

A Glass Floor in Microentrepreneurship? Gender, Care Work, and a Widening Profit Gap

Mauricio Villena

`mauricio.villena@udp.cl`

Universidad Diego Portales

María José Quinteros

Universidad de Santiago de Chile

Research Article

Keywords: gender gap, microentrepreneurship, care work, informality, entrepreneurial performance, Chile

Posted Date: September 30th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10988479/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Glass Floor in Microentrepreneurship? Gender, Care Work, and a Widening Profit Gap

Abstract

Why can gender gaps in entrepreneurial performance persist—and even widen—when women match or exceed men on many observable capabilities usually associated with business success? Using four waves of Chile’s nationally representative Microentrepreneurship Survey (EME, 2017–2025; $N = 28,174$), we examine how gender differences in microenterprise performance have evolved and what accounts for them. In the waves in which each measure is available, women match or exceed men in education, internet use, digital-sales participation, and business training, yet their businesses earn substantially lower profits. The gender profit gap widened by roughly 50% between 2019 and 2025, and the increase remains statistically significant after adjusting for observable characteristics. Decomposition exercises indicate that the widening reflects primarily changes in relative returns rather than shifts in the composition of male and female entrepreneurs. The disadvantage is also concentrated among the lowest-profit businesses—a “glass floor” pattern that remains significant in direct distributional tests. Women report 1.9–2.8 times as many hours of direct unpaid care as men. Measured care time is strongly associated with necessity-driven entry into entrepreneurship but accounts for little of the profit or employer-status gaps. The findings suggest that convergence in observable entrepreneurial capabilities is not sufficient to generate convergence in entrepreneurial returns, highlighting the conditions under which women and men convert similar measured resources into business outcomes.

Keywords: gender gap, microentrepreneurship, care work, informality, entrepreneurial performance, Chile

JEL Classification: J16, J46, L26, O17

1 Introduction

Women make up roughly 40% of Chile’s 2.1 million microentrepreneurs, running businesses that span similar sectors, adopt the same digital tools, and — if anything — receive more business training than the businesses men run. Yet, in every wave for which net profit is consistently measurable (2019, 2022, and 2025), women’s businesses generate roughly half the profit of men’s. The central question is therefore not simply whether a gender performance gap exists, but why it has persisted — and widened — even as several measured entrepreneurial capabilities have converged.

Several conventional explanations receive little support in the data. Occupational segregation, as measured by the Duncan dissimilarity index, shows no economically meaningful trend from 2017 to 2025. Women are not less educated, less likely to use the internet, or less likely to have received business training; on training and digital adoption specifically, women exceed men. Nor does women’s higher informality rate alone account for the profit differential, although informality is associated with a larger profit penalty for women. After three complementary decomposition methods are applied, most of the profit gap remains unexplained by the business characteristics, observable human capital, and digital-engagement measures available in the survey.

This puzzle is consistent with gender-aware approaches to entrepreneurship that distinguish access to resources from the ability to convert those resources into entrepreneurial outcomes. Entrepreneurial-ecosystem research cautions against assuming that entrepreneurs with similar observable resources face equal opportunities to deploy them ([Brush et al., 2019](#)); recent evidence also questions a simple “female underperformance” interpretation by showing that lower realized returns need not imply lower underlying productivity ([Mickiewicz & Nguyen, 2025](#)). This distinction resonates with a broader argument that visible measures of gender inequality often obscure deeper institutional and social barriers not captured by individual-level characteristics ([Agarwal, 2026](#)). From a policy perspective, this distinction matters because interventions that focus on correcting presumed individual skill deficits may leave broader structural constraints untouched ([Foss et al., 2019](#)). Our empirical design speaks directly to this capability-to-returns gap. This paper makes four contributions. First, we build and validate a pooled repeated cross-section spanning four waves of EME (2017, 2019, 2022, 2025) — to our knowledge the longest such harmonized series constructed for gender analysis of Chilean microentrepreneurship. This series shows that the gender profit gap is not a static feature of the microentrepreneurial sector: it has grown by roughly half in six years, spanning a pre-pandemic baseline, the pandemic shock, and the subsequent recovery. This directly extends [Lauterbach & Durán-Sandoval \(2026\)](#), whose closely related

single-wave (2019) analysis of the same underlying survey cannot speak to whether the gap they document is stable, closing, or widening. We formally test this widening with a Female $\times$ Wave interaction. Because EME is a repeated cross-section rather than a panel, we further decompose the 2019–2025 change in the gap into a portion attributable to compositional change in who is entering microentrepreneurship and a portion attributable to changing relative returns. This decomposition matters because it keeps the widening claim from resting on an unstated assumption that the same women are being tracked over time.

Second, we situate the profit gap within a broader life-cycle account of gendered microentrepreneurship, spanning selection into entrepreneurship (necessity versus opportunity motivation), sectoral segregation, the transition from own-account work to becoming an employer, and access to training, rather than treating performance as an isolated outcome.

Third, we subject the headline profit-gap decomposition to three independent robustness checks: mean Oaxaca-Blinder (Oaxaca, 1973; Blinder, 1973), an unconditional quantile (RIF-OLS) decomposition (Firpo et al., 2009), and a Ñopo (Ñopo, 2008) matching decomposition that addresses the common-support critique of mean-based methods. The near-total “unexplained” finding survives all three. We further show that the gap is substantially larger among the lowest earners than the highest, and we confirm this distributional pattern with a direct bootstrap test of the difference between the 10th- and 90th-percentile gaps, rather than relying on a visual comparison of separately estimated quantile effects. Borrowing the “floor” terminology from the gender wage-gap literature (de la Rica et al., 2008), we characterize this as a “glass floor.” The pattern itself is not without precedent in the informal-enterprise literature: Islam & Amin (2023), studying labor productivity across informal firms in 14 developing economies, similarly document “sticky floors” rather than glass ceilings. Our contribution relative to that literature lies less in the qualitative pattern — productivity and profit gaps concentrated among the lowest performers appear to be a recurring feature of gender gaps in micro- and informal enterprise — than in two additional elements we bring to bear on a single, nationally representative setting: a repeated cross-section spanning eight years rather than a single wave, which allows a formal test of whether the gap is widening over time together with a composition/returns/interaction decomposition of that change; and a direct statistical test of the floor pattern itself, via bootstrap inference on the $\text{Gap}_{p10} - \text{Gap}_{p90}$ difference, rather than a description based on point estimates alone.

Fourth, we harmonize the EME time-use module fielded in three of the four waves and use its consistently comparable measure of direct unpaid-care time to test a theoretically motivated candidate mechanism for the gap: the “double burden” of unpaid care work documented elsewhere in the Latin American labor-informality literature (Berniell et al., 2021). Women report 1.9 to 2.8 times as many measured hours of direct unpaid care as

men. Measured care time is a statistically significant predictor of necessity-driven entry into entrepreneurship, but it accounts for only a small fraction of the profit and employer-status gaps. This specific, hours-based measure therefore helps account for the *entry* gap but not the *performance* gap, pointing toward other mechanisms — including access to capital, markets, and networks, as well as possible discriminatory treatment — that household survey data of this kind cannot identify directly.

The remainder of the paper proceeds as follows. Section 2 reviews the relevant literature on gender gaps in entrepreneurship, informality, and unpaid care work, situating our contribution against the closest existing studies. Section 3 describes the EME data, our harmonization procedure, and the construction of the pooled sample. Section 4 sets out the empirical strategy, including the regression specifications and the three decomposition methods. Section 5 presents results across the five dimensions of the gendered entrepreneurial pathway. Section 6 discusses limitations, policy implications, and directions for future research.

2 Literature Review

2.1 Gender gaps in entrepreneurial performance: endowments versus returns

The canonical tool for decomposing group differences in economic outcomes is the Oaxaca-Blinder decomposition (Oaxaca, 1973; Blinder, 1973), building on the earlier rate-decomposition method of Kitagawa (1955), which separates a mean gap into a portion attributable to differences in observable characteristics (“endowments”) and a residual portion attributable to differences in the estimated returns to those characteristics. The residual is conventionally interpreted, with appropriate caution, as an upper-bound proxy for discrimination or other unobserved structural disadvantage, since it also absorbs any relevant unobserved endowments (Jann, 2008). Applied to entrepreneurship, this framework has consistently found that a substantial share of gender earnings gaps among the self-employed is unexplained by observable human capital (Hundley, 2001; Lechmann & Schnabel, 2012; Simon & Way, 2016). More recently, Mickiewicz & Nguyen (2025) revisit the female-underperformance hypothesis and show that realized profitability and underlying productivity need not move together, underscoring the distinction between entrepreneurial capabilities and the economic returns entrepreneurs are able to realize.

