



ISEG – Lisbon School of Economics Management

Master's in Applied Econometrics and Forecasting

Gender Activity Rate Gaps in Morocco

A Spatial Econometric Perspective

Dissertation submitted in partial fulfillment of the requirements for the
degree of Master in Applied Econometrics and Forecasting

Majda El Oumami

Supervisors

Luís Silveira Santos

Esmeralda Arranhado

June 2026

Contents

Introduction	1
1 Literature Review	4
1.1 Theoretical Frameworks and Their Critiques	4
1.2 Empirical Evidence Across Countries	6
1.3 The Gender Gap in Morocco	7
1.4 Gaps in the Literature	8
2 Data Description	10
2.1 Data Sources and Geographical Framework	10
2.2 Variable Construction	11
2.2.1 Outcome Variable	11
2.2.2 Explanatory Variables	11
2.3 Exploratory Analysis	14
2.3.1 Descriptive Statistics	14
2.3.2 Urban and Rural Comparisons	15
2.4 Spatial Distribution	17
2.4.1 Geographic Patterns	17
2.4.2 Global Spatial Autocorrelation	19
2.4.3 Local Spatial Autocorrelation	21
3 Methodology	23
3.1 OLS Baseline	23
3.2 Fractional Regression Model	24
3.3 Spatial Weights Matrix	25
3.4 Spatially Lagged X Model	26
3.5 Fractional Logit with Spatially Lagged Regressors	27
3.6 Model Selection and Comparison	28
4 Results	30
4.1 Model Selection	30

4.2	Alternative Models: Direct and Indirect Effects	32
4.3	Comparison with Previous Empirical Results	34
5	Discussion and Conclusion	36
5.1	Findings, Theory, and the Literature	36
5.2	Policy Implications	37
5.3	Summary and Limitations	37

Abstract

This study provides the first spatial econometric analysis of the gender activity rate gap at the commune level in a Middle East and North Africa context, using data from Morocco's 2024 General Census of Population and Housing across 1,499 communes. Despite sustained increases in female educational attainment, Morocco's female labour force participation rate has declined over the past two decades, presenting a fundamental challenge to standard development theory. Using a Spatially Lagged X model with heteroskedasticity and autocorrelation consistent standard errors as the primary specification, and a fractional logit SLX as a robustness check, the study documents significant spatial dependence in the activity gap and identifies both direct and indirect determinants of the gap across communes. Secondary and higher education shares are the strongest negative direct determinants, while household size and married population share exert significant positive direct effects consistent with household bargaining models. Among spillover effects, working-age population share generates a significant positive indirect effect, and higher education a significant negative indirect effect, pointing to the geographic diffusion of both agricultural male employment and human capital externalities across commune boundaries. The near-perfect similarity of average partial effects across linear and nonlinear specifications confirms the adequacy of the linear model. The findings highlight the importance of subnational spatial analysis for understanding gender inequality in developing economies and carry direct implications for the geographic targeting of education and labour market policies in Morocco.

Keywords: gender activity rate gap, spatial econometrics, fractional regression, Morocco, commune-level analysis

JEL Classification: J16, J21, O15, R12, C21

Introduction

Gender inequality in labour markets remains one of the most persistent structural challenges in the developing world. Globally, women's labour force participation rates lie behind men's by nearly 25 percentage points, with female participation at 47.4% against male participation at 72.3%, a gap that has barely narrowed over the past three decades despite sustained economic growth and significant advances in female educational attainment ([International Labour Organization, 2023](#)). Among women aged 25–54, the gap widens further when children are present. Mothers of children under six participate at only 53.1%, compared to 95.7% for fathers in the same age group ([International Labour Organization, 2024](#)). Moreover 2.7 billion women worldwide remain legally restricted from having the same choice of jobs as men, and in 2023, more than one-third of 190 assessed economies still had laws constraining women's decision to work ([UN Women, 2023](#)). The consequences extend well beyond individual welfare. The World Bank estimates that closing the gender gap in employment and entrepreneurship could increase global GDP by more than 20%, while McKinsey Global Institute estimates that advancing women's equality could add up to \$12 trillion to annual global GDP ([McKinsey Global Institute, 2015](#); [World Bank, 2024](#)). In response, international frameworks from the United Nations Sustainable Development Goals (SDG 5 on gender equality and SDG 8 on decent work) to the Convention on the Elimination of All Forms of Discrimination Against Women (CEDAW) have placed gender-inclusive labour markets at the centre of development policy. In the Middle East and North Africa (MENA) region, female participation rates average below 20%, the lowest of any world region ([OECD, 2020](#)), and have remained broadly stagnant even as education gaps have closed. Critically, closing gender employment gaps in MENA could raise GDP per capita by as much as 49% in the long run, more than double the global average gain, making the region where the cost of inaction is highest ([World Bank, 2023](#)).

Morocco presents this challenge in its sharpest form. Despite two decades of sustained economic growth, significant expansion of female educational attainment, and structural transformation away from agriculture, female labour force participation declined from approximately 24% in 2000 to around 21% in 2018 ([Lopez-Acevedo et al., 2021](#)), a trend that has continued into the 2020s with ILO modelled estimates placing

the rate at 19.5% in 2024, the lowest recorded value in the country's history ([TheGlobalEconomy.com, 2024](#)), while male participation remained stable above 70% throughout the same period ([Lopez-Acevedo et al., 2021](#)). Morocco's female participation rate sits below the MENA regional average and substantially below neighbouring countries such as Tunisia (26%) and Lebanon (29%), despite Morocco having achieved higher levels of female educational attainment than many of its peers ([Think Global Health, 2023](#)). Approximately 60% of employed Moroccan women work in the agricultural sector, a far higher share than in comparable MENA economies, reflecting the absence of compatible formal employment in the service sector ([Center for Macroeconomics and Development, 2023](#)). The resulting gender activity rate gap of approximately 50 percentage points places Morocco among the countries with the widest gender labour market disparities globally. This divergence is particularly striking given Morocco's stated commitments. The country ratified CEDAW in 1993, adopted a Family Code reform in 2004 expanding women's legal rights, and enshrined gender equality in its 2011 Constitution ([World Bank, 2023](#)). Understanding the determinants of this gap is therefore not merely an academic exercise. It carries direct implications for the design of labour market and education policies in a country where female economic inclusion remains a central development challenge and where the distance between legal commitments and labour market outcomes demands explanation.

This study draws on Morocco's 2024 General Census of Population and Housing, conducted by the Haut-Commissariat au Plan (HCP) and covering the entire Moroccan population. Morocco's administrative geography operates at three hierarchical levels: regions, provinces and prefectures, and communes, the smallest administrative unit at which labour markets, social norms, and household decisions actually operate. The final analytical sample comprises 1,499 communes across Morocco's 12 administrative regions, of which 204 are urban (13.6%) and 1,295 are rural (86.4%). Analysing the gender activity rate gap at this granular level reveals that the national aggregate of approximately 50 percentage points conceals a distribution ranging from 19.5 to 100 percentage points across communes, with systematic regional clustering that national-level analysis would be structurally incapable of capturing.

Existing research on this paradox has relied predominantly on decomposition methods that attribute observed gaps to differences in endowments or returns. This approach is not applicable here: the HCP census provides commune-level aggregate tabulations rather than individual microdata, making decomposition methods inapplicable. The methodological implications of this data structure are discussed in detail in Section 1.4.

This study addresses these gaps through three contributions. First, it constructs a commune-level dataset from Morocco's 2024 General Census of Population and Housing, providing the most granular and comprehensive evidence base available

for this analysis. Second, it applies spatial econometric methods, specifically a Spatially Lagged X model and a fractional logit counterpart, to jointly estimate the direct determinants of the gap within communes and the spillover effects operating across commune boundaries. Third, it provides formal evidence on the adequacy of the linear specification through the similarity of average partial effects across linear and nonlinear models, establishing that the substantive conclusions are robust to functional form assumptions.

The main findings are as follows. The gender activity rate gap exhibits significant positive spatial autocorrelation (Moran's $I = 0.351$, $p < 0.001$), confirming that geographic clustering of gender inequality is a structural feature of Morocco's commune landscape. Secondary and higher education shares are the strongest negative determinants of the gap, while household size and married population share exert positive effects consistent with household bargaining theory. Spatial spillovers are significant for working-age population share, higher education, and the urban indicator, indicating that neighbouring commune characteristics independently shape local gender gaps beyond the direct effects of own-commune covariates.

The remainder of the study is structured as follows. Chapter 1 reviews the theoretical and empirical literature on gender activity rate gaps and motivates the spatial approach. Chapter 2 describes the data and presents exploratory spatial analysis. Chapter 3 sets out the empirical methodology. Chapter 4 presents the results. Chapter 5 discusses the findings in light of the theoretical frameworks and concludes with policy implications and directions for future research.

Chapter 1

Literature Review

Understanding the gender activity rate gap remains a central challenge for policy makers and researchers in developing economies. Morocco presents a particularly striking paradox. Despite sustained increases in female educational attainment and significant structural transformation over the past two decades, female participation in the labour force has stagnated and in some periods declined ([Lopez-Acevedo et al., 2021](#)). This divergence from standard development theory raises fundamental questions about the mechanisms linking economic growth, structural change, and gender inclusion.

The theoretical literature on female labour force participation draws on three interconnected frameworks: the U-shaped development hypothesis, standard labour supply theory, and collective household bargaining models. Each offers partial explanations for observed gender gaps, yet none fully accounts for Morocco's trajectory. The following sections review each framework, its empirical support, and its limitations, before turning to the empirical evidence across countries and the specific case of Morocco.

