Lisbon School of Economics & Management
Universidade de Lisboa

MASTER
MSC ECONOMICS

MASTER´S FINAL WORK
DISSERTATION

JOBLESS GROWTH IN MENA: A DEMOGRAPHIC DECOMPOSITION OF OKUN'S LAW

SEIF SULEIMAN ALI ALORAN

JULY - 2026

Lisbon School of Economics & Management
Universidade de Lisboa

MASTER
MSC ECONOMICS

MASTER´S FINAL WORK
DISSERTATION

JOBLESS GROWTH IN MENA: A DEMOGRAPHIC DECOMPOSITION OF OKUN'S LAW

SEIF SULEIMAN ALI ALORAN

SUPERVISION:
VÍTOR ESCÁRIA

JULY - 2026

To my late father,

my loving mother,

my supportive partner,

and my dear brother,

whose love, encouragement, and unwavering

belief in me made my journey possible.

GLOSSARY

CRVE — Cluster-Robust Variance Estimator

FE — Fixed Effects

GCC — Gulf Cooperation Council

GDP — Gross Domestic Product

ILO — International Labour Organization

ILOEST — ILO Modelled Estimates

LIFO — Last-In, First-Out

MENA — Middle East and North Africa

OECD — Organization for Economic Co-operation and Development

OLS — Ordinary Least Squares

UAE — United Arab Emirates

WDI — World Development Indicators

ABSTRACT, KEYWORDS AND JEL CODES

This dissertation examines whether Okun's Law holds in six Middle East and North African economies: Algeria, Egypt, Jordan, Morocco, Saudi Arabia, and the United Arab Emirates, and whether the responsiveness of unemployment to output growth differs across demographic groups. Using annual data for 2000–2024 from the ILO Modelled Estimates and the World Bank World Development Indicators, first-difference panel regressions relate changes in unemployment rates to real GDP growth. The models include country fixed effects and country-clustered standard errors, with additional robustness checks addressing the small number of country clusters, common year effects, the COVID-19 period, and individual-country influence.

The pooled aggregate Okun coefficient is negative but modest: a one-percentage-point increase in real GDP growth is associated with an approximately 0.14-percentage-point decline in unemployment. However, this aggregate relationship is fragile and loses statistical significance when common year effects are included or the pandemic period is excluded. Country-specific estimates also vary widely, from -0.06 in Saudi Arabia to -0.45 in Egypt, and the pooled estimate is pulled toward the flatter responses of the Gulf economies. The more substantial finding concerns age. Youth unemployment is considerably more responsive to output growth than adult unemployment, with a youth-adult coefficient difference of approximately −0.19. The full-sample age differential survives certain inference and fixed-effect checks but becomes smaller and statistically insignificant when pandemic observations are excluded. The estimated gender difference is smaller and not statistically significant.

A demographic decomposition indicates that adult men, who represent approximately 68% of the labor force on average across the six countries, account for over half of the implied aggregate coefficient. The weak aggregate response therefore appears consistent with labor-force composition: the dominant demographic group has a comparatively modest cyclical response, while more responsive youth groups carry substantially smaller weights.

Keywords: Okun's Law; Jobless growth; Unemployment; Demographic heterogeneity; MENA; Panel data.

JEL Codes: E24; E32; J21; J64; O53; C33.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my supervisor, Professor Vítor Escária, for his guidance, insight, and support throughout the development of this dissertation. I am also grateful to the faculty of ISEG – Lisbon School of Economics and Management, whose teaching over the course of my studies shaped the foundation on which this work rests. Finally, I thank my friends and my family for their encouragement and patience throughout this process.

Artificial Intelligence Statement

In accordance with the School's guidelines on the use of generative artificial intelligence, I declare generative AI tools (large-language-models): Anthropic's Claude and OpenAI's ChatGPT, were used during the preparation of this dissertation in the following ways. First, they assisted with R programming, including drafting, debugging, and refining the code used for the panel estimation, the wild-cluster bootstrap, and the robustness checks. Second, they were used as critical reviewers, to identify weaknesses, analytical gaps, and potential objections, which I then evaluated and addressed. Third, they assisted with rewriting and editing text for clarity and grammar. All analytical decisions, all interpretations of results, and all substantive content are my own; every AI-generated suggestion was independently reviewed, verified, and edited before inclusion, and all empirical results were reproduced and checked against the raw data. The responsibility for the final content of this dissertation, including any errors, rests entirely with me.

1. INTRODUCTION

One of the central expectations in macroeconomics is that economic growth should generate employment. Yet this relationship does not always hold in practice. Economic growth does not always create jobs, and in many economies, output can expand while unemployment remains stubbornly high. This phenomenon, commonly referred to as jobless growth, has become a central concern in the Middle East and North Africa (MENA), where labor market outcomes have often failed to improve despite periods of economic expansion.