The closest single study to the present paper is Lauterbach & Durán-Sandoval (2026), who apply a Kitagawa-Blinder-Oaxaca decomposition to the sixth wave of EME (2019,

$N = 7,340$) and find that women’s lower net microenterprise income is driven overwhelmingly (approximately 75%) by lower returns to observable characteristics such as education and access to bank credit, rather than by differences in those characteristics themselves. Our analysis both extends and, in one respect, diverges from theirs. We use the same underlying survey infrastructure, but pool four waves of entrepreneurial data (2017–2025) and three waves of harmonized profit data (2019–2025) rather than relying on a single cross-section. This lets us show that both the raw profit gap and its unexplained share have grown since the wave they study — a pattern a single wave cannot reveal. [Hazarika \(2020\)](#) applies an unconditional quantile decomposition to the gender income gap among rural Indian handloom micro-enterprises and finds the gap varies substantially across the income distribution — the direct methodological precedent for the quantile results we report in Section 5.

Sophisticated decomposition and quantile methods applied to gender gaps in developing-economy enterprise and labor outcomes are well established, including in this journal. Closest to our distributional results, [Islam & Amin \(2023\)](#) combine Kitagawa-Oaxaca-Blinder mean and quantile decomposition to study the gender labor-productivity gap across informal firms in 14 developing economies and document evidence of “sticky floors,” a pattern we discuss further below. [Hoang et al. \(2021\)](#) apply an endogenous-switching framework alongside an Oaxaca decomposition to innovation and gender-productivity gaps among Vietnamese SMEs, while [Lo Bue et al. \(2022\)](#) use harmonized global microdata and decomposition methods to document gender gaps in vulnerable employment across developing countries. [Vial & Hantou \(2015\)](#) apply quantile methods to a genuine four-wave household panel to show that the welfare returns to Indonesian micro-entrepreneurship are highest, in relative terms, for the poorest households — a heterogeneity-across-the-distribution result methodologically related to, though substantively distinct from, the glass-floor pattern we document. None of these studies, however, combines a repeated cross-section spanning multiple years with a formal test of whether a gender gap is widening over time, and none applies the composition/returns/interaction decomposition we use to separate compositional change from changing returns as an explanation for that widening.

A separate strand of the literature focuses on the gender gap in *entry* into entrepreneurship rather than performance conditional on entry. [Bernat et al. \(2017\)](#) apply a Fairlie-type decomposition ([Fairlie, 2003](#)) to entrepreneurial participation across Latin America and find that observable characteristics — education, risk tolerance, access to transport, prior exposure to a parent’s business — explain 23–38% of the entry gap, a larger explained share than we find for the performance gap conditional on entry. This is consistent with a broader pattern in the literature in which the barriers to becoming an entrepreneur and the barriers to succeeding as one are analytically distinct and only partially overlapping ([Guzmán &](#)

Kacperczyk, 2019; Rietveld & Patel, 2022).

2.2 Necessity entrepreneurship and the entrepreneurial life cycle

Women are disproportionately represented among necessity-driven, as opposed to opportunity-driven, entrepreneurs, particularly in developing-economy contexts (Vossenbergh, 2013; Strawser et al., 2021). Rietveld & Patel (2022) show across 97 countries that higher national gender inequality is associated with a larger gap in necessity- versus opportunity-based entrepreneurial motivation. Using detailed United States administrative firm data spanning 14 years and two recessions, Wheat & Mac (2024) find that necessity entrepreneurship explains little of the persistent gender revenue gap among American firms. We corroborate this finding in a different institutional context: even the strongest candidate mechanism we test (unpaid care work) leaves the great majority of the profit gap unexplained, echoing their conclusion that policies targeting entrepreneurial motivation alone are unlikely to close performance gaps.

2.3 Informality and the “motherhood effect”

Informality is a defining feature of Latin American labor markets, with close to half of the region’s employed population working informally as of 2025 (ILO, 2025). A substantial literature documents that women are more likely than men to operate informally, and that the gender gap in informality is closely tied to unpaid care responsibilities. Using harmonized household survey data across several Latin American countries, Berniell et al. (2021) show that motherhood substantially raises women’s, but not men’s, probability of informal employment — what they term the “motherhood effect.” Silupu et al. (2024), focusing specifically on Peruvian women microentrepreneurs, find that businesses run from the home — which facilitate the combination of paid and unpaid work — are systematically less likely to formalize. Particularly relevant for our mechanism test, Arráiz (2018) studies 2,408 micro and small businesses in Ecuador and finds that unpaid household work is the largest single factor explaining the gender profit differential, accounting for 33% of the observed gap; among otherwise comparable owners, the profit difference disappears. Our Chilean results provide a sharp contrast: measured unpaid care hours are strongly gendered but explain only a small fraction of the profit gap.

These findings motivate our use of the EME time-use module to test directly whether unpaid care-work burden mediates the gender gaps we document in entry, informality, and performance. To our knowledge, no published study using EME data has incorporated this module, including Lauterbach & Durán-Sandoval (2026). The theoretical grounding for this test comes from the feminist-economics literature on the “double burden” or “second shift”

— the observation that women’s labor-market participation is typically additive to, rather than a substitute for, unpaid domestic and care labor (Chant, 1994; Elson, 1999), with direct implications for the scale, sector, location, and formality of the businesses women are able to operate.

2.4 Digitalization and the gender gap

A separate literature asks whether digital technology adoption narrows or widens gender gaps in entrepreneurship. Cross-country evidence in Moeini Gharagozloo et al. (2023) finds that a country’s digital readiness is associated with greater female participation in opportunity-driven entrepreneurship. Policy reports similarly frame closing the “digital divide” as central to closing gender gaps in entrepreneurial outcomes more broadly (Trumbic & O’Connor, 2025; GEM, 2023), arguing that digital technologies expand the opportunity space available to women entrepreneurs by mitigating constraints in information, finance, and networks. Closest to our setting, Cirera et al. (2026), using novel firm-level technology-adoption data across developing economies, find that female-managed firms lag in adopting advanced, sector-specific production technologies but earn *higher* productivity gains conditional on adopting them, partially offsetting the initial gap; an Oaxaca-Blinder decomposition attributes part of the adoption gap itself to differences in managerial quality, access to government support, and sectoral composition.

Our results offer a complementary, and in one respect contrasting, picture. In pooled available-wave comparisons, Chilean women microentrepreneurs match or exceed men on internet use and digital-sales participation. The wave-specific pattern is more nuanced: internet use is close to parity in the earlier waves and higher among women in the later waves, while comparable digital-sales measures are available only in the later waves. Yet the profit gap has *grown*, not shrunk, over the same period, and an interaction test (Section 5) finds no evidence that internet use moderates the gender profit gap. Where Cirera et al. (2026) identify technology adoption as a channel through which the productivity gap narrows once women adopt advanced technologies, our setting is one in which adoption gaps are largely absent yet the performance gap persists and widens regardless. This suggests that closing observable digital-adoption gaps is not, on its own, a sufficient condition for closing performance gaps in every context.

2.5 Positioning of this paper

Taken together, the existing literature establishes that (i) gender performance and productivity gaps among entrepreneurs and informal-firm owners are typically dominated by an

unexplained residual rather than by differences in observable characteristics, and are frequently concentrated among the lowest performers rather than uniformly distributed (Islam & Amin, 2023; Hoang et al., 2021; Lo Bue et al., 2022); (ii) necessity entrepreneurship and informality are more prevalent among women and linked to unpaid care responsibilities; and (iii) technology and digital adoption gaps are an active margin for narrowing gender productivity differences (Cirera et al., 2026), though their relationship to performance gaps that persist *despite* adoption parity is less studied. Against this backdrop, the four contributions outlined in Section 1 — the eight-year pooled repeated cross-section and its widening test, the life-cycle framing that jointly models selection, structure, resources, and performance, the three-way robustness of the decomposition, and the direct test of the unpaid-care-work mechanism — are, to our knowledge, new to this literature in combination, even where individual elements have precedent in related work. In doing so, the paper links the gendered-ecosystem perspective of Brush et al. (2019) to the policy concern raised by Foss et al. (2019): if women already match or exceed men on several observable capabilities, closing remaining performance gaps may require attention to the conditions under which those capabilities are converted into returns, rather than additional skills-oriented interventions alone.

3 Data

3.1 Source and coverage

We use four waves of Chile’s official Microentrepreneurship Survey (Encuesta de Microemprendimiento, EME), conducted jointly by the Instituto Nacional de Estadísticas (INE) and the Ministry of Economy: EME V (fielded 2017), EME VI (2019, pre-pandemic), EME VII (2022, post-pandemic), and EME VIII (2025, the first wave to include an artificial-intelligence module). Each wave is a nationally representative, complex-design survey of individuals identified as microentrepreneurs in Chile’s National Employment Survey sampling frame, with sample sizes ranging from 6,579 (EME VII) to 7,808 (EME VI) respondents.