1.1 Theoretical Frameworks and Their Critiques

The U-Shaped Development Hypothesis

The U-shaped hypothesis, articulated by [Goldin \(1995\)](#), posits a nonlinear relationship between economic development and female participation. In the early stages of industrialisation, participation declines as agricultural economies transition toward manufacturing. This initial decline reflects several reinforcing mechanisms. Factory work is physically demanding and socially stigmatised for women, rising male wages generate household income effects that reduce women's need to work, and urban labour markets initially offer few occupations compatible with women's domestic responsibilities. As development advances further, participation recovers and rises as education expands, fertility declines, and service sector employment creates opportunities

more compatible with female participation.

Empirical support for this hypothesis is, however, mixed. [Klasen and Gaddis \(2014\)](#) demonstrate that evidence for the U-shape is highly sensitive to data sources and estimation choices, raising doubts about its structural universality across development contexts. At the microeconomic level, [Klasen \(2019\)](#) documents a U-shaped relationship between education and participation within developing countries, with intermediate educated women often displaying lower participation rates than both uneducated and highly educated women, a pattern in line with the stigma and reservation wage mechanisms operating during transitional stages. A further limitation of the U-shaped framework is its treatment of space. It operates at the national or cross-country level and cannot account for the substantial within-country geographic heterogeneity in participation that characterises countries like Morocco. The mechanisms it identifies, namely sectoral transition, stigma, and income effects, operate locally and may diffuse spatially across commune boundaries in ways the aggregate framework cannot capture.

Standard labour supply Theory

On the other hand, standard labour supply theory frames participation decisions as reflecting the balance between substitution and income effects. Rising market wages increase the opportunity cost of non-participation, drawing women into the labour force through substitution effects. Conversely, rising household income may reduce participation through income effects as families can afford to have women withdraw from the labour market. This framework connects directly to the U-shape mechanism. The initial declining phase of the U can be understood as income effects dominating substitution effects during early industrialisation, when male wages rise rapidly but female-compatible jobs remain scarce. Cross-country evidence from [Klasen \(2019\)](#) documents substantial heterogeneity in how these effects operate across contexts, suggesting that local labour market conditions, wage structures, and social norms mediate theoretical predictions.

A fundamental limitation of the canonical labour supply framework is its implicit spatial neutrality. It treats all locations as equivalent, ignoring the geographic clustering of labour demand and the spatial diffusion of social norms that shape female participation decisions in practice. In a country like Morocco, where communes range from dense urban agglomerations to isolated rural settlements, the assumption that all locations face equivalent labour market conditions is particularly untenable.

Collective Household Bargaining Models

Finally, collective household bargaining models extend the analysis beyond individual optimisation by emphasising intrahousehold power dynamics. Fertility and participation decisions are jointly determined through bargaining processes in which women's outside options, shaped by education, inheritance rights, and local labour market conditions, influence outcomes. Instrumental variable approaches robustly identify negative causal effects of additional children on participation, though magnitudes vary substantially across contexts ([Klasen, 2019](#)). For Morocco, where household size averages 4.23 persons and married women represent over 60% of the working age female population across communes, bargaining and fertility dynamics are likely central determinants of the observed gender gap.

In Middle East and North Africa (MENA)¹ contexts this bargaining framework must be extended to account for identity norms and social expectations surrounding female employment. [Hayo and Caris \(2013\)](#) show that attitudes toward gender roles are a significant and independent predictor of female labour force participation across Middle Eastern and North African countries, operating through women's own preferences as well as through household and community pressure. This identity channel helps explain why standard income and substitution effects may be weaker or reversed in Morocco relative to predictions from canonical models.

A shared limitation of all three frameworks is their implicit assumption of spatial homogeneity. They model participation decisions as if geographic context were irrelevant, or at best captured by a rural/urban indicator. None accounts for the spatial diffusion of norms, the geographic clustering of labour demand, or the spillover effects through which conditions in neighbouring communes influence local participation decisions. This spatial blindness is the central theoretical gap that motivates the spatial econometric approach adopted in this study.

1.2 Empirical Evidence Across Countries

A substantial empirical literature documents the determinants of female labour force participation across developing and middle-income countries, providing the comparative context within which Morocco's trajectory can be evaluated.

Cross-country studies systematically identify education, fertility, and structural transformation as the primary drivers of female participation. [Tsani et al. \(2013\)](#) examine female labour force participation across MENA countries and find that education

¹The MENA region comprises the countries of the Middle East and North Africa, including Morocco, Algeria, Tunisia, Egypt, Jordan, Lebanon, and the Gulf states, among others. It is characterised by the lowest female labour force participation rates of any world region ([World Bank, 2023](#)).

expansion alone is insufficient to raise participation in the presence of strong identity norms and limited compatible employment opportunities, a finding directly relevant to the Moroccan paradox. [International Labour Organization \(2023\)](#) document that the global gender participation gap has narrowed only marginally over three decades despite large gains in female education, pointing to the primacy of demand-side and normative constraints.

At the subnational level, a growing European literature employs spatial econometric methods to analyse gender employment gaps. [Falk and Leoni \(2010\)](#) combine spatial lag models with geographically weighted regression across Austrian districts and find that the determinants of female labour force participation vary substantially across space, with education and childcare availability generating heterogeneous effects across regions. [Noback et al. \(2013\)](#) document significant spatial spillovers in female employment across Dutch municipalities, showing that local labour market conditions in neighbouring areas independently shape local female participation beyond own-municipality characteristics. [Chica-Olmo et al. \(2021\)](#) identify systematic spatial clustering of gender employment gaps across Andalusian municipalities, demonstrating that non-spatial regressions yield biased coefficients and underestimate spillover effects. Taken together, these studies establish that spatial methods are not merely a technical refinement but a substantive necessity when analysing gender gaps at the subnational level. No comparable analysis exists for Morocco or the broader MENA region.

1.3 The Gender Gap in Morocco

Morocco's labour market trajectory presents a fundamental challenge to standard development theory. Despite sustained economic growth, significant expansion of female educational attainment, and structural transformation away from agriculture, female labour force participation has not followed the predicted upward path. Instead, participation declined from approximately 24% in 2000 to around 21% in 2018, while male participation remained stable above 70% throughout the same period ([Lopez-Acevedo et al., 2021](#)). This divergence, namely rising education alongside declining participation, constitutes what the literature has termed the Moroccan paradox and motivates the subnational spatial investigation undertaken in this study.

The education-participation disconnect is particularly striking. While tertiary education increases the probability of female participation, its marginal impact has weakened considerably over time ([Lopez-Acevedo et al., 2021](#)). Educated women increasingly face job-matching constraints. Their qualifications exceed what the local labour market supplies in terms of compatible, adequately remunerated positions. This supply-demand mismatch generates a reservation wage effect: educated women prefer in-

activity over accepting positions below their qualification level, which paradoxically suppresses aggregate participation even as education expands.

1.4 Gaps in the Literature

Despite extensive national-level research on Morocco's labour market and a growing European literature on spatial gender employment gaps, four major gaps persist in the existing knowledge base that this study directly addresses.

First, Moroccan research has overwhelmingly relied on national aggregates or broad urban/rural distinctions, masking the substantial subnational variation that characterises the country's commune landscape. Analytical frameworks operating at the national level are structurally incapable of capturing this heterogeneity or identifying commune-level determinants.

Second, the dominant methodological tradition in Moroccan and MENA labour market research relies on decomposition approaches, principally Oaxaca–Blinder and recentered influence function decompositions ([Fortin et al., 2011](#)), to attribute observed gaps to differences in endowments versus returns. This tradition is inapplicable here. The HCP census provides commune-level tabulations rather than individual micro-data, so applying decomposition methods to aggregate data would commit the ecological fallacy identified by [Robinson \(1950\)](#). More fundamentally, the research question here is not about attributing the within-commune gap to endowments versus returns, but about explaining spatial variation in the gap across communes, a question that spatial regression is designed to answer and that decomposition methods are not.

Third, the mechanisms generating spatial clustering of gender gaps in Morocco have not been empirically tested. Theoretical channels, including commuting network integration, cultural norm diffusion, shared exposure to regional economic shocks, and spatially correlated omitted variables, are theoretically well motivated but empirically unexamined in the Moroccan context.

Fourth, spatial heterogeneity in the determinants of female participation remains unexplored at the subnational level in Morocco. If coefficients on education, urbanisation, and household structure vary across space, as [Falk and Leoni \(2010\)](#) document for Austria, uniform national policy recommendations derived from spatially homogeneous models may be systematically misleading.

This study addresses these gaps by integrating microeconomic labour supply theory with spatial econometric modelling applied to commune-level census data from Morocco's 2024 census. It explicitly tests for spatial dependence using global and local Moran's I statistics, motivates the introduction of spatial models through Lagrange Multiplier tests, and estimates a Spatially Lagged X model alongside a fractional logit specification with spatially lagged covariates. The similarity of average partial effects

across linear and nonlinear specifications provides evidence on the adequacy of the linear model and the robustness of the substantive findings.

Chapter 2

Data Description

2.1 Data Sources and Geographical Framework

This study draws on Morocco's 2024 General Census of Population and Housing, conducted by the Haut-Commissariat au Plan (HCP; [Haut-Commissariat au Plan, 2024](#)) between September and November 2024. Unlike survey-based data sources, which are subject to sampling error and are typically representative only at the national or regional level, this census provides exhaustive coverage of all 1,503 administrative communes, enabling subnational analysis at the finest available geographic resolution. This is the first general census conducted in Morocco since 2014, making it the most current and comprehensive source of commune-level demographic and labour market data available. The census covers the entire Moroccan population across all administrative communes and provides commune-level indicators spanning labour market participation, unemployment, educational attainment, household structure, and demographic composition, all disaggregated by gender.

