Every input to (5.2) is measured in the same survey. The child dependency ratio is the weighted population aged 0–15 relative to that aged 15–64, on the convention of Section 5. Female participation is the weighted participation rate of women aged 15–64, the population used in (5.2). The share parameter ξ_j is calculated state by state from the female share of the labor bill among the employed with strictly positive earnings, through (C.4), so that a state’s ξ_j is measured rather than borrowed from the national value.

As a check on the construction, the sixteen-state aggregate reproduces the national inversion: it returns $\rho = 0.37$ for 2023 against 0.43 for Brazil in Section 5.2. The residual is the difference between modeled and survey participation and between projected and surveyed dependency, not a difference in method; the survey measures participation about four points higher than the ILO modeled series does. Aggregates across states in this appendix weight each state by its survey sample size, which tracks but does not exactly equal its share of the population.

Table F1: gender norms across Brazilian states, 2025

	ρ_j	Female LFP	Child dep.	ξ_j	Informality
Minas Gerais	0.266	0.719	0.266	0.444	0.225
Paraná	0.286	0.707	0.270	0.452	0.224
São Paulo	0.302	0.702	0.266	0.464	0.247
Rio Grande do Sul	0.304	0.702	0.268	0.466	0.222
Goiás	0.312	0.684	0.273	0.446	0.280
Distrito Federal	0.312	0.692	0.282	0.468	0.237
Mato Grosso do Sul	0.321	0.674	0.320	0.467	0.239
Espírito Santo	0.357	0.643	0.302	0.445	0.245
Rio de Janeiro	0.382	0.644	0.260	0.447	0.345
Sergipe	0.419	0.617	0.313	0.471	0.294
Bahia	0.425	0.654	0.245	0.499	0.351
Pernambuco	0.448	0.591	0.302	0.448	0.360
Pará	0.481	0.596	0.258	0.456	0.485
Acre	0.528	0.555	0.350	0.481	0.298
Ceará	0.550	0.563	0.275	0.467	0.419
Maranhão	0.569	0.559	0.301	0.491	0.447

Sources: PNAD Contínua 2025; authors' calculations. Notes: ρ_j is recovered from (5.2) at $\gamma = 0.45$ and $\pi = 0.25$, with ξ_j measured state by state through (C.4). Female LFP is the participation rate of women aged 15–64. Child dependency is the population aged 0–15 relative to that aged 15–64. Informality is the share of employed women who do not contribute to social security. States are ordered by ρ_j ; a lower value denotes a smaller share of child-related home time borne by women.

Within-country dispersion. The dispersion inside one country is wider than the dispersion across the region: ρ_j runs from 0.27 in Minas Gerais to 0.57 in Maranhão, a ratio of 2.1, which exceeds the ratio of about 2 across the seven countries of Section 5.2. Maranhão and Ceará lie above every country in our sample and Minas Gerais below all but Peru.

Informality Section 5.2 attributes part of the regional cross-section to informal, unpaid family and subsistence activity entering measured participation, which would lower ρ where informality is high. Across Brazilian states the association runs the other way: the correlation between $\ln \rho_j$ and the female informality rate is +0.85 in 2025 and +0.56 pooled across all state-years with year effects removed, and the correlation between informality and female participation is -0.78 . Within Brazil, informality goes with *less* measured female participation, not more. Because informality is correlated with income and urbanization, these results do not resolve the cross-country interpretation. They do, however, suggest that a negative relationship between informality and ρ is not a general feature of the data.. The mechanism supported by the data operates through the age composition of participation rather than through informality per: measuring participation on the population aged 15 and over would add activity at ages where pension coverage rather than gender norms governs the decision to work (see section 5.2).

Diverging trends. Regressing the annual log change in ρ_j over 2012–2025 on its initial log level gives a coefficient of -0.002 with a standard error of 0.026 , and the cross-state standard deviation of $\ln \rho_j$ rises from 0.16 to 0.25 . Fourteen years is a short window and it contains the pandemic, so this is not decisive but it supports holding ρ fixed in the counterfactuals of Section 5.1 rather than projecting convergence.

Counterfactual . The counterfactuals in Section ??5. neither estimate corresponds to the long-run elasticity identified from the census-based evidence in Section 5. dvanced-economy values as a benchmark. The state-level distribution provides an alternative benchmark drawn from within the region. Holding each state’s child dependency ratio and ξ_j at their 2025 values and moving every state’s ρ_j down to a domestic benchmark, aggregate female participation in the sixteen states rises from 65.4 percent to 67.4 percent at the median ρ , to 69.5 percent at the twenty-fifth percentile, and to 72.3 percent at the level of Minas Gerais. Applying $d\ln Y = \bar{\xi} d\ln L_f$, these correspond to output gains of $1.4, 2.8$ and 4.6 percent. The largest is of the same order as the demographic effect this paper quantifies for the region as a whole, and it is measured against a Brazilian state rather than against Nordic countries.

Limitations Two limitations are worth noting, and both are properties of the data rather than of the method. First, ρ is a residual. Across states the correlation between $\ln \rho_j$ and female participation is -0.98 , so the cross-section of ρ_j is the cross-section of participation after the demographic and share terms have been removed. Second, this panel identifies ρ but not γ . Within-state variation in child dependency over 2012–2025 is small, short-run and dominated by migration and cyclical composition, and a two-way fixed-effects regression on these data returns an elasticity of female participation to child dependency of $+0.13$ and of the female-to-male ratio of $+0.02$, neither estimate corresponds to the long-run elasticity identified from the census-based evidence in Section 5. Instead, γ is pinned from the elasticity of the female-to-male *ratio* and then applied to the law of motion for the female rate.