The EME is explicitly *not* a panel. It does not follow the same individual microentrepreneur across waves; each wave draws an independent cross-sectional sample. We therefore construct a *pooled repeated cross-section* rather than a panel, and all within-wave and across-wave comparisons in this paper should be interpreted accordingly: we observe how the gender gap looks in a representative sample of microentrepreneurs at each point in time, not how individual entrepreneurs’ outcomes evolve.

3.2 Harmonization

Because question wording, response coding, and even variable availability differ across the four survey waves, we constructed a harmonized crosswalk covering 45 source variables across the four waves. The resulting analysis files additionally contain constructed outcomes, identifiers, sample- design fields (strata, primary sampling unit (PSU), and calibrated survey weights), and derived variables used in the empirical analysis. Key harmonization decisions, documented in full in an accompanying technical appendix, include:

- **Weight recalibration.** EME V and EME VI are weighted using a recalibrated expansion factor (`fact_cal`, rebased to 2017) rather than the original survey weight, so that population totals are comparable across all four waves; EME VII and EME VIII use their native calibrated weight.
- **Missing-not-random treatment of AI variables.** Awareness and use of artificial intelligence tools were fielded only in EME VIII; these variables are coded as missing, not zero, in earlier waves.
- **Informality scale reversal.** EME VIII inverted the informality coding relative to EME VI/VII; we harmonize all waves to a common 1 = informal, 0 = formal convention.
- **Monetary deflation.** All monetary variables are deflated to real 2025 Chilean pesos using the INE consumer price index (base 2023=100), and both nominal and real series are retained.
- **Profit transformation and trimming.** Reported nominal monthly net profit is first deflated to 2025 Chilean pesos and then transformed using the inverse hyperbolic sine (IHS): $IHS(x) = \text{asinh}(x) = \ln(x + \sqrt{x^2 + 1})$. Thus, the construction order is $\text{profit}^{\text{nominal}} \rightarrow \text{profit}^{\text{real}, 2025} \rightarrow \text{ihs_profit}$. The outcome variable used throughout this paper’s profit regressions, `ihs_profit`, is *not* trimmed or winsorized. Separately, the harmonized data also provides `profit_trimmed_1_99`, which sets real monthly profit to missing (rather than capping it) outside the 1st/99th percentile of the *pooled*, not within-wave, distribution; this trimmed variable is available for robustness checks but is not used in any result reported here.
- **Variable-specific missing codes.** Missing-value codes (88/96/99/888/999, which vary by wave and variable) are recoded to `NA` individually rather than through a single global rule.

We report two analytic samples throughout. **Pooled V–VIII** ($N = 28,174$ active microentrepreneurs) spans the full 2017–2025 period but includes EME V, the wave with the weakest variable coverage — notably lacking a harmonizable informality measure, net profit, and hours worked. **Pooled VI–VIII** ($N = 20,682$), spanning 2019–2025, is our *recommended* and primary analytic sample, since all core outcome variables are consistently available across its three constituent waves. We validated the harmonized data against six officially published EME VIII indicators (active population, internet use, formality, AI awareness and use) and reproduce each exactly. A separate twelve-indicator cross-wave validation drawn from the official published summaries of EME VI, VII, and VIII — including three EME VIII benchmarks that overlap with the first exercise — is matched to within rounding error (twelve of twelve checks passed).

3.3 The care-work extension

For this paper we additionally harmonize the time-use module (Section L of the EME questionnaire), fielded in EME VI, VII, and VIII but not EME V, which contains separate items on unpaid domestic work and direct unpaid care. Because the domestic-work items are not fully comparable across waves, the consistently harmonized measure used in the analysis, `care_hrs_week`, is restricted to *direct unpaid-care time*. It is recorded through a single-day diary in EME VII/VIII and direct weekly-hours recall in EME VI. We convert the day-diary care measure to a weekly-hour equivalent by multiplying by seven and use each wave’s explicit non-participation flag to code structural zeros rather than treating survey skip patterns as missing data. Because the underlying elicitation method changes across waves (weekly recall in EME VI versus single-day diary $\times 7$ in EME VII/VIII), we treat cross-wave comparisons of *absolute* care-work hours with caution and focus inference on the male–female gap *within* each wave, which is measured consistently. The resulting variable, `care_hrs_week`, is merged into the pooled base by reconstructing each wave’s row-level identifier and verifying the merge against sex concordance (100% agreement for all three waves before the merged variable was used in any analysis).

3.4 Sample restriction and weighting

All analyses restrict to respondents identified as active microentrepreneurs (`microent_active = 1`; in EME VI–VIII this corresponds to `num_micro = 1` in the raw questionnaire coding). All descriptive statistics and regressions are weighted using the harmonized survey weight described above; standard errors are clustered at the primary sampling unit (PSU) level throughout. Because raw PSU identifiers are not unique across survey waves in the source

files (106 codes are shared between EME VII and EME VIII, and 8 between EME V and EME VI), we cluster on a wave-qualified PSU identifier (wave $\times$ PSU) rather than on the raw PSU code, so that observations from different waves are never pooled into the same cluster. We verified this correction changes no reported coefficient (clustering does not affect point estimates) and no reported standard error by more than the last decimal place shown, since the affected codes are a small share (under 3%) of all PSUs; we nonetheless apply the correction throughout for correctness.

3.5 Descriptive statistics by sex

Table 1: Weighted descriptive statistics by sex, pooled EME VI–VIII

Variable	Men	Women	Difference (M–F)
Real monthly profit (2025 CLP)	1,067,907	525,420	542,487
IHS profit	13.55	12.68	0.87
Age (years, band midpoint)	49.3	47.1	2.2
Education (1–6 scale)	3.25	3.29	–0.03
Informal (share)	0.522	0.597	–0.075
Internet use (share)	0.633	0.656	–0.023
Digital sales (share)	0.304	0.375	–0.070
Trained (share)	0.224	0.315	–0.091
Employer status (share)	0.164	0.086	0.078
Necessity entrepreneur (share) [†]	0.417	0.597	–0.180
Direct unpaid care (hrs/week)	2.51	5.60	–3.10
<i>N</i>	11,927	8,754	

[†]EME VII–VIII only (variable not fielded in EME VI). $N = 20,681$ (11,927 men + 8,754 women) is one observation fewer than the full pooled VI–VIII active sample ($N = 20,682$) used elsewhere in the paper because one record reports an implausible monthly profit of –\$90 million (against a sample mean near \$1 million). We treat this observation as a likely data-entry error and exclude it from this table and from all profit-related analyses.

Table 1 makes explicit the pattern discussed narratively throughout the paper: women have substantially lower profits and employer rates and higher rates of necessity entrepreneurship and informality, while matching or modestly exceeding men in the pooled available-wave means for education, internet use, digital sales, and business training, and reporting more than twice as many hours of direct unpaid care.

4 Empirical Strategy

4.1 Regression specifications

We estimate a common set of gender-gap regressions for several entrepreneurial outcomes. Let y_{it} denote the outcome observed for entrepreneur i in survey wave t ; the dependent variable therefore differs across specifications. For the performance analysis, y_{it} is monthly net business profit expressed in real 2025 Chilean pesos and transformed using the inverse hyperbolic sine (IHS). As described in Section 3, reported nominal monthly net profit is first deflated to 2025 prices and then transformed as

$$\text{IHS}(\text{profit}_{it}) = \ln\left(\text{profit}_{it} + \sqrt{\text{profit}_{it}^2 + 1}\right).$$

The IHS transformation behaves approximately like a logarithm for large positive values (specifically, $\text{IHS}(x) \approx \ln(2x)$) but remains defined at zero and for negative profits.¹ Because a consistently harmonized profit measure is available only from EME VI onward, profit regressions use the 2019, 2022, and 2025 waves. The analysis uses the untrimmed IHS-profit measure, apart from the single implausible profit record described in Section 3.5.

For the other stages of the entrepreneurial pathway, y_{it} is a binary indicator for necessity-driven entry, employer status, or receipt of business training, as defined from the harmonized survey responses described in Section 3. The availability of these outcomes differs across waves; in particular, necessity motivation is observed only in EME VII and EME VIII.

Our baseline specification uses the following linear index:

$$y_{it} = \alpha + \beta \cdot \text{Female}_i + \mathbf{X}_{it}'\boldsymbol{\gamma} + \delta_t + \varepsilon_{it}, \quad (1)$$

where Female_i equals one for women entrepreneurs, $\mathbf{X}_{it}$ contains the relevant observable characteristics — age band, education level, sector (six harmonized categories), macrozone and, where appropriate, internet use, informality, and training — and δ_t denotes survey-wave fixed effects. For IHS profit, Equation (1) is estimated by weighted least squares. For binary outcomes, the same right-hand-side index enters a weighted logit model. All specifications use the harmonized survey weights, with standard errors clustered at the wave-qualified primary sampling unit (PSU) level described in Section 3.4.