Morocco's administrative geography comprises three hierarchical levels: regions, provinces and prefectures, and communes. Regions are the highest-level administrative units, of which there are 12 following the 2015 regionalisation reform ([Haut-Commissariat au Plan, 2015](#)). Each region is subdivided into provinces and prefectures, of which there are 62 in total, and each province or prefecture is further subdivided into communes, the smallest administrative unit and the unit of analysis in this study.¹ Communes are the level at which local governments operate, at which labour markets and social norms are most directly shaped by local economic and cultural conditions, and at which household decisions about female labour force participation are most proximately determined. This makes the commune the natural unit for an

¹A commune is the basic unit of local government in Morocco, governed by an elected municipal council under the framework of the 2015 Organic Law on Communes (Law No. 113-14). Communes exercise autonomous competences over local development, infrastructure, and social services, making them the appropriate level at which to study the interaction between local economic conditions, social norms, and female labour market outcomes.

analysis of the spatial determinants of the gender activity rate gap.

A commune-level shapefile from the HCP provides geographic boundaries for spatial analysis. Commune polygons are matched to census indicators through a name-based merging procedure using Unicode decomposition to resolve prefix variations and diacritical inconsistencies that arise from the transliteration of Arabic and Amazigh place names into French. Four communes in the Southern Provinces Mijik, Zoug, Aghouinite, and Lagouira are excluded due to systematically missing data across all key indicators, likely reflecting administrative incompleteness in these sparsely populated border territories. The final analytical sample comprises 1,499 communes covering all 12 administrative regions of Morocco.

Of the 1,499 communes retained, 204 are urban (13.6%) and 1,295 are rural (86.4%), reflecting the predominantly rural character of Morocco's commune landscape despite ongoing urbanisation.

2.2 Variable Construction

2.2.1 Outcome Variable

The outcome variable is the *gender activity rate gap*, defined as the difference between male and female labour force participation rates among individuals aged 15 and above, expressed in percentage points:

$$\text{Gap}_i = \text{LFPR}_i^{\text{male}} - \text{LFPR}_i^{\text{female}}, \quad i = 1, 2, \dots, 1499. \quad (2.1)$$

A positive value indicates that the male rate exceeds the female rate, which is the case for all 1,499 communes in the sample. The gap is defined at the commune level as an aggregate measure of gender inequality in economic participation, capturing the combined effect of all factors that lead women to participate in the labour force at lower rates than men within a given commune (Elborgh-Woytek et al., 2013; Klasen, 2019). For the fractional regression specifications, the outcome is rescaled to the unit interval by dividing by 100, placing it in $[0, 1]$ as required by the fractional logit estimator (Papke and Wooldridge, 1996).

2.2.2 Explanatory Variables

Eight explanatory variables are included, selected on the basis of theoretical motivation from the frameworks reviewed in Chapter 1 and data availability at the commune level in the 2024 census.²

²Several variables identified by theory as potentially relevant were not included due to unavailability at the commune level in the 2024 census tabulations. These include local wage levels disaggregated

Table 2.1: Variable Definitions

Variable	Definition	Unit
Activity gap	Male minus female labour force participation rate (ages 15+)	pp
Illiteracy rate	Share of working-age population unable to read or write	%
Secondary education	Share of working-age population with secondary attainment	%
Higher education	Share of working-age population with tertiary attainment	%
Unemployment rate	Aggregate (non-gender-specific) unemployment rate	%
Urban	1 if commune classified urban by HCP, 0 otherwise	binary
Household size	Mean number of persons per household	persons
Working-age pop. share	Share of total population aged 15–64	%
Married pop. share	Share of working-age female population that is married	%

Note: All variables constructed from the 2024 HCP census ([Haut-Commissariat au Plan, 2024](#)).

Three education indicators capture the human capital environment. The *aggregate illiteracy rate* measures the share of the working-age population unable to read or write, capturing the baseline stock of human capital deprivation that the U-shaped hypothesis and labour supply theory identify as associated with low female participation in early development stages ([Goldin, 1995](#); [Boserup, 1970](#)). The *secondary education share* and *higher education share* measure the proportions of the working-age population with secondary and tertiary educational attainment respectively, capturing the upward segment of the human capital-participation relationship ([Psacharopoulos and Patrinos, 2018](#)). These three indicators are included separately rather than as a single index because they capture qualitatively different stages of the educational distribution, and their effects on the activity gap are theoretically and empirically distinct. Illiteracy captures deprivation, secondary education captures intermediate attainment, and higher education captures the reservation wage and job-matching dynamics that [Lopez-Acevedo et al. \(2021\)](#) identify as central to the Moroccan paradox.

The *aggregate unemployment rate* proxies local labour market conditions ([Discenza and Faccini, 2025](#)). A higher unemployment rate may reflect a discouraged worker

by sector and gender, commune-level fertility rates, childcare availability, female inheritance and property rights indicators, attitudinal measures of gender role norms, and commuting infrastructure indices. The implications of these omissions for the interpretation of the results are discussed in the Conclusion.

environment that suppresses female entry, or alternatively a demand-constrained setting in which both male and female participation are depressed (Benati, 2001). The aggregate rather than gender-specific rate is used to capture the general local labour market environment, and to avoid the circularity that would arise from including a female-specific labour market outcome as a regressor for the gender gap.

The *binary urban classification indicator* distinguishes urban from rural communes as classified by the HCP (Haut-Commissariat au Plan, 2024). Urban communes are characterised by higher population density, greater availability of formal sector employment in services and industry, and different social norms surrounding female employment (Assaad, 2014). Its inclusion controls for the structural differences between urban and rural labour markets documented below in Section 2.3.2.

The *married population share* measures the proportion of the working-age female population that is married. This variable is the primary empirical proxy for the household bargaining dynamics identified in Chapter 1. Higher rates of marriage are associated with greater domestic responsibilities and stronger social expectations of female withdrawal from the labour force, consistent with the identity norms documented by Hayo and Caris (2013) for MENA contexts and the bargaining framework of Chiappori (1992).

The *working-age population share* measures the proportion of the total population aged 15–64, capturing the demographic structure of the commune (Bloom et al., 2003). Communes with a higher working-age share have a larger potential labour force relative to dependants. In the Moroccan context, where male agricultural and pastoral employment dominates rural economies, a larger working-age population tends to be associated with higher male activity rates, which mechanically widens the gap when female participation remains constrained by social norms and domestic responsibilities. This variable therefore captures both the demographic dividend channel and the gender-asymmetric sectoral composition of rural employment.

Mean household size captures the average number of persons per household (García-Maíñar et al., 2018). Larger households impose greater domestic burdens on women through care and home production responsibilities, reducing their outside options in household bargaining and suppressing female labour force participation (Doss, 2013). This variable is the primary empirical proxy for the fertility and care burden mechanisms identified in Chapter 1.

The female activity rate is excluded from the regressor set to avoid circularity, since it is algebraically related to the outcome variable. Multicollinearity among the retained regressors is assessed through variance inflation factors (VIF), reported in Table 2.2. All VIF values remain below 5 the highest being 3.95 for secondary education confirming the absence of problematic multicollinearity that could destabilise the coefficient estimates.

Table 2.2: Variance Inflation Factors

Variable	VIF	Tolerance
Secondary education	3.95	0.253
Illiteracy rate	3.51	0.285
Higher education	3.14	0.318
Urban	1.93	0.518
Working-age pop. share	1.46	0.686
Married pop. share	1.29	0.776
Household size	1.15	0.873
Unemployment rate	1.11	0.899

Note: VIF computed from OLS on the full regressor set.

Tolerance = $1/\text{VIF}$. All values below the conventional threshold of 5.

2.3 Exploratory Analysis

2.3.1 Descriptive Statistics

Table 2.3: Descriptive Statistics

Variable	Mean	Median	Std. Dev.	Min–Max
Activity gap (pp)	57.75	58.90	8.53	19.5–100.0
Illiteracy rate (%)	36.95	38.50	10.33	3.3–75.0
Secondary education (%)	7.15	6.20	3.87	0.0–27.5
Higher education (%)	5.28	4.00	3.89	0.0–47.9
Unemployment rate (%)	28.00	26.15	13.51	0.0–92.7
Household size	4.23	4.20	0.66	1.3–7.6
Working-age pop. share (%)	57.15	57.00	3.70	47.2–100.0
Married pop. share (%)	60.35	60.90	3.98	34.9–70.9

Note: $N = 1,499$ communes. All variables constructed from the 2024 HCP census ([Haut-Commissariat au Plan, 2024](#)).

The mean activity gap of 57.75 percentage points is nearly 10 points wider than the national headline figure of approximately 50 points ([Haut-Commissariat au Plan, 2024](#)), reflecting the extremely low female labour force participation that persists across the overwhelming majority of communes. The distribution spans from 19.5 to 100 percentage points, indicating that the national average conceals considerable subnational heterogeneity. The Oum Dreyga commune, located in the Guelmim-Oued Noun re-

gion in southern Morocco, reaches the absolute boundary of 100 percentage points. This extreme value reflects a situation of complete female exclusion from the labour force, where no women in the working-age population are recorded as economically active. Oum Dreyga is characterised by a predominantly agricultural and pastoral economy, extreme geographic isolation, and social norms that severely constrain female economic participation conditions that make it an extreme but not isolated case along the distributional tail. The standard deviation of 8.53 percentage points and the near-symmetry of the distribution around the median (58.90 pp) indicate that the activity gap is broadly distributed across communes rather than driven by a small number of outliers, lending credibility to the regression analysis that follows.