To test whether gaps have widened or narrowed over time, we interact Female_i with wave indicators, using EME VI (2019) as the omitted baseline wave:

¹We report differences in this transformed outcome as “IHS-profit points” throughout, rather than “log points,” since the two coincide only approximately.

$$y_{it} = \alpha + \beta \cdot \text{Female}_i + \sum_{t \neq 2019} \theta_t \cdot (\text{Female}_i \times \mathbb{I}[\text{Wave} = t]) + \mathbf{X}'_{it} \boldsymbol{\gamma} + \delta_t + \varepsilon_{it}. \quad (2)$$

Here β is the female effect in the 2019 baseline wave, and θ_{2022} and θ_{2025} measure how the gender gap in each later wave differs from its 2019 level; the widening hypothesis tested in Section 5 is $H_0 : \theta_{2025} = 0$. We estimate Equation (2) both without $\mathbf{X}'_{it} \boldsymbol{\gamma}$ (unadjusted) and with it (covariate-adjusted), to verify that the widening result is not an artifact of the covariate specification. To test whether digitalization or informality moderates the gender profit gap, we estimate variants of Equation (1) augmented with $\text{Female}_i \times \text{Internet}_i$ and $\text{Female}_i \times \text{Informal}_i$ interaction terms, respectively.

4.2 Mean decomposition (Oaxaca-Blinder)

To decompose the raw mean gender gap in `lhs_profit` into a portion explained by group differences in observable characteristics and a residual, we estimate separate weighted OLS models for men ($g = M$) and women ($g = F$), $y_i = \mathbf{X}'_i \boldsymbol{\beta}_g + \varepsilon_i$, and decompose

$$\bar{y}_M - \bar{y}_F = \underbrace{(\bar{\mathbf{X}}_M - \bar{\mathbf{X}}_F)' \boldsymbol{\beta}_M}_{\text{explained}} + \underbrace{\bar{\mathbf{X}}_F' (\boldsymbol{\beta}_M - \boldsymbol{\beta}_F)}_{\text{unexplained}} \quad (3)$$

using male coefficients as the reference (non-discriminatory) structure, following standard practice when one group is treated as disadvantaged (Oaxaca, 1973). For the mean decompositions reported in Table 4, $\mathbf{X}_i$ includes age band, education, sector, macrozone, internet use, informality, training, and `care_hrs_week`. The pooled row decomposes the pooled raw male–female mean gap and therefore does not add wave indicators to the Oaxaca covariate vector; wave is instead handled explicitly in the wave-specific rows and in the pooled regression and distributional analyses.

4.3 Unconditional quantile decomposition

The mean decomposition in Equation (3) can mask heterogeneity in the gap across the outcome distribution. We therefore apply the recentered-influence-function (RIF) approach of Firpo et al. (2009). For quantile τ , the RIF of y is

$$\text{RIF}(y; q_\tau) = q_\tau + \frac{\tau - \mathbb{I}[y \leq q_\tau]}{f_Y(q_\tau)} \quad (4)$$

where q_τ is the (weighted) τ -th quantile of the *pooled* outcome distribution and $f_Y(q_\tau)$ is the (weighted, kernel-estimated) pooled density at that reference quantile. We regress the

resulting pooled-reference RIF on $\mathbf{X}_i$ separately by sex and apply the same Oaxaca-style decomposition as in Equation (3) at $\tau \in \{0.10, 0.50, 0.90\}$. Thus, the RIF gaps reported below are pooled-reference unconditional-distribution contrasts, not literal differences between the male and female sex-specific quantiles. As a robustness check, we also compute those sex-specific weighted quantile gaps directly and bootstrap them by resampling PSUs within wave-by-stratum cells.

4.4 Matching decomposition (Ñopo)

Mean and quantile decompositions based on linear regression implicitly extrapolate outside the region of common covariate support between men and women. Following Ñopo (2008), we additionally construct exact covariate cells (age band $\times$ education $\times$ sector $\times$ macrozone $\times$ wave) and restrict comparison to cells populated by both sexes, decomposing the total gap as

$$\bar{y}_M - \bar{y}_F = \Delta_0 + \Delta_X + \Delta_M + \Delta_F \quad (5)$$

where Δ_0 is the gap within common support at matched-cell means (the unexplained component), Δ_X is the portion of the within-support gap attributable to differing covariate distributions, and Δ_M, Δ_F capture, respectively, men and women whose covariate cells have no counterpart in the opposite sex's support.

4.5 Mechanism-accounting exercise: unpaid care work

To assess whether unpaid care-work hours help account for the gender gaps documented in Section 5, we compare the estimated effect of Female_i in Equation (1) with and without `care_hrs_week` included as a covariate, across three outcomes: net profit, necessity-driven entry, and employer status. For the continuous profit outcome, estimated by linear WLS, we compare the Female_i coefficient directly, since linear coefficients are unaffected by the rescaling issues that complicate nonlinear models. For the two binary outcomes, estimated by weighted logistic regression, raw logit coefficients from nested specifications are *not* directly comparable: adding a covariate can shift the estimated scale of the latent index even absent any genuine mediating relationship (Karlson et al., 2012). We therefore report the survey-weighted *average marginal effect* (AME) of Female_i ,

$$\text{AME}_{\text{Female}} = \frac{\sum_{i=1}^N w_i \left[\widehat{\Pr}(y_i = 1 \mid \text{Female}_i = 1, \mathbf{X}_i) - \widehat{\Pr}(y_i = 1 \mid \text{Female}_i = 0, \mathbf{X}_i) \right]}{\sum_{i=1}^N w_i}, \quad (6)$$

computed on the probability scale from the fitted logit model, with and without `care_hrs_week`, for the binary outcomes, where w_i is the survey weight. For consistency of units within the resulting table, we report the effect of `care_hrs_week` itself on the binary outcomes as an AME as well — the weighted average change in predicted probability from a one-hour increase in weekly care work, holding Female_i and all other covariates at observed values — rather than as a raw logit coefficient. For the linear profit outcome, the WLS coefficient on `care_hrs_week` is already on a directly interpretable scale and is reported as-is. We additionally report the implied *percentage attenuation*, $100 \times (\hat{\beta}_{\text{without}} - \hat{\beta}_{\text{with}}) / \hat{\beta}_{\text{without}}$ (or the AME analogue), as a descriptive summary rather than a claim of any specific proportion causally mediated. A substantial attenuation, together with a statistically significant coefficient on care-work hours itself, is consistent with — though it does not establish — a mechanistic role for care work. We interpret the results with explicit caution, given the repeated cross-sectional, non-experimental design and the coarseness of an hours-based care-work measure, discussed further in Section 6.

4.6 Statistical inference for decomposition estimates

Point estimates from Equations (3)–(5) are accompanied by 95% confidence intervals constructed via a cluster (wave $\times$ PSU) bootstrap with 300–500 replications: in each replication we resample clusters with replacement (preserving all observations within a resampled cluster), recompute the full decomposition, and take the 2.5th and 97.5th percentiles of the resulting distribution as the confidence interval bounds. This respects the survey’s clustered design and avoids overstating precision from treating individual respondents as independent draws. For the “glass floor” claim specifically — that the gap is larger at low quantiles than at high quantiles — point-estimate comparison of separately bootstrapped quantile gaps is not sufficient to establish statistical significance of the *difference*; we therefore additionally compute $\text{Gap}_{p10}^{(b)} - \text{Gap}_{p90}^{(b)}$ within each bootstrap replication b and report the resulting percentile confidence interval directly, following standard practice for testing a difference between two dependent statistics.

4.7 Decomposing the change in the gap over time: composition, returns, and interaction

A natural further question is *why* the gap widened between 2019 and 2025: does it reflect a changing population of women entering microentrepreneurship (composition), worsening relative returns for a comparable population (returns), or both? To distinguish these channels without treating changes in the conventional Oaxaca–Blinder explained and unexplained

components as pure composition and return effects, we use a standard two-group, two-period composition/returns/interaction decomposition. Let $G_t \equiv \bar{\mathbf{X}}'_{Mt}\boldsymbol{\beta}_{Mt} - \bar{\mathbf{X}}'_{Ft}\boldsymbol{\beta}_{Ft}$ denote the period- t gap (period 0 = 2019, 1 = 2025), estimated from separate weighted regressions by sex *and* period on the common covariate set from Equation (3). Then

$$G_1 - G_0 = C + R + I, \quad (7)$$

where

$$\begin{aligned} C &= (\bar{\mathbf{X}}_{M1} - \bar{\mathbf{X}}_{M0})'\boldsymbol{\beta}_{M0} - (\bar{\mathbf{X}}_{F1} - \bar{\mathbf{X}}_{F0})'\boldsymbol{\beta}_{F0} && \text{(composition),} \\ R &= \bar{\mathbf{X}}'_{M0}(\boldsymbol{\beta}_{M1} - \boldsymbol{\beta}_{M0}) - \bar{\mathbf{X}}'_{F0}(\boldsymbol{\beta}_{F1} - \boldsymbol{\beta}_{F0}) && \text{(returns),} \\ I &= (\bar{\mathbf{X}}_{M1} - \bar{\mathbf{X}}_{M0})'(\boldsymbol{\beta}_{M1} - \boldsymbol{\beta}_{M0}) - (\bar{\mathbf{X}}_{F1} - \bar{\mathbf{X}}_{F0})'(\boldsymbol{\beta}_{F1} - \boldsymbol{\beta}_{F0}) && \text{(interaction).} \end{aligned}$$

where C is evaluated at period-0 (2019) returns for both sexes, R at period-0 characteristics for both sexes, and I is the remainder.

Unlike the male-reference Oaxaca-Blinder split, C isolates pure compositional change (evaluated at unchanging, period-0 returns for both sexes), R isolates pure return/coefficient change (evaluated at unchanging, period-0 characteristics for both sexes), and I captures the remainder attributable to composition and returns changing together. A large C relative to ΔGap would indicate the widening reflects a changing population of entrepreneurs (consistent with, though not proof of, a compositional artifact of repeated cross-sectional sampling); a large R indicates worsening relative returns for a comparable population, which a purely compositional story cannot explain.

5 Results

We present results along the five stages of the gendered entrepreneurial pathway motivated in the introduction: selection into entrepreneurship, structural gaps in sector and employer status, access to resources such as training and digital tools, the resulting performance gap, and the candidate unpaid-care-work mechanism.

5.1 Selection into entrepreneurship

Motivation for starting a business (necessity versus opportunity) was fielded only in EME VII and EME VIII. Women are substantially more likely to report necessity-driven motivation: the logistic coefficient on Female is 0.81 ($p < 0.001$), corresponding to roughly 2.2 times the

odds, controlling for age, education, sector, and macrozone. An interaction with wave shows no significant narrowing between 2022 and 2025 ($p = 0.17$): women remain substantially more likely to be pushed into entrepreneurship by necessity in the most recent wave as in the earlier one.

5.2 Structural gaps

Having shown that women enter microentrepreneurship differently, we next ask whether the businesses themselves differ structurally — in the sectors they operate in and in whether their owners become employers.

Table 2: Sectoral segregation (Duncan dissimilarity index) by wave, with bootstrap 95% CIs

	EME V (2017)	EME VI (2019)	EME VII (2022)	EME VIII (2025)
Dissimilarity index	0.364	0.367	0.379	0.357
95% CI	[0.336, 0.387]	[0.342, 0.390]	[0.356, 0.396]	[0.337, 0.379]

The sectoral dissimilarity index — the share of women (or men) who would need to switch sectors to equalize the sectoral distribution across sexes — shows no economically meaningful trend across the four waves: point estimates range narrowly from 0.357 to 0.379, and the bootstrap confidence intervals overlap substantially across all four waves (Table 2), providing no evidence of convergence in the type of business men and women operate over an eight-year period spanning a global pandemic and a technology-adoption wave.

Women are also substantially less likely to be employers rather than own-account workers (logistic coefficient -0.63 , $p < 0.001$, pooled). Interacting with wave (baseline 2017) shows the gap stable through 2019 and 2022, then widening significantly by 2025 (additional coefficient -0.34 , $p = 0.035$): fewer women are transitioning into employer status in the most recent wave than would be predicted by the earlier trend.

5.3 Resource access

Structural differences aside, the businesses women and men run also draw on similar — or, on some margins, more favorable — resources. Women are, if anything, *more* likely to have received business training than men (logistic coefficient $+0.52$, $p < 0.001$, EME VI–VIII pooled), with no statistically significant Female $\times$ wave shifts relative to 2019. Internet use is close to parity in the earlier waves and higher among women in the later waves; where digital-sales participation is comparably observed (EME VII–VIII), women’s participation

is at least as high as men’s (Table 1). Financing-related variables (loan requests, initial financing source) are flagged in our crosswalk as non-comparable across waves due to differing question wording and are therefore excluded from formal cross-wave modeling.

5.3.1 Interaction effects: does digitalization or informality moderate the gender gap?

Table 3: Interaction effects on the gender profit gap (pooled VI–VIII, WLS)

Interaction	Coefficient	SE	95% CI	<i>p</i> -value
Female $\times$ Internet use	−0.127	0.090	[−0.303, 0.049]	0.158
Female $\times$ Informal	−0.221	0.100	[−0.416, −0.026]	0.027
Female $\times$ EME VII (2022)	−0.056	0.108	[−0.267, 0.155]	0.604
Female $\times$ EME VIII (2025)	−0.303	0.119	[−0.536, −0.069]	0.011

Table 3 reports the interaction estimates referenced narratively elsewhere in the paper. Internet use does not significantly moderate the gender profit gap (point estimate negative but not distinguishable from zero), consistent with our claim that digital adoption does not close the performance gap even though women match or exceed men on adoption itself. Informality is significantly *more* costly for women than for men ($p = 0.027$): the profit penalty associated with operating informally is larger for women. The Female $\times$ wave rows formally establish the widening documented throughout the paper: the gender gap in 2025 is significantly larger than in 2019 ($p = 0.011$), while the 2022 gap is not significantly different from 2019.

5.4 Performance gap

None of the entry, structural, or resource-access patterns above translate into converging outcomes: we now turn to the profit gap itself, its evolution over time, and how much of it observable characteristics can explain.

Table 4: Oaxaca-Blinder decomposition of the gender profit gap (ihs_profit), by wave

Wave	Raw gap	Explained	Unexplained	% unexplained
EME VI (2019)	0.712	0.005	0.706	99.3%
EME VII (2022)	0.824	0.028	0.796	96.6%
EME VIII (2025)	1.059	0.108	0.951	89.8%
Pooled VI–VIII	0.868	0.056	0.811	93.5%

Note: Figures include `care_hrs_week` as a covariate (see Section 5.5). Explained and unexplained columns need not sum exactly to the raw gap due to rounding. Cluster (PSU) bootstrap 95% CI for the pooled VI–VIII row: raw gap [0.772, 0.969]; % unexplained [86.9%, 100.1%] ($B = 500$). The pooled row uses $N = 20,631$ complete cases for the listed covariates, including direct care hours.

The raw gender gap in IHS-profit grew by approximately 49% between 2019 and 2025 (Table 4). Across all mean Oaxaca-Blinder specifications in Table 4, approximately 90% of the gap is unexplained by sector, education, age, formality, internet use, training, or measured direct care (ranging from 89.8% in 2025 to 99.3% in 2019; the RIF median and Ñopo estimates reported below are in a similar but not identical range, 85–108%) — the observable characteristics most plausibly available to a policymaker seeking to close the gap through skills- or access-oriented interventions leave the overwhelming majority of the gap untouched.

5.4.1 Is the widening statistically significant?

Table 5: Wave-specific gender profit gap, with 95% confidence intervals

Wave	Gap coefficient	SE	95% CI	p -value
EME VI (2019)	−0.668	0.086	[−0.836, −0.500]	< 0.001
EME VII (2022)	−0.771	0.086	[−0.940, −0.602]	< 0.001
EME VIII (2025)	−0.973	0.112	[−1.192, −0.753]	< 0.001




Figure 1: Covariate-adjusted gender profit gap by wave, with 95% confidence intervals (Table 5). The gap is expressed as a positive quantity (men minus women) for ease of reading; error bars are the same 95% confidence intervals reported in Table 5.

Table 5 and Figure 1 report the covariate-adjusted gender gap estimated separately by wave. To formally test whether the 2025 gap differs from the 2019 baseline, we estimate a pooled specification with Female $\times$ Wave interactions (Table 3 above; this is Model 2, covariate-adjusted, in Table 6 below). The coefficient on Female $\times$ EME VIII(2025) is -0.303 (SE 0.119, 95% CI $[-0.536, -0.069]$, $p = 0.011$): we reject the null that the gap is unchanged between 2019 and 2025 at conventional significance levels. The 2022 coefficient is not significant ($p = 0.604$), indicating the widening is concentrated in the most recent wave rather than a steady trend across all three.

Table 6: Widening test: unadjusted vs. covariate-adjusted

Model	Wave	Coefficient	SE	95% CI	p -value
Model 1: unadjusted	EME VII (2022)	-0.113	0.112	$[-0.332, 0.107]$	0.314
Model 1: unadjusted	EME VIII (2025)	-0.347	0.123	$[-0.588, -0.105]$	0.005
Model 2: covariate-adjusted	EME VII (2022)	-0.056	0.108	$[-0.267, 0.155]$	0.604
Model 2: covariate-adjusted	EME VIII (2025)	-0.303	0.119	$[-0.536, -0.069]$	0.011

Note: Model 1 = Female $\times$ Wave only, no covariates. Model 2 = adds internet use, informality, training, age band, education, sector, and macrozone controls (Equation (1)).