Among the education regressors, illiteracy averages 36.95% and reaches 75% in the most educationally deprived communes, reflecting the persistence of severe human capital deficits in remote rural areas. The secondary education share averages only 7.15%, with a median of 6.20%, indicating that the majority of communes have very low shares of working-age residents with secondary education. The higher education share averages 5.28% with a median of 4.00% and a maximum of 47.9%, the latter reflecting the concentration of university-educated populations in major urban communes such as Rabat and Casablanca. The substantial right-skewness of both education shares mean above median reflects the concentration of educational attainment in a relatively small number of urban communes, with the majority of communes having very low shares.

Unemployment exhibits the widest relative dispersion of any regressor (mean 28.0%, range 0.0–92.7%, standard deviation 13.51%), reflecting the extreme heterogeneity of local labour market conditions across Morocco's commune landscape. Some communes record zero unemployment likely reflecting either full absorption of the local workforce in subsistence agriculture or measurement artefacts while others exceed 90%, consistent with the near-complete absence of formal employment in economically marginalised areas. The married population share averages 60.4% with remarkably low dispersion (standard deviation 3.98%), suggesting that the high prevalence of marriage among working-age women is a broadly uniform feature of Morocco's demographic landscape rather than a characteristic specific to particular regions. Mean household size averages 4.23 persons, consistent with the household bargaining dynamics identified in Chapter 1 as central determinants of female participation.

2.3.2 Urban and Rural Comparisons

Figure 2.1 shows the geographic distribution of urban and rural communes across Morocco's 12 administrative regions.




Figure 2.1: Urban and rural commune classification

The mean activity gap is 49.5 percentage points in urban communes and 59.1 points in rural ones, a statistically significant difference of 9.6 points ($t = -16.15, p < 0.001$). This urban–rural differential reflects several reinforcing mechanisms. Urban communes have substantially higher education shares and lower illiteracy, consistent with the negative association between educational attainment and the activity gap. Urban labour markets offer a wider range of formal sector employment in services and industry that is more compatible with female participation than the agricultural employment that dominates rural economies (Assaad, 2014). However, even in urban communes the mean gap of 49.5 percentage points is extremely large by international standards (International Labour Organization, 2023), indicating that urbanisation alone is far from sufficient to close the gender gap in Morocco.

Unemployment rates do not differ significantly between urban and rural communes ($t = -1.16, p = 0.247$), a finding that is itself informative. It suggests that the activity gap is driven primarily by participation constraints social norms, household bargaining, and care responsibilities rather than by differential access to employment opportunities as measured by the unemployment rate. This motivates the inclusion

of household structure and demographic variables alongside labour market proxies in the regression specifications.

Figure 2.2 presents the distributions of key variables by urban and rural classification.




Figure 2.2: Distribution of key variables by urban and rural classification

The activity gap distribution is approximately bell-shaped for both groups but shifted rightward by nearly 10 percentage points for rural communes, which also exhibit greater dispersion consistent with the greater heterogeneity of economic and social conditions across rural Morocco. The illiteracy distributions diverge sharply: rural communes span 10%–75% while urban communes cluster at substantially lower values, confirming the concentration of human capital deprivation in rural areas. The household size distribution is similarly divergent, with rural communes displaying higher means and greater right-skewness, reflecting the persistence of larger family structures in rural Morocco.

2.4 Spatial Distribution

2.4.1 Geographic Patterns

Figure 2.3 presents choropleth maps of the activity gap and key regressors across Morocco's 1,499 communes.




Figure 2.3: Spatial distribution of key variables across Moroccan communes

The activity gap displays a pronounced geographic gradient that reflects Morocco's underlying economic geography. Smaller gaps ranging from approximately 20 to 45 percentage points concentrate along the northern Atlantic corridor, encompassing the urban agglomerations of Casablanca, Rabat, Tangier, and Fès, where formal sector employment in services, industry, and public administration creates compatible employment opportunities for women (Assaad, 2014; Lopez-Acevedo et al., 2021). The Souss-Massa region around Agadir also displays relatively smaller gaps, reflecting the development of agri-food industry and tourism that generate female-compatible employment (Center for Macroeconomics and Development, 2023). By contrast, larger gaps ranging from 65 to 100 percentage points dominate the southern regions, the pre-Saharan zone, the eastern Atlas, and parts of the Atlantic coast south of Safi, areas characterised by predominantly agricultural and pastoral economies, extreme geographic isolation, and limited formal employment opportunities for women (Center for Macroeconomics and Development, 2023).

The illiteracy map mirrors the activity gap pattern closely, with high illiteracy concentrating in the same southern and eastern regions where the gap is largest, reflect-

ing the geographic co-location of educational deprivation and gender inequality documented in Section 2.3.1. The unemployment rate map displays a different pattern, with high unemployment clustering in urban and peri-urban areas the reverse of the activity gap pattern reinforcing the finding that unemployment is not the primary driver of the gap. The married population share map shows relatively uniform coverage across the commune landscape, with somewhat higher values in the south, consistent with the higher prevalence of traditional household structures in more economically marginalised areas.

2.4.2 Global Spatial Autocorrelation

Figure 2.4 presents Moran scatterplots for the activity gap, illiteracy rate, and unemployment rate under the primary KNN $k = 5$ weights. Each plot shows the standardised value of the variable on the horizontal axis against its spatial lag, the average value among the five nearest neighbours on the vertical axis. The positive slopes of all fitted lines indicate positive spatial autocorrelation. Communes with above-average values of each variable tend to be surrounded by neighbours that also have above-average values.




Figure 2.4: Moran scatterplots KNN $k = 5$ weights

Table 2.4 reports global Moran's I statistics (Moran, 1950) for the outcome and all continuous regressors under the primary k -nearest neighbours weights ($k = 5$). Under row-standardisation, Moran's I does not span the theoretical $[-1, +1]$ interval exactly; its bounds depend on the structure of the weights matrix, and inference is conducted via the standardised z -score (Cliff and Ord, 1981). A value of zero corresponds to the expectation under spatial randomness; positive values indicate positive spatial autocorrelation, meaning that similar values tend to cluster geographically. All variables exhibit statistically significant positive spatial autocorrelation, confirming that the geographic patterns visible in the choropleth maps are not random.

Table 2.4: Global Moran's I KNN $k = 5$ (Primary Weights)

Variable	Moran's I	z -score	
Activity gap	0.351	22.68	***
Illiteracy rate	0.479	30.87	***
Secondary education	0.447	28.86	***
Higher education	0.333	21.52	***
Unemployment rate	0.280	18.11	***
Household size	0.691	44.53	***
Working-age pop. share	0.652	42.06	***
Married pop. share	0.483	31.17	***

Note: *** $p < 0.001$. Inference based on the standardised normal approximation.

Weights matrix: KNN $k = 5$, row-standardised.

The activity gap has Moran's $I = 0.351$ ($z = 22.68$, $p < 0.001$), confirming significant positive spatial autocorrelation. Communes with large gaps systematically neighbour other high-gap communes, and communes with small gaps systematically neighbour other low-gap communes. The two variables with the strongest spatial autocorrelation are household size ($I = 0.691$) and working-age population share ($I = 0.652$), reflecting the strong geographic persistence of demographic structures across neighbouring communes consistent with the slow-moving nature of fertility transitions and population age structures that tend to evolve gradually across geographic space (Bloom et al., 2003). Educational variables display intermediate spatial autocorrelation (I between 0.333 and 0.479), reflecting the geographic concentration of educational investment documented in the choropleth maps. Unemployment shows the weakest spatial autocorrelation ($I = 0.280$), consistent with the more idiosyncratic nature of local labour market conditions relative to the structural demographic and educational variables.

The spatial autocorrelation of both the outcome and all regressors has two important methodological implications. First, OLS standard errors will be downward biased under spatial error dependence, necessitating the Spatial HAC correction described in Chapter 3. Second, the spatial autocorrelation of the regressors implies that conditions in neighbouring communes may independently influence the activity gap in a given commune, leading to the use of a model with spatial lags of the explanatory variables (or simply, an SLX model), which is designed to capture these spatial spillover effects.

2.4.3 Local Spatial Autocorrelation

Global Moran's I establishes that the activity gap is spatially autocorrelated across Morocco's commune landscape, but does not identify *where* that autocorrelation is concentrated. Local Indicators of Spatial Association (LISA; [Anselin, 1995](#)) decompose the global statistic into commune-level contributions, identifying statistically significant local clusters and spatial outliers.

Figure 2.5 presents the LISA cluster map for the activity gap under the primary KNN $k = 5$ weights, based on 999 conditional permutations with a significance threshold of $p < 0.05$. Of the 1,499 communes, 387 (25.8%) exhibit statistically significant local spatial autocorrelation, classified into four cluster types. High-high (HH) communes are those where a large gap is surrounded by similarly large gaps in neighbouring communes; low-low (LL) communes are those where a relatively smaller gap is surrounded by similarly smaller gaps; and high-low (HL) and low-high (LH) spatial outliers mark transition zones between these regimes.