Table 6 confirms the widening is not an artifact of the covariate adjustment: the *raw*, unadjusted 2025 gender differential is also significantly larger than 2019 ($p = 0.005$), with a

slightly larger point estimate (-0.347) than the adjusted version (-0.303), consistent with covariates absorbing a modest share of the raw widening. Both the raw and covariate-adjusted gender profit gaps therefore increased significantly between 2019 and 2025.

As an additional survey-design robustness check, we resample PSUs within wave-by-stratum cells while retaining the survey weights. The 95% bootstrap interval for Female $\times$ EME VIII remains entirely below zero in both the unadjusted specification ($[-0.585, -0.128]$) and the covariate-adjusted specification ($[-0.526, -0.097]$; $B = 500$). Thus the widening result does not depend on ignoring the survey stratification structure in the main cluster-robust sandwich calculation.

5.4.2 Decomposing the widening: composition, returns, or both?

Because EME is a repeated cross-section, the widening documented above could in principle reflect either (i) a changing population of women entering microentrepreneurship over time (“composition”), or (ii) worsening relative returns for a comparable population of women (“returns”), or both. A naive split of the *change* in the Oaxaca-Blinder explained and unexplained components is not a valid way to separate these — as shown in Section 4.7, $\Delta\text{Explained}$ is not pure composition and $\Delta\text{Unexplained}$ is not pure returns, since each contains a term where the other factor also changes. We instead apply the composition/returns/interaction (C+R+I) decomposition of Equation (7) directly.

Table 7: Intertemporal composition/returns/interaction (C+R+I) decomposition, 2019–2025

Component	Value	% of ΔG	95% CI (bootstrap)
ΔG ($G_{2025} - G_{2019}$)	0.347	100.0%	[0.113, 0.629]
C (composition, at 2019 returns)	0.113	32.6%	[0.007, 0.220]
R (returns, at 2019 characteristics)	0.356	102.4%	[0.083, 0.643]
I (interaction)	-0.122	-35.1%	$[-0.277, 0.025]$
<i>Symmetric allocation (interaction split equally)</i>			
$C^S = C + I/2$	0.052	15.1%	[–]
$R^S = R + I/2$	0.295	84.9%	[–]

Note: Cluster (wave $\times$ PSU) bootstrap, $B = 500$. Symmetric-allocation confidence intervals are not reported separately, as C^S and R^S are affine transformations of C , R , and I ; the substantive conclusion is unchanged under either allocation.

The C+R+I decomposition (Table 7) separates the widening into changes in characteristics, changes in coefficients, and their interaction. The pure returns component R alone

(102.4% of ΔG) exceeds the entire observed widening: had characteristics not changed at all, worsening relative returns for women would have produced an even larger increase in the gap than what we actually observe. Pure composition C also points toward widening (32.6%), but a sizable negative interaction term ($I = -35.1\%$) — reflecting that the characteristics which changed most and the coefficients which changed most partially offset each other across the two periods — brings the total back down to the observed $\Delta G = 0.347$. Bootstrap inference shows both C and R are individually statistically distinguishable from zero (their 95% confidence intervals exclude zero), while I is not (its interval spans zero); their confidence intervals also overlap one another, so while the *point estimate* for R substantially exceeds that for C , we do not treat the precise 32.6%/102.4%/−35.1% split as more exact than the data support.

As a robustness check on the reference-period dependence inherent in any threefold decomposition — C and R above are each evaluated holding the *other* factor at its 2019 level — we also report a symmetric allocation of the interaction term, $C^S = C + I/2$ and $R^S = R + I/2$. This gives $C^S = 0.053$ (15.1% of ΔG) and $R^S = 0.295$ (84.9%). The substantive conclusion is unchanged under this alternative allocation: changing returns account for 102.4% of the widening in the conventional threefold decomposition and 84.9% under the symmetric allocation. Thus the qualitative finding — changing returns dominate compositional change — does not depend on the particular convention used to assign the interaction term. Taken together with the bootstrap results, the paper’s central intertemporal claim is: changing returns are the dominant force behind the widening, not a changing population of entrepreneurs, though composition is not negligible, the two interact in a way not statistically distinguishable from zero, and the exact percentage split should be read as indicative rather than precise. We return to the limits of this claim, given the absence of true panel data, in Section 6.

5.4.3 Robustness: distributional and matching decompositions

Table 8: Unconditional quantile (RIF-OLS) decomposition, pooled VI–VIII

Quantile	Gap (ihs_profit)	% unexplained	Gap 95% CI	%unexpl. 95% CI
p10 (bottom earners)	1.126	96.3%	–	–
p50 (median)	0.898	84.7%	[0.836, 0.963]	[80.6%, 88.6%]
p90 (top earners)	0.503	92.2%	–	–

Note: Bootstrap CIs (cluster/wave×PSU, $B = 500$) reported for the median as a representative quantile.

Table 9: Direct bootstrap test of the “glass floor”: $\text{Gap}_{p10} - \text{Gap}_{p90}$

Statistic	Point estimate	95% CI	Excludes zero?
$\text{Gap}_{p10} - \text{Gap}_{p90}$	0.623	[0.491, 0.760]	Yes

Note: Cluster (wave $\times$ PSU) bootstrap, $B = 500$, computing $\text{Gap}_{p10}^{(b)} - \text{Gap}_{p90}^{(b)}$ within each replication (see Section 4).




Figure 2: The gender profit gap across the profit distribution (Table 8), with the direct bootstrap test of the $\text{Gap}_{p10} - \text{Gap}_{p90}$ difference (Table 9) annotated.

The gap is more than twice as large at the bottom of the earnings distribution as at the top (Table 8; Figure 2), using the pooled-reference RIF gap of Equation (4): 1.13 IHS-profit points at the 10th percentile versus 0.50 at the 90th. Table 9 confirms this difference is statistically significant, not merely a difference in point estimates that could be within noise of one another: the 95% confidence interval for $\text{Gap}_{p10} - \text{Gap}_{p90}$, computed directly within each bootstrap replication, is [0.491, 0.760] and excludes zero. Because these RIF gaps use pooled-reference quantiles, we also calculate the literal sex-specific weighted unconditional quantiles. That check yields male-minus-female gaps of 1.133 at p10, 0.829 at p50, and 0.642 at p90; a PSU-within-wave-by-stratum bootstrap gives a 95% CI of [0.274, 0.603] for the p10-minus-p90 difference. The distributional conclusion is therefore unchanged under the direct quantile definition. This is a **glass floor**, not a glass ceiling, in a formally testable sense: the mean-based result in Table 4 obscures a penalty that falls significantly harder on the most economically vulnerable women than on those already performing well.

Table 10: Ñopo (2008) matching decomposition, pooled VI–VIII

Component	Value (% of gap)
Δ_0 — unexplained, within common support	+0.936 (107.9%)
Δ_X — explained, covariate distribution	−0.090 (−10.3%)
Δ_M — men outside women’s support	+0.018 (2.1%)
Δ_F — women outside men’s support	+0.003 (0.4%)
Common support, unweighted sample share (men / women)	75.5% / 97.6%

Note: Cluster (wave×PSU) bootstrap 95% CI ($B = 300$): raw gap [0.773, 0.970]; Δ_0 share [98.7%, 115.9%].

Covariate overlap between men and women is substantial but not complete: 97.6% of female observations and 75.5% of male observations fall in cells shared with the opposite sex (Table 10), so roughly a quarter of male observations lie outside the region of common support with women. This is a caveat worth bearing in mind when interpreting the matching results, even though the extrapolation it could induce in a linear Oaxaca-Blinder model is evidently not driving our headline result. Within common support, the covariate-distribution component is slightly *negative* — women in shared cells have, if anything, marginally more favorable observable characteristics than men — while the unexplained component Δ_0 alone exceeds the entire raw gap, with a bootstrap confidence interval entirely above 98%.

Three independent decomposition methods, each with bootstrap confidence intervals excluding low values, converge on a common conclusion: the gender profit gap reflects predominantly different returns to observably similar characteristics among entrepreneurs with comparable measured profiles, rather than differences in those characteristics themselves. We deliberately avoid the stronger claim that this establishes the businesses themselves are otherwise identical. The covariates available in EME are a specific, finite set, and the decomposition cannot speak to unobserved differences in capital, collateral, prior experience, networks, or business quality, discussed further in Section 6.

5.5 Mechanism test: unpaid care work

The remaining question is whether unpaid care work — the mechanism motivated in Section 4.5 — can account for the entry, structural, and performance gaps documented above.