Figure 2.5: LISA cluster map of the gender activity rate gap (Local Moran's I , KNN $k = 5$, 999 permutations, $p < 0.05$). Colours indicate the LISA cluster classification: High-High (red), Low-Low (dark blue), Low-High (light blue), High-Low (orange), and not statistically significant (grey).

The 153 HH communes form two geographically coherent hotspots of gender inequality. The first and largest concentrates in the Southern Provinces and Guelmim-

Oued Noun, encompassing communes such as Oum Dreyga where agricultural and pastoral economies and extreme geographic isolation combine to produce near-complete female exclusion from the labour force. The second runs along the Atlantic coast south of Safi and into the pre-Saharan fringe of Drâa-Tafilalet, where similar economic and social conditions prevail. The 167 LL communes form a geographically coherent coldspot of relative gender equality concentrated along the northern Atlantic corridor the Casablanca–Rabat–Tangier axis and the Souss-Massa region around Agadir where formal sector employment and higher educational attainment are most developed. It bears emphasis that these LL communes are not characterised by gender equality in any absolute sense. Their mean activity gap remains substantial, but it is significantly lower than the national commune-level mean of 57.75 percentage points. The 36 HL and 31 LH spatial outliers are sparsely distributed and mark transition zones between the two dominant regimes, consistent with the gradual geographic diffusion of both economic development and social norm change.

These LISA results confirm that the spatial dependence documented by the global Moran's I is not uniformly distributed but is driven by two geographically coherent clusters whose locations align precisely with Morocco's underlying economic geography a southern and pre-Saharan hotspot of entrenched gender inequality, and a northern Atlantic coldspot of relative female economic participation. This geographic structure directly motivates the spatial econometric specifications described in Chapter 3, which are designed to account for both local spatial dependence and cross-commune spillover effects.

Chapter 3

Methodology

This chapter sets out the empirical strategy employed to analyse the determinants of the gender activity rate gap across Morocco's 1,499 communes. The analysis proceeds in four stages, progressing from a non-spatial baseline to spatial specifications that account for the geographic dependence documented in Chapter 2. The gender activity rate gap is rescaled to the unit interval ($y_i^* = y_i/100$, clipped to $[\varepsilon, 1 - \varepsilon]$ with $\varepsilon = 10^{-6}$), which ensures consistency of scale across linear and nonlinear specifications and facilitates direct comparison of partial effects. This rescaling is stated once here and is not repeated in the equations that follow, where the outcome is denoted y^* . The single commune reaching the boundary value of 100 percentage points is handled by the clipping procedure, which avoids undefined log-odds in the fractional specifications without materially affecting any estimated coefficient. The analysis is observational and identifies conditional associations rather than causal effects. The cross-sectional census data do not support causal identification, and all estimated coefficients should be interpreted accordingly.

3.1 OLS Baseline

The OLS baseline model takes the form

$$y^* = \mathbf{X}\beta + \varepsilon \quad (3.1)$$

where y^* is the $n \times 1$ vector of rescaled gender activity rate gaps, $\mathbf{X}$ is the $n \times 9$ matrix of commune-level covariates including a column of ones, and ε is the error vector. The covariate matrix includes a constant, the illiteracy rate, secondary education share, higher education share, unemployment rate, urban indicator, household size, working-age population share, and married population share. The female activity rate is excluded to avoid circularity; the aggregate unemployment rate is retained as it captures local labour market conditions rather than being algebraically determined

by the outcome.

OLS serves as the natural starting point for three reasons. First, the linear probability model provides directly interpretable marginal effects in percentage point units, facilitating substantive interpretation and comparison across specifications. Second, OLS residuals form the basis of the Lagrange Multiplier tests that guide spatial model selection, as described in Section 3.6. Third, OLS provides a useful benchmark against which the average partial effects from the nonlinear specifications can subsequently be compared.

The cross-sectional nature of the data, together with the existence of spatial dependence documented in Chapter 2, makes two forms of non-spherical errors likely. First, heteroskedasticity arises because commune-level outcomes are aggregates of populations that vary substantially in size, generating variance that is likely correlated with commune characteristics. Second, spatial error autocorrelation arises because spatially adjacent communes share unobserved factors common infrastructure, cultural environments, and regional economic shocks that generate residual correlation across space. Standard errors are therefore estimated using a heteroskedasticity and autocorrelation consistent (HAC) kernel estimator following [Kelejian and Prucha \(2007\)](#), with bandwidth $h = \lfloor n^{0.25} \rfloor = 6$ and a triangular kernel. This estimator is used precisely to address both heteroskedasticity and spatial autocorrelation in the residuals, without imposing a parametric structure on the error covariance, ensuring valid inference even when the precise form of either is unknown. The presence of heteroskedasticity is confirmed by Breusch–Pagan and White tests reported in Chapter 4. Spatial HAC standard errors correct for its inferential consequences without affecting the consistency of the coefficient estimates, which hold under heteroskedasticity by standard OLS theory.

3.2 Fractional Regression Model

The gender activity rate gap rescaled to $[0, 1]$ is a fractional outcome. It is bounded, and OLS does not respect this constraint. Because OLS ignores the scale of the outcome, its predicted values can fall outside $[0, 1]$, an inadmissible range for a variable defined as a proportion. In addition, the conditional variance of a bounded outcome is necessarily a function of its conditional mean, violating the homoskedasticity assumption. Following [Papke and Wooldridge \(1996\)](#), a fractional logit model is estimated by quasi-maximum likelihood (QMLE), specifying the conditional mean as

$$E[y_i^* \mid \mathbf{x}_i] = G(\mathbf{x}_i\boldsymbol{\gamma}) = \frac{\exp(\mathbf{x}_i\boldsymbol{\gamma})}{1 + \exp(\mathbf{x}_i\boldsymbol{\gamma})}, \quad i = 1, 2, \dots, 1499. \quad (3.2)$$

where $G(\cdot)$ is the logistic function, which by construction maps the linear index $\mathbf{x}_i\gamma$ into the unit interval. The QMLE estimator maximises the Bernoulli log-likelihood

$$\mathcal{L}(\gamma) = \sum_{i=1}^n [y_i^* \log G(\mathbf{x}_i\gamma) + (1 - y_i^*) \log(1 - G(\mathbf{x}_i\gamma))] \quad (3.3)$$

even though y_i^* does not follow a Bernoulli distribution. Consistency of the QMLE requires only correct specification of the conditional mean $E[y_i^* | \mathbf{x}_i] = G(\mathbf{x}_i\gamma)$ and does not require y_i^* to follow any particular distribution (Papke and Wooldridge, 1996). This robustness to distributional misspecification is the key advantage of QMLE over maximum likelihood estimation for fractional outcomes.

Standard errors are estimated using White standard errors, corrected for small sample size. Since the logistic function is nonlinear, the coefficient vector γ does not directly measure marginal effects. Average Partial Effects (APEs) are therefore computed as

$$\widehat{\text{APE}}_k = \frac{1}{n} \sum_{i=1}^n \hat{\gamma}_k G(\mathbf{x}_i\hat{\gamma}) [1 - G(\mathbf{x}_i\hat{\gamma})], \quad k = 1, 2, \dots, 8. \quad (3.4)$$

which average the derivative of the conditional mean with respect to covariate k across all observations, yielding effects in the same percentage point units as the OLS coefficients and facilitating direct comparison.

Functional form adequacy is assessed through two specification tests. The link test regresses the outcome on the fitted index and its square; a significant coefficient on the square indicates link function misspecification. The generalised RESET test augments the conditional mean with powers of the fitted values; rejection indicates that the logistic link does not adequately capture the nonlinearity in the conditional mean. Moran's I computed on the response residuals $\hat{u}_i = y_i^* - G(\mathbf{x}_i\hat{\gamma})$ establishes whether spatial dependence persists after nonlinear conditioning on the covariates, motivating the fractional logit with spatially lagged regressors specification described in Section 3.5.

3.3 Spatial Weights Matrix

The spatial weighting matrix $\mathbf{W}$ formalises the notion of neighbourhood and is a prerequisite for all spatial specifications. The choice of spatial weighting matrix is consequential. It determines which communes are considered neighbours, how strongly their characteristics influence each other, and therefore what the spatial lag of any variable represents. If $\mathbf{W}$ is row-standardised, the spatial lag of a variable $\mathbf{x}$ for commune i , given by $(\mathbf{W}\mathbf{x})_i = \sum_j w_{ij}^* x_j$, is a weighted average of the values of $\mathbf{x}$ in commune i 's neighbouring communes. Several specifications are constructed to ensure that results

are not artefacts of a particular neighbourhood definition.

The primary specification employs a k -nearest neighbours spatial weighting matrix with $k = 5$, constructed from geodesic distances between commune centroids.¹ Under this criterion, every commune is assigned exactly five neighbours regardless of its geographic configuration, ensuring a balanced and complete spatial structure with no “islands”. This matters particularly for Morocco, where 81 communes share no boundary with any other commune under queen contiguity: a contiguity-based criterion requires the existence of at least one neighbour, and would therefore exclude these 81 communes from any spatial analysis entirely. The choice of $k = 5$ reflects a standard benchmark in the spatial econometrics literature for commune-level analyses (LeSage and Pace, 2009). Sensitivity to this choice is assessed by re-estimating all spatial models under $k = 3$, $k = 8$, and $k = 10$, as well as under queen contiguity, as robustness checks.