Table 11: Gender gap in weekly direct unpaid-care hours, by wave

Wave	Men	Women	Gap (hrs)	Ratio
EME VI (2019)	1.73	4.92	3.19	2.8×
EME VII (2022)	2.85	6.05	3.20	2.1×
EME VIII (2025)	3.00	5.83	2.83	1.9×

Women microentrepreneurs report 1.9 to 2.8 times as many weekly hours of direct unpaid care as men in every wave for which the time-use module is available (Table 11); the absolute gap is stable at roughly three hours per week across the period. As discussed in Section 3, the change in elicitation method across waves (weekly recall in EME VI versus day-diary $\times 7$ in EME VII/VIII) means the narrowing *ratio* should be interpreted cautiously, but the within-wave gender gap, measured consistently, is unambiguous. The underlying EME module contains items on both domestic tasks and direct care of household members. However, because the domestic-work measures are not fully comparable across waves, the pooled variable used here captures *direct care hours only*. It does not capture domestic work outside that harmonized care item, the fragmentation of the working day, the unpredictability of care demands, the intensity or mental load of caregiving, or whether care responsibilities coincide with business hours — an important scope limitation we return to below.

Table 12: Mechanism-accounting exercise: does care work help explain the gender gaps?

Outcome	Female effect (w/o care)	Female effect (w/ care)	% atten.	<code>care_hrs_week</code> effect (per +1 hr/wk)
Profit (lhs coef., WLS)	−0.800***	−0.782***	2.3%	−0.0066 WLS coef. ($p = 0.08$)
Necessity entry (AME, pp)	+19.4***	+18.2***	6.3%	+0.44 pp AME ($p < 0.001$)
Employer status (AME, pp)	−7.4***	−7.4***	0.4%	−0.01 pp AME ($p = 0.80$)

Note: For the binary outcomes (necessity entry, employer status), both the Female effect and the `care_hrs_week` effect are reported as average marginal effects (AME) in percentage points, so that all binary-outcome entries in this table are on a consistent probability-scale unit; raw logit coefficients are not directly comparable across nested specifications due to latent-scale rescaling (Karlson et al., 2012). For the linear profit model, both effects are WLS coefficients, which are unaffected by this issue. % attenuation is descriptive: the percentage reduction in the Female effect once `care_hrs_week` is added, not a formal proportion-mediated estimate from a causal mediation framework.

Table 12 reports the central mechanism-accounting exercise, using the average-marginal-effect construction in Equation (6). Including measured care hours attenuates the female average marginal effect on necessity-driven entry by approximately 6.3% (from 19.4 to 18.2 percentage points), and the care-work effect itself is large and highly significant on the probability scale: each additional weekly hour of unpaid care work is associated with a 0.44

percentage-point higher probability of necessity-driven entry ($p < 0.001$). This is consistent with a genuine, if partial, role for care burden in *why* women enter entrepreneurship. By contrast, adding care-work hours to the profit and employer-status models leaves the female effect almost entirely unchanged (2.3% and 0.4% attenuation respectively), and the care-work effect is marginal ($p = 0.08$) or null ($p = 0.80$). In the Oaxaca-Blinder decomposition, care-work hours alone contribute only 1.6% of the pooled profit gap, despite the large raw difference in hours (men average 2.5 hours/week, women 5.6).

We interpret this pattern cautiously. The measured hours-based care-work variable does not explain the profit or employer-status gaps, while it does help account for the necessity-entry gap. This is a meaningful and, we think, informative result — it is inconsistent with the simplest version of the resource-constraint story for *performance*, while supporting it for *entry* — but we stop short of claiming it rules out care work as a mechanism for performance more broadly. As noted above, our measure captures weekly hours only. It does not capture the timing, flexibility, predictability, or intensity of care demands, any of which could matter for business performance in ways an hours count cannot detect. Nor can a repeated cross-section rule out reverse causality, since entrepreneurs who select into flexible, necessity-driven, home-based business models may subsequently have more capacity to absorb care work rather than care burden being purely a prior constraint.

What the exercise does establish is that the majority of the profit and employer-status gaps documented in this paper is attributable to neither observable business or human-capital characteristics (Section 5) nor to this specific, theoretically well-motivated time-constraint measure. This leaves differential access to capital, markets, networks, or discriminatory treatment as plausible — though not directly identified — candidate explanations for the residual, a distinction we return to in Section 6.

5.6 Synthesis

Taken together, the five dimensions examined above summarize a consistent pattern. Women entering Chilean microentrepreneurship between 2017 and 2025 are more likely to be necessity-driven, report two to three times as many measured hours of direct unpaid care as men, and are no less — indeed somewhat more — engaged with digital tools and formal training than men. None of this translates into narrowing gaps: sectoral segregation shows no meaningful trend, the employer-status gap has widened significantly, and the profit gap has grown by roughly half, an increase confirmed statistically significant by a Female $\times$ Wave test, with roughly 85–108% of the level gap unexplained by observable characteristics across three independent decomposition methods and concentrated most heavily among the lowest earners.

6 Concluding Remarks

Using a newly harmonized and validated four-wave pooled repeated cross-section of Chile’s Microentrepreneurship Survey (2017–2025), this paper documents a gender gap in microentrepreneurial outcomes that is large, persistent, and concentrated among the lowest earners. The employer-status gap widens significantly over 2017–2025, while the profit gap — consistently measurable from 2019 onward — increases significantly between 2019 and 2025. Both unadjusted and covariate-adjusted Female $\times$ Wave specifications confirm that the 2025 profit gap is significantly larger than its 2019 baseline. A composition/returns/interaction decomposition further indicates that changing relative returns, rather than observable compositional change alone, are the dominant force behind this widening, although the precise allocation between components should not be interpreted mechanically.

The distributional evidence sharpens this result. The gender profit gap is substantially larger at the bottom than at the top of the profit distribution, and a direct bootstrap test confirms that the difference between the 10th- and 90th-percentile gaps is statistically significant. This “glass floor” pattern suggests that gender disparities are particularly pronounced among the most economically vulnerable microentrepreneurs rather than being concentrated among high-performing firms. At the same time, women in the EME are not less educated, less digitally engaged, or less likely to receive business training than men; on several of these margins they match or exceed them. Observable differences in sector, human capital, formality, digital adoption, and training therefore account for only a limited share of the observed performance gap.

The unpaid care-work evidence adds a further dimension to this puzzle. Women report substantially more directly measured unpaid care time than men in every wave for which the time-use module is available. Measured care hours are strongly associated with necessity-driven entry into entrepreneurship, but including them changes the estimated gender profit and employer-status gaps very little. Thus, the hours-based care measure available in EME appears more relevant to the circumstances under which women enter entrepreneurship than to the performance gap conditional on operating a business. This result contrasts with evidence from other Latin American settings in which household responsibilities explain a substantially larger share of gender differences in business profitability, underscoring the importance of institutional context.

Three limitations qualify these findings. First, EME is a repeated cross-section rather than a panel. Although the widening of the gender profit gap is statistically significant and the composition/returns/interaction decomposition indicates that changing relative returns are quantitatively more important than observable compositional change, we cannot rule

out changes over time in unobserved characteristics of the entrepreneurs entering the sample. Longitudinal or synthetic-cohort evidence would be required to distinguish these mechanisms more sharply.

Second, the care-work analysis is associational rather than causal. The harmonized EME measure captures weekly hours of direct unpaid care, not a fully comparable total of domestic plus care work, and it does not capture timing, predictability, fragmentation, or intensity; the elicitation method also differs across survey waves. Reverse causality is also possible if entrepreneurs who select into more flexible forms of self-employment subsequently absorb greater care responsibilities. We therefore interpret the results as evidence about the explanatory power of the measured hours-based constraint, not as a causal mediation estimate.

Third, the large “unexplained” component in the decomposition exercises remains a residual. It may reflect unobserved differences in wealth, collateral, prior experience, household resources, business networks, business quality, measurement error, or discriminatory treatment. Accordingly, the analysis establishes that the characteristics observed in EME explain only a limited fraction of the gender performance gap; it does not identify any single mechanism as the cause of the remaining differential.

These findings have implications for entrepreneurship policy. Policies directed exclusively at closing presumed deficits in digital skills, business training, or other observable capabilities may have limited effects on gender performance gaps when women already match or exceed men on these margins. In this respect, the results reinforce the concern raised by [Foss et al. \(2019\)](#) that women’s entrepreneurship policy can place excessive emphasis on correcting individual-level deficits while giving less attention to the broader conditions under which entrepreneurs convert resources and capabilities into economic returns. The concentration of the gap among lower-performing entrepreneurs further suggests that policies aimed at the most economically vulnerable businesses may be particularly relevant.

The care-work results point to a related distinction. Policies that reduce care constraints — through childcare provision or greater flexibility in the organization of paid and unpaid work — may expand women’s entrepreneurial choice set and reduce necessity-driven entry, even if they are insufficient on their own to close performance gaps among existing businesses. Understanding those performance gaps requires direct evidence on mechanisms that EME cannot observe consistently, including access to capital, customers, markets, and business networks.