The matrix is row-standardised so that each row sums to unity:

$$w_{ij}^* = \frac{w_{ij}}{\sum_j w_{ij}}, \quad i = 1, 2, \dots, 1499. \quad (3.5)$$

where $w_{ij} = 1$ if commune j is among the five nearest neighbours of commune i and zero otherwise. Row-standardisation ensures that the spatial lag of any variable is a weighted mean of that variable across the five nearest neighbours, giving the spatial lag a natural and interpretable units-preserving meaning. It also ensures that the largest eigenvalue of $\mathbf{W}$ is unity, which is a standard regularity condition for spatial models (LeSage and Pace, 2009).

The robustness specification also employs a first-order queen contiguity matrix, under which communes i and j are neighbours if they share at least one boundary point. This matrix produces 81 island communes and a mean of 5.13 neighbours per commune, and is used to verify that the spatial parameters and covariate effects are not sensitive to the neighbourhood definition. Both matrices are row-standardised prior to estimation.

3.4 Spatially Lagged X Model

LeSage and Pace (2009) provide the theoretical foundations for the family of spatial regression models employed in this study. The spatial autocorrelation documented in Chapter 2, together with the existence of spatial dependence in all regressors, implies that spillover effects operating across commune boundaries may be substantively important and that OLS will not exhibit the usual statistical properties if such depen-

¹Centroid distances are computed on the projected coordinate system EPSG:26191 (Merchich / Nord Maroc), which is appropriate for distance calculations across Morocco.

dence is left unaddressed. The Spatially Lagged X (SLX) model addresses this by augmenting the OLS baseline with spatial lags of all covariates:

$$y^* = \mathbf{X}\beta + \mathbf{WX}\theta + \varepsilon \quad (3.6)$$

where $\mathbf{W}$ is the row-standardised k -nearest neighbours spatial weighting matrix, $\mathbf{WX}$ is the $n \times 8$ matrix of spatially lagged covariates (the constant is not lagged), and θ is the 8×1 vector of spillover (indirect-effect) coefficients associated with each lagged covariate. The constant is absorbed into $\mathbf{X}$ as in equation (3.1).

The SLX model is preferred over the Spatial Lag Model (SLM) and Spatial Error Model (SEM) for three reasons. First, the SLX model does not introduce a lagged dependent variable $\mathbf{W}y^*$, which would generate simultaneity and require instrumental variable or maximum likelihood estimation. Under the maintained exogeneity assumptions, estimation by OLS with spatial HAC standard errors provides consistent coefficient estimates and valid inference. Second, in the SLX model, the direct effects, β_k , and the indirect effects, θ_k , do not require simulation from a spatial multiplier matrix, unlike the SLM, where simulation is required to obtain the standard errors of the spatial effects (LeSage and Pace, 2009). Third, omitting spatially lagged covariates from a non-spatial model is directly analogous to omitting a relevant lagged regressor in a time-series setting: it violates dynamic completeness and induces omitted-variable bias in the remaining coefficients, rather than implying that the true data-generating process must be an SEM or a Spatial Durbin Model. Respecifying the conditional mean to include $\mathbf{WX}$ addresses this omitted-variable concern directly. Whether this is sufficient to absorb the spatial structure present in the OLS residuals is then confirmed empirically by the Lagrange Multiplier tests described in Section 3.6, rather than assumed at the outset.

The interpretation of the SLX coefficients is as follows. The coefficient β_k measures the direct effect of covariate k in commune i on its own activity gap, holding the values of all other covariates and all neighbouring covariates constant. The coefficient θ_k measures the indirect effect of the average value of covariate k across commune i 's five nearest neighbours on commune i 's own activity gap, holding own-commune covariates constant. The total effect of a one unit change in covariate k that occurs uniformly across a commune and all its neighbours is $\beta_k + \theta_k$. HAC standard errors with the same kernel bandwidth as the OLS baseline are used throughout.

3.5 Fractional Logit with Spatially Lagged Regressors

Building on Papke and Wooldridge (1996) and the spatial framework of LeSage and Pace (2009), a fractional logit model with spatially lagged regressors is estimated as

a robustness check on the SLX model. This specification addresses two concerns simultaneously. One, the bounded nature of the outcome, which OLS does not respect. Two, the possible nonlinearity of the conditional mean, which the linear specification assumes away. If the linear SLX provides an adequate approximation of the true conditional mean, the average partial effects from the fractional logit with spatially lagged regressors should be similar to the direct and indirect effects from the linear SLX.

$$\mathbb{E}(y_i^* \mid \mathbf{X}, \mathbf{W}) = G(\mathbf{x}_i' \boldsymbol{\gamma} + (\mathbf{W}\mathbf{X})_i' \boldsymbol{\theta}), \quad i = 1, 2, \dots, 1499. \quad (3.7)$$

where $G(\cdot)$ denotes the logistic cumulative distribution function, $\mathbf{x}_i$ is the vector of explanatory variables for commune i , $(\mathbf{W}\mathbf{X})_i$ is the corresponding vector of spatially lagged explanatory variables, $\boldsymbol{\gamma}$ contains the direct-effect coefficients, and $\boldsymbol{\theta}$ contains the spillover-effect coefficients.

The model is estimated by QMLE with White standard errors, corrected for small sample size. Average partial effects (APEs) are computed separately for the direct covariates $\mathbf{X}$ and the spatially lagged covariates $\mathbf{W}\mathbf{X}$, yielding direct and indirect APEs in percentage-point units that are directly comparable to the $\hat{\beta}$ and $\hat{\theta}$ from the linear SLX.

3.6 Model Selection and Comparison

Spatial model selection follows the Florax–Folmer–Rey decision rule ([Florax et al., 2003](#)), a sequential testing procedure that uses Lagrange Multiplier (LM) tests on OLS residuals to identify the most appropriate spatial specification. The procedure operates in two steps.

In the first step, two standard LM tests are computed on the OLS residuals. The LM-Error test assesses whether the residuals exhibit significant spatial error autocorrelation, consistent with a Spatial Error Model (SEM), where unobservables in commune i are directly influenced by unobservables in neighbouring communes. The LM-Lag test assesses whether a spatially lagged dependent variable $\mathbf{W}y^*$ belongs in the model, that is, whether the outcome in commune i is directly influenced by the outcomes in neighbouring communes, consistent with a Spatial Lag Model (SLM). Both tests are asymptotically $\chi^2(1)$ under the null hypothesis of no spatial dependence.

In the second step, when both standard tests strongly reject the corresponding null hypothesis as is typically the case in practice, since SEM and SLM dependence are not mutually exclusive the robust versions of the tests are used to discriminate between the two structures. The Robust LM-Error test conditions on the presence of a spatial lag, so that a rejection indicates residual error dependence beyond what an SLM would explain. Conversely, the Robust LM-Lag test conditions on the presence of a spatial

error, so that a rejection indicates lag dependence beyond what a SEM would explain. The preferred spatial structure is the one corresponding to the test with the largest test statistic, or equivalently the smallest p-value. A dominant Robust LM-Error points toward a SEM, and a dominant Robust LM-Lag points toward an SLM. These tests can also be conducted on the model extended with spatial lags of the covariates, WX , to assess whether residual spatial dependence is fully captured once these are included.

The numerical results of these tests are reported and discussed in Chapter 4, where they are used to motivate the adoption of the SLX as the primary spatial specification. Model fit across all four specifications is also compared using the Akaike Information Criterion (AIC) and Moran's I on the corresponding residuals. A lower AIC indicates a better-fitting model. Moran's I on the residuals of each specification serves as a descriptive diagnostic of residual spatial dependence, but it does not by itself discriminate between dependence absorbed by WX and global spatial autocorrelation arising from factors common to neighbouring communes that no covariate set can fully capture. The decisive test of whether the spatially lagged covariates absorb the spatial structure identified in the OLS residuals is therefore whether the Robust LM tests, recomputed on the WX -augmented residuals, are no longer significant, irrespective of whether Moran's I itself declines. Heteroskedasticity is assessed throughout using Breusch-Pagan and White tests. Its presence motivates the use of Spatial HAC standard errors for the spatial linear models and White standard errors for the fractional logit model, rather than invalidating the coefficient estimates, which remain consistent under heteroskedasticity.

Chapter 4

Results

This chapter presents the results from the four estimated models, following the sequence set out in Chapter 3. Section 4.1 presents the specification diagnostics and the Lagrange Multiplier tests that motivate the choice of spatial model, before any coefficients are interpreted. Section 4.2 reports the direct and indirect effects from all four specifications in a single results table and interprets the estimates from the selected model. Section 4.3 situates these findings against the existing empirical literature on subnational gender employment gaps.

4.1 Model Selection

The non-spatial OLS baseline explains 48.6% of the variance in the rescaled gender activity rate gap ($R^2 = 0.486$, adjusted $R^2 = 0.483$). Seven of the eight regressors are significant at the 1% level, the urban indicator being the sole exception ($p = 0.273$). All education coefficients carry the expected negative sign. A one percentage point increase in the secondary education share is associated with a 0.88 pp reduction in the activity gap, and a one pp increase in higher education with a 0.70 pp reduction. Household size and married population share are positive and significant, consistent with household bargaining models in which larger and more married communes impose greater domestic constraints on female participation. Household size carries the largest coefficient after the constant (1.99 pp), followed by the working-age population share (0.96 pp), the latter plausibly reflecting higher male agricultural activity rates in communes with a larger working-age share.

Three diagnostic tests establish the need for both robust standard errors and spatial modelling. The Breusch–Pagan test rejects homoskedasticity ($\chi^2 = 126.35$, $p < 0.001$), as does the White test ($\chi^2 = 288.47$, $p < 0.001$). The Ramsey RESET test detects functional form misspecification ($F = 7.51$, $p = 0.006$). Moran’s I on the OLS residuals equals 0.244 ($z = 15.75$, $p < 0.001$), indicating that substantial spatial dependence

remains after conditioning on the eight regressors.