Future research can therefore build on the patterns documented here by combining survey evidence with administrative, credit, or transaction data capable of observing these channels directly. Panel or synthetic-cohort designs could further distinguish changes in the

composition of women entering microentrepreneurship from persistent changes in the returns earned by otherwise comparable entrepreneurs, while richer time-use information could identify whether the timing and intensity of care responsibilities matter in ways that aggregate weekly hours do not capture.

Taken together, the results point to a broader conclusion. The gender gap in microentrepreneurial performance in Chile cannot be understood simply as a deficit in women's observable entrepreneurial capabilities. Women have reached or surpassed men on several measured dimensions conventionally associated with business performance, yet the profit gap has widened and is largest among those at the bottom of the distribution. The central challenge is therefore not only whether women obtain entrepreneurial resources and capabilities, but also whether they are able to convert them into comparable economic returns.

References

- Agarwal, B. (2026). Development as equality: A gender lens on progress and its hidden barriers. *World Development*, 200, 107280. <https://doi.org/10.1016/j.worlddev.2025.107280>.
- Arráiz, I. (2018). Time to share the load: Gender differences in household responsibilities and business profitability. *Small Business Economics*, 51(1), 57–84. <https://doi.org/10.1007/s11187-017-9925-z>.
- Bernat, L.F., Lambardi, G., & Palacios, P. (2017). Determinants of the entrepreneurial gender gap in Latin America. *Small Business Economics*, 48(3), 727–752.
- Berniell, I., Berniell, L., de la Mata, D., Edo, M., & Marchionni, M. (2021). Gender gaps in labor informality: The motherhood effect. *Journal of Development Economics*, 150, 102599. <https://doi.org/10.1016/j.jdeveco.2020.102599>.
- Blinder, A.S. (1973). Wage discrimination: Reduced form and structural estimates. *Journal of Human Resources*, 8(4), 436–455. <https://doi.org/10.2307/144855>.
- Brush, C., Edelman, L.F., Manolova, T., & Welter, F. (2019). A gendered look at entrepreneurship ecosystems. *Small Business Economics*, 53(2), 393–408. <https://doi.org/10.1007/s11187-018-9992-9>.
- Chant, S. (1994). Women, work and household survival strategies in Mexico, 1982–1992: Past trends, current tendencies and future research. *Bulletin of Latin American Research*, 13(2), 203–233.

- Cirera, X., Cruz, M., Martins-Neto, A., Lee, K.M., & Nogueira, C. (2026). The role of technology in reducing the gender gap in productivity. *World Development*, 197, 107197. <https://doi.org/10.1016/j.worlddev.2025.107197>.
- de la Rica, S., Dolado, J.J., & Llorens, V. (2008). Ceilings or floors? Gender wage gaps by education in Spain. *Journal of Population Economics*, 21(3), 751–776. <https://doi.org/10.1007/s00148-006-0128-1>.
- Elson, D. (1999). Labor markets as gendered institutions: Equality, efficiency and empowerment issues. *World Development*, 27(3), 611–627.
- Fairlie, R.W. (2003). An extension of the Blinder-Oaxaca decomposition technique to logit and probit models. Economic Growth Center Discussion Paper No. 873, Yale University.
- Firpo, S., Fortin, N.M., & Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3), 953–973.
- Foss, L., Henry, C., Ahl, H., & Mikalsen, G.H. (2019). Women’s entrepreneurship policy research: A 30-year review of the evidence. *Small Business Economics*, 53(2), 409–429. <https://doi.org/10.1007/s11187-018-9993-8>.
- Global Entrepreneurship Monitor (2023). *Women’s Entrepreneurship Report 2023*. Babson College / GEM Consortium.
- Guzmán, J., & Kacperczyk, A. (2019). Gender gap in entrepreneurship. *Research Policy*, 48(7), 1666–1680.
- Hazarika, B. (2020). Gender income gap in rural informal micro-enterprises: An unconditional quantile decomposition approach in the handloom industry. *Eurasian Business Review*, 10(3), 441–473.
- Hoang, N., Nahm, D., & Dobbie, M. (2021). Innovation, gender, and labour productivity: Small and medium enterprises in Vietnam. *World Development*, 146, 105619. <https://doi.org/10.1016/j.worlddev.2021.105619>.
- Hundley, G. (2001). Why women earn less than men in self-employment. *Journal of Labor Research*, 22(4), 817–829. <https://doi.org/10.1007/s12122-001-1054-3>.
- International Labour Organization (2025). *Labour Overview 2025: Latin America and the Caribbean*. ILO Regional Office for Latin America and the Caribbean.

- Islam, A.M., & Amin, M. (2023). The gender labor productivity gap across informal firms. *World Development*, 167, 106229. <https://doi.org/10.1016/j.worlddev.2023.106229>.
- Jann, B. (2008). The Blinder-Oaxaca decomposition for linear regression models. *Stata Journal*, 8(4), 453–479.
- Karlson, K.B., Holm, A., & Breen, R. (2012). Comparing regression coefficients between same-sample nested models using logit and probit: A new method. *Sociological Methodology*, 42(1), 286–313.
- Kitagawa, E.M. (1955). Components of a difference between two rates. *Journal of the American Statistical Association*, 50(272), 1168–1194.
- Lauterbach, R., & Durán-Sandoval, D. (2026). Women at the helm: Gender gaps in business income among micro-entrepreneurs in Chile. *Eurasian Business Review*, 16(1), 25–50. <https://doi.org/10.1007/s40821-026-00343-z>.
- Lechmann, D.S.J., & Schnabel, C. (2012). Why is there a gender earnings gap in self-employment? A decomposition analysis with German data. *IZA Journal of European Labor Studies*, 1(1), Article 6. <https://doi.org/10.1186/2193-9012-1-6>.
- Lo Bue, M.C., Le, T.T.N., Santos Silva, M., & Sen, K. (2022). Gender and vulnerable employment in the developing world: Evidence from global microdata. *World Development*, 159, 106010. <https://doi.org/10.1016/j.worlddev.2022.106010>.
- Mickiewicz, T., & Nguyen, B. (2025). “In the company of cheerful ladies”: Whether female entrepreneurs are more productive? *Small Business Economics*, 64, 837–861. <https://doi.org/10.1007/s11187-024-00929-9>.
- Moeini Gharagozloo, M.M., Forghani Bajestani, M., Moeini Gharagozloo, A., Amini Sedeh, A., & Askarzadeh, F. (2023). The role of digitalization in decreasing gender gap in opportunity driven entrepreneurship. *Information Technology for Development*, 29(4), 645–664. <https://doi.org/10.1080/02681102.2023.2237470>.
- Ñopo, H. (2008). Matching as a tool to decompose wage gaps. *Review of Economics and Statistics*, 90(2), 290–299. <https://doi.org/10.1162/rest.90.2.290>.
- Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 14(3), 693–709.

- Rietveld, C.A., & Patel, P.C. (2022). Gender inequality and the entrepreneurial gender gap: Evidence from 97 countries (2006–2017). *Journal of Evolutionary Economics*, 32, 1205–1229. <https://doi.org/10.1007/s00191-022-00780-9>.
- Silupu, B., Usero, B., & Montoro, M.A. (2024). Why formalize? Examining the context of Peruvian women microentrepreneurs. *Journal of Enterprising Communities: People and Places in the Global Economy*, 18(6), 1364–1383. <https://doi.org/10.1108/JEC-07-2023-0124>.
- Simon, J.K., & Way, M.M. (2016). Why the gap? Determinants of self-employment earnings differentials for men and women millennials in the US. *Journal of Family and Economic Issues*, 37(2), 297–312.
- Strawser, J.A., Hechavarría, D.M., & Passerini, K. (2021). Gender and entrepreneurship: Research frameworks, barriers and opportunities for women entrepreneurship worldwide. *Journal of Small Business Management*, 59(S1), S1–S15. <https://doi.org/10.1080/00472778.2021.1965615>.
- Trumbic, T., & O’Connor, D. (2025). Digital inclusion: How to unlock the \$5 trillion opportunity for women entrepreneurs. *World Economic Forum*.
- Vial, V., & Hanoteau, J. (2015). Returns to micro-entrepreneurship in an emerging economy: A quantile study of entrepreneurial Indonesian households’ welfare. *World Development*, 74, 142–157. <https://doi.org/10.1016/j.worlddev.2015.04.008>.
- Vossenbergh, S. (2013). Women entrepreneurship promotion in developing countries: What explains the gender gap in entrepreneurship and how to close it. Maastricht School of Management Working Paper Series, 8(1), 1–27.
- Wheat, C., & Mac, C. (2024). Does necessity entrepreneurship explain the gender revenue gap? *JPMorganChase Institute*. <https://www.jpmorganchase.com/institute/all-topics/business-growth-and-entrepreneurship/does-necessity-entrepreneurship-explain-the-gender-revenue-gap>.