The Lagrange Multiplier tests guide spatial model selection. Both standard tests reject strongly the corresponding null hypothesis (LM-Error = 245.87, $p < 0.001$; LM-Lag = 184.93, $p < 0.001$). Among the robust versions, LM-Error dominates (Robust LM-Error = 67.66, $p < 0.001$; Robust LM-Lag = 6.72, $p = 0.010$), nominally pointing toward a SEM. As discussed in Section 3.4, however, this pattern is equally consistent with omitted spatially lagged covariates rather than an SEM data-generating process. The LM and Robust LM tests applied to the model extended with WX test this directly by checking whether spatial dependence is fully captured once these covariates are included.

A non-spatial fractional logit, estimated as a functional-form check on the OLS baseline, supports the same conclusion. The link test finds no evidence of link function misspecification ($p = 0.835$), and the generalised RESET test finds no evidence of functional form misspecification ($p = 0.971$ for the squared term, $p = 0.958$ for the cubic term). The average partial effects from this fractional logit are, in sign, magnitude, and significance pattern, essentially indistinguishable from the OLS coefficients (McFadden pseudo $R^2 = 0.016$), consistent with the asymptotic equivalence of OLS and QMLE when the dependent variable is concentrated in the interior of the unit interval (Papke and Wooldridge, 1996). Moran's I on the fractional logit residuals equals 0.248 ($p < 0.001$), confirming that the spatial dependence identified above is not an artefact of functional form.

Adding eight spatially lagged covariates (WX) to the OLS baseline — the SLX specification increases R^2 from 0.486 to 0.497 and reduces the AIC from -4104.41 to -4122.74 . The Ramsey RESET test on the SLX specification continues to detect functional form misspecification, albeit less strongly than for the OLS baseline ($F = 3.98$, $p = 0.046$, versus $F = 7.51$, $p = 0.006$ for OLS). Moran's I on the SLX residuals is 0.246 ($p < 0.001$). The Robust LM tests on the SLX residuals, however, find no significant residual spatial dependence (Robust LM-Error $p = 0.243$, Robust LM-Lag $p = 0.915$), confirming that the spatially lagged covariates absorb the spatial structure identified in the OLS residuals without requiring a lagged dependent variable. Table 4.1 summarises fit statistics across all four estimated models.

Table 4.1: Model Comparison and Selection Criteria

Model	N	R^2 / pseudo- R^2	AIC	Moran's I (resid.)	SE type	WX
OLS	1,499	0.486	-4104.4	0.244***	HAC	No
Fractional Logit	1,499	0.016	1350.9	0.248***	White	No
SLX	1,499	0.497	-4122.7	0.246***	HAC	Yes
Fractional Logit SLX	1,499	0.017	1366.4	0.249***	White	Yes

*** $p < 0.001$. AIC not comparable between linear and nonlinear models.

On the basis of this evidence, the SLX model is selected as the primary specification, on the grounds of a lower AIC relative to OLS (-4122.7 versus -4104.4), the absorption of residual spatial dependence confirmed by the Robust LM tests, and the addition of substantively meaningful indirect effects. The persistence of Moran's I on residuals across all four models reflects global spatial autocorrelation in omitted variables rather than structural misspecification, given that the Robust LM tests on the SLX residuals find no actionable spatial dependence left unaddressed. The fractional logit with spatially lagged regressors (converged in three iterations, McFadden pseudo $R^2 = 0.017$, AIC = 1366.44) is retained as a robustness check on the functional form of the selected SLX specification — directly addressing the residual RESET rejection noted above — and is interpreted alongside it in Section 4.2. As shown there, the direct and indirect effects of this nonlinear specification are near-identical to the linear SLX estimates in sign, magnitude, and significance, indicating that the residual functional form misspecification detected by the RESET test does not materially affect the substantive conclusions drawn from the SLX model.

4.2 Alternative Models: Direct and Indirect Effects

Table 4.2 reports direct and indirect effects, with significance stars, across all four estimated specifications: OLS, the non-spatial fractional logit, the selected SLX model, and the fractional logit with spatially lagged regressors.

Table 4.2: Average Partial Effects All Specifications

Variable	OLS		Frac. Logit		SLX (direct)		Frac. SLX (direct)	
	Coef.	Sig.	APE	Sig.	Direct	Sig.	Direct APE	Sig.
Illiteracy	−0.0013	***	−0.0014	***	−0.0014	**	−0.0016	***
Secondary education	−0.0088	***	−0.0093	***	−0.0092	***	−0.0098	***
Higher education	−0.0070	***	−0.0071	***	−0.0059	***	−0.0061	***
Unemployment	0.0004	***	0.0004	***	0.0004	**	0.0004	**
Urban	−0.0069		−0.0054		−0.0063		−0.0054	
Household size	0.0199	***	0.0181	***	0.0206	***	0.0185	***
Working-age pop.	0.0096	***	0.0106	***	0.0065	***	0.0076	***
Married pop.	0.0054	***	0.0055	***	0.0051	***	0.0054	***
					SLX (indirect)		Frac. SLX (indirect)	
Illiteracy					0.0004		0.0008	
Secondary education					0.0005		0.0009	
Higher education					−0.0024	*	−0.0022	*
Unemployment					0.0002		0.0002	
Urban					0.0254	*	0.0232	*
Household size					0.0021		0.0030	
Working-age pop.					0.0050	***	0.0047	***
Married pop.					0.0003		−0.0002	

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. OLS and SLX: HAC SE. Frac. Logit and Frac. SLX: White SE, small-sample corrected.

Only the selected SLX model is interpreted in depth below. The remaining three specifications OLS, the non-spatial fractional logit, and the fractional logit with spatially lagged regressors yield estimates that are essentially indistinguishable from the SLX results in sign, magnitude, and significance across every regressor, as Table 4.2 shows directly. The largest difference across any pair of comparable effects is 0.0010 pp (working-age population share, OLS versus SLX), which is economically negligible. This consistency indicates that the linear model provides an adequate approximation of the conditional mean, that the bounded nature of the outcome does not distort inference, and that the substantive conclusions are robust to both functional form and spatial specification.

Among the SLX direct effects, secondary education remains the largest significant negative driver (−0.92 pp), followed by higher education (−0.59 pp). Illiteracy carries a smaller negative direct effect (−0.14 pp), reflecting the fact that its variation is partly absorbed by the secondary and higher education shares, with which it is correlated. Household size (+2.06 pp) and married population share (+0.51 pp) retain their sign and significance, reinforcing the household bargaining interpretation, while the urban indicator remains statistically insignificant once the other covariates are controlled for, implying that the urban/rural gap documented in Chapter 2 is fully mediated by the included covariates.

Among the indirect effects, working-age population share shows the strongest

and most precisely estimated positive spillover, significant at the 1% level (+0.50 pp, $p < 0.01$). Communes surrounded by neighbours with a higher working-age share tend to have larger activity gaps, consistent with the spillover of agricultural male employment across adjacent communes. The indirect effect of higher education is negative and marginally significant (-0.24 pp, $p < 0.10$), suggesting that communes with highly educated neighbours have modestly smaller gaps, plausibly through norm diffusion or commuting access to formal employment. The urban indirect effect is positive and marginally significant (+2.54 pp, $p < 0.10$), indicating that communes adjacent to urban centres do not inherit the lower gaps of those centres, consistent with the spatial concentration of formal employment within rather than around urban communes. The corresponding APEs from the fractional logit SLX (+0.47 pp, -0.22 pp, and +2.32 pp respectively) confirm the same pattern of significant spillovers.

4.3 Comparison with Previous Empirical Results

The SLX results connect to several strands of the existing empirical literature on sub-national gender employment gaps reviewed in Chapter 1. The positive indirect effect of working-age population share mirrors the labour market spillovers documented by [Noback et al. \(2013\)](#) across Dutch municipalities, providing the first such evidence of cross-boundary labour market spillovers in a MENA context. The significant negative indirect effect of higher education is consistent with the norm-diffusion and commuting-access channels emphasised by [Hayo and Caris \(2013\)](#) for MENA identity norms and by [Noback et al. \(2013\)](#) for neighbouring labour market conditions more generally.

Unlike the spatially heterogeneous education effects that [Falk and Leoni \(2010\)](#) document across Austrian districts, where coefficients on education and childcare availability vary substantially by region, the direct effects of education estimated here are stable across the Moroccan commune landscape. This is plausibly because job-matching barriers for educated women operate broadly across Morocco rather than being concentrated in specific regions.

The gap between the OLS direct-effects-only estimates and the additional indirect effects uncovered by the SLX model echoes the finding of [Chica-Olmo et al. \(2021\)](#) for Andalusian municipalities, that non-spatial regressions systematically underestimate spillover effects and yield biased coefficients when spatial dependence is present but unmodelled. The present results extend this finding to a MENA, commune-level setting for the first time.

Finally, the strong negative direct effects of secondary and higher education, set against Morocco's declining national female participation rate over a period of sustained educational expansion, are consistent with the reservation-wage mechanism

documented by [Lopez-Acevedo et al. \(2021\)](#): as women's qualifications increasingly exceed what the local labour market supplies in compatible, adequately remunerated positions, education raises participation within the cross-sectional variation captured here without resolving the aggregate national paradox. The commune-level association estimated in this study should therefore be read as evidence of cross-sectional heterogeneity in the gap, not as a guarantee that further education alone will close it at the national level.

Chapter 5

Discussion and Conclusion

This chapter interprets the findings of Chapter 4 in light of the theoretical frameworks of Chapter 1, develops policy implications, and summarises the contributions and limitations of the study. The detailed comparison of these results with the existing empirical literature on subnational gender employment gaps is presented in Section 4.3; this chapter instead focuses on what the findings imply for the theoretical frameworks reviewed in Chapter 1.

5.1 Findings, Theory, and the Literature

The strong negative direct effects of secondary and higher education (0.59–0.92 pp across specifications) support the U-shaped development curve’s upward segment, yet national participation has declined even as education expanded, the Moroccan paradox resolved by a reservation wage mechanism in which educated women increasingly face job-matching constraints and prefer inactivity over under-qualified positions ([Lopez-Acevedo et al., 2021](#)). The commune-level association thus reflects cross-sectional heterogeneity rather than a guarantee that education alone closes the gap, consistent with the within-country U-shape documented elsewhere ([Klasen, 2019](#)).

The positive direct effects of household size (+2.06 pp) and married population share (+0.51 pp) support a bargaining channel in which higher household income and limited outside options suppress participation ([Klasen, 2019](#)). With household size averaging 4.23 persons and married women exceeding 60% of the working-age female population, these dynamics are central rather than peripheral to the Moroccan case. The identity channel emphasised for MENA contexts ([Hayo and Caris, 2013](#)) best explains why the urban indicator is insignificant once other covariates are controlled for, yet its marginal positive indirect effect (+2.54 pp) shows that neighbouring communes do not inherit urban centres’ lower gaps a spatially sticky constraint.

The spatial dimension adds a layer that the three canonical frameworks reviewed

in Chapter 1 cannot capture by construction, since none accounts for spillover effects across commune boundaries (Section 1.1). The positive indirect effect of working-age population share and the negative indirect effect of higher education, reported in Section 4.2 and compared against the European spatial labour-economics literature in Section 4.3, confirm that gender inequality in Morocco is a spatially propagating phenomenon rather than one that is fully determined within commune boundaries. Substantively, this means that own-commune policy levers alone raising local education or formalising local employment will understate their true effect, since part of the benefit diffuses to, or is constrained by, neighbouring communes.

5.2 Policy Implications

Education investment is the single most effective lever, but the reservation wage effect ([Lopez-Acevedo et al., 2021](#)) means it must be paired with demand-side measures, such as incentives for firms locating in high-gap communes, to avoid outpacing compatible job creation. Bargaining constraints call for expanded childcare, continued enforcement of the 2004 Family Code's property reforms, and wider microcredit access for women. Because high-gap communes neighbour other high-gap communes, isolated interventions will underestimate the dosage required. Spatially coordinated interventions across entire clusters, embedded in Morocco's regional development frameworks, are likely more effective. Transport and commuting infrastructure could extend urban employment benefits to adjacent rural communes, and the negative education spillover implies long-run returns to education that exceed the direct effects estimated here, which community-based awareness and role-model programmes could help amplify.

5.3 Summary and Limitations

This study provides the first spatial econometric analysis of gender labour market gaps at the commune level in a MENA context, decomposing the gender activity rate gap into direct and indirect spillover effects across Morocco's 1,499 communes using the 2024 census. The gap exhibits strong positive spatial autocorrelation, confirming that geographic clustering of gender inequality is structural. Direct determinants are consistent across all four specifications education is the strongest negative driver, while household structure is the strongest positive one. and the SLX model identifies spillovers, for working-age population share, higher education, and the urban indicator, that OLS misses entirely. Throughout, inference relies on heteroskedasticity- and autocorrelation-consistent (HAC) standard errors that jointly correct for heteroskedasticity and residual spatial dependence, as detailed in Chapter 3; without this correc-

tion, the OLS and SLX standard errors reported in Chapter 4 would be downward biased and inference overly optimistic. The near-perfect similarity of average partial effects across specifications confirms the linear model adequately approximates the conditional mean and that conclusions are robust to functional form and spatial specification.

Three limitations apply. The cross-sectional census precludes the examination of the time dimension of the problem, and although Robust LM tests confirm no actionable spatial dependence is left unaddressed and results are robust to the choice of weighting matrix, the persistence of Moran's I on residuals indicates spatially correlated omitted variables remain. The absence of individual-level microdata prevents analysis of within-commune heterogeneity. Finally, several theoretically central variables disaggregated local wages, fertility rates, childcare availability, property rights indicators, and attitudinal gender-norm data are unavailable in the 2024 HCP tabulations, limiting the study to the observable dimensions of a richer process. Future research could combine multiple census rounds in panel spatial models, examine spatial heterogeneity in coefficients via geographically weighted regression, and use survey microdata to bridge the macro patterns documented here with their micro-level mechanisms.

References

- Anselin, L. (1995). Local indicators of spatial association—LISA. *Geographical Analysis* 27.
- Assaad, R. (2014). *The Jordanian Labour Market in the New Millennium*. Oxford: Oxford University Press.
- Benati, L. (2001). The cyclical behaviour of equilibrium unemployment and vacancies revisited. *Oxford Bulletin of Economics and Statistics* 63(5).
- Bloom, D. E., D. Canning, and J. Sevilla (2003). The demographic dividend: A new perspective on the economic consequences of population change. Technical report, RAND Corporation, Santa Monica, CA.
- Boserup, E. (1970). *Woman's Role in Economic Development*. New York: St. Martin's Press.
- Center for Macroeconomics and Development (2023). Gender inequality in the labor market: The case of Morocco. Technical report, Center for Macroeconomics and Development.
- Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy* 100(3).
- Chica-Olmo, P. et al. (2021). Spatial analysis of gender employment gaps in andalusia. *Applied Spatial Analysis and Policy*.
- Cliff, A. D. and J. K. Ord (1981). *Spatial Processes: Models and Applications*. London: Pion.
- Discenza, A. R. and R. Faccini (2025). Local labour market conditions and female labour force participation: Evidence from Italian municipalities. *Regional Science and Urban Economics* 110.
- Doss, C. (2013). Intrahousehold bargaining and resource allocation in developing countries. *World Bank Research Observer* 28(1).

- Elborgh-Woytek, K. et al. (2013). Women, work, and the economy: Macroeconomic gains from gender equity. Technical report, International Monetary Fund, Washington, DC. IMF Staff Discussion Note SDN/13/10.
- Falk, M. and T. Leoni (2010). Determinants of female labour force participation: A spatial analysis. *Regional Studies* 44.
- Florax, R. J., H. Folmer, and S. J. Rey (2003). Specification searches in spatial econometrics: The relevance of Hendry's methodology. *Regional Science and Urban Economics* 33(5).
- Fortin, N., T. Lemieux, and S. Firpo (2011). Decomposition methods in economics. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4. Elsevier.
- García-Maíñar, A., J. A. Molina, and J. Velilla (2018). Household size and composition and women's labour force participation in Spain. *Feminist Economics* 24(1).
- Goldin, C. (1995). The u-shaped female labor force function in economic development. *NBER Working Paper*.
- Haut-Commissariat au Plan (2015). Loi organique no. 113-14 relative aux communes. Official Gazette of the Kingdom of Morocco.
- Haut-Commissariat au Plan (2024). Recensement général de la population et de l'habitat 2024: Résultats préliminaires.
- Hayo, B. and T. Caris (2013). Female labour force participation in the MENA region: The role of identity. *Review of Middle East Economics and Finance* 9(3).
- International Labour Organization (2023). World employment and social outlook: Trends 2023.
- International Labour Organization (2024). New data shine light on gender gaps in the labour market.
- Kelejian, H. H. and I. R. Prucha (2007). HAC estimation in a spatial framework. *Journal of Econometrics* 140(1).
- Klasen, S. (2019). What explains uneven female labor force participation levels and trends in developing countries? *World Bank Research Observer* 34.
- Klasen, S. and I. Gaddis (2014). Economic development, structural change, and women's labor force participation. *Journal of Population Economics* 27.
- LeSage, J. and R. K. Pace (2009). *Introduction to Spatial Econometrics*. CRC Press.

- Lopez-Acevedo, G. et al. (2021). *Morocco's Jobs Landscape*. World Bank.
- McKinsey Global Institute (2015). The power of parity: How advancing women's equality can add \$12 trillion to global growth. Technical report, McKinsey & Company.
- Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika* 37.
- Noback, I., L. Broersma, and J. Van Dijk (2013). Gender-specific spatial interactions on regional labour markets. *Regional Studies* 47(8).
- OECD (2020). Changing laws and breaking barriers for women's economic empowerment in Egypt, Jordan, Morocco and Tunisia. Technical report, OECD Publishing, Paris.
- Papke, L. E. and J. M. Wooldridge (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics* 11(6).
- Psacharopoulos, G. and H. A. Patrinos (2018). Returns to investment in education: A decennial review of the global literature. *Education Economics* 26(5).
- Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review* 15(3).
- TheGlobalEconomy.com (2024). Morocco: Female labor force participation rate.
- Think Global Health (2023). Women's labor-force participation in the Middle East.
- Tsani, S., L. Paroussos, C. Fragiadakis, I. Charalambidis, and P. Capros (2013). Female labour force participation and economic growth in the MENA region. *Economics Letters* 120(2).
- UN Women (2023). Facts and figures: Economic empowerment.
- World Bank (2023). Women, business and the law 2023. Technical report, World Bank, Washington, DC.
- World Bank (2024). Women, business and the law 2024. Technical report, World Bank, Washington, DC.