One of the most widely used empirical frameworks for analyzing the relationship between growth and unemployment is Okun's Law, which describes the negative relationship between changes in output and changes in unemployment. In many advanced economies this relationship appears relatively stable. Periods of stronger economic growth tend to coincide with falling unemployment, while economic slowdowns lead to rising unemployment. However, empirical evidence suggests that this relationship is considerably weaker in many developing economies, including those in the MENA region.

This raises a question. If economic growth can occur while unemployment remains high, how strong is the relationship between output growth and unemployment in these economies? The issue is particularly relevant in MENA, where unemployment rates, especially among young people, remain among the highest in the world. At the same time, several countries in the region have experienced sustained economic growth over the past decades. The coexistence of growth and persistent unemployment has therefore led researchers to question whether the traditional Okun relationship holds in these economies.

Previous studies examining Okun's Law in MENA countries frequently find weak or statistically insignificant relationships between output growth and unemployment. These findings are often linked to structural features of these labor markets, including the large role of the public sector, labor market rigidities, and the prevalence of informal employment. In addition, demographic pressures associated with large youth populations may further complicate labor market adjustment. As a result, unemployment dynamics in

the region may be influenced by institutional and structural factors that weaken the connection between output growth and employment creation.

A key limitation of much of the existing literature is that it relies on aggregate unemployment measures. Labor markets are heterogeneous, and different groups within the labor force may respond differently to the same economic fluctuations. Evidence from other regions suggests that youth unemployment tends to be more sensitive to economic cycles than adult unemployment. Gender differences may also arise due to differences in sectoral employment patterns and labor force participation. If such heterogeneity exists, aggregate unemployment statistics may mask cyclical responses occurring within particular demographic groups.

This dissertation examines whether demographic heterogeneity helps explain the weak aggregate Okun relationship observed in MENA economies. The central research question is therefore:

Does Okun's Law hold in selected MENA economies, and does the responsiveness of unemployment to output growth differ across demographic groups defined by age and gender?

To address this question the empirical analysis focuses on six MENA economies: Algeria, Egypt, Jordan, Morocco, Saudi Arabia, and the United Arab Emirates. These countries were selected to capture meaningful structural variation in how labor markets are organized across the region. The analysis applies to a first-difference specification of Okun's Law, which relates changes in unemployment rates to real GDP growth. In addition to estimating the aggregate relationship, the study estimates Okun coefficients across demographic series defined by sex and age, including the four mutually exclusive cells formed by their intersection (male youth, male adult, female youth, female adult). Comparing these estimates, and testing the differences between them formally, allows for an assessment of whether cyclical labor market responses differ across segments of the labor force, and whether age or gender is the more important source of any such difference.

The contribution of this dissertation is to connect two strands of literature that have largely developed separately. A large empirical literature examines the validity of Okun's Law across countries, while another body of research highlights demographic differences

in labor market outcomes. Relatively limited work has combined these perspectives within the context of the MENA region. By disaggregating unemployment by age and gender, this study aims to determine whether cyclical labor market sensitivity may be hidden when unemployment is measured only at the aggregate level. The analysis then combines the four cell-specific coefficients with their labor-force shares to assess how demographic composition relates to the aggregate estimate.

The results can be briefly summarized. At the aggregate level, the estimated Okun relationship is negative but weak and specification-sensitive: in the baseline model, a one-percentage-point increase in real GDP growth is associated with an approximately 0.14-percentage-point decline in unemployment, but this association loses statistical significance when common year effects are included or pandemic observations are excluded. The demographic analysis provides stronger full-sample evidence of age heterogeneity: youth unemployment is substantially more responsive to output growth than adult unemployment, with an estimated coefficient difference of approximately −0.19 that is robust to wild-cluster-bootstrap inference and to the exclusion of any individual country from the sample, but weakens and loses significance when the pandemic period is removed. The estimated gender difference is smaller and not statistically significant. Finally, a decomposition indicates that the weak aggregate coefficient is consistent with the composition of the labor force: adult men, roughly two-thirds of the labor force, display a comparatively modest cyclical response and dominate the weighted average, while the strongly cyclical male-youth group carry too little weight to move it.

The remainder of the dissertation is organized as follows. Chapter 2 reviews the literature on Okun's Law, including global empirical evidence, studies focusing on the MENA region, and research on demographic heterogeneity in unemployment dynamics. Chapter 3 describes the data sources and econometric methodology used in the analysis. Chapter 4 presents empirical results for both aggregate and demographic-specific Okun relationships. Chapter 5 concludes by summarizing the main findings and discussing their implications for labor market dynamics in the MENA region.

2. LITERATURE REVIEW

2.1 Theoretical Foundations and Global Evidence

The empirical relationship between aggregate output and unemployment is most commonly associated with Okun (1962), who identified a systematic negative association between cyclical movements in output and changes in the unemployment rate in the post-war United States. In its most cited formulation, Okun suggested that a one percentage point increase in cyclical unemployment was associated with roughly a two to three percent decline in real output relative to potential. Importantly, Okun did not interpret this relationship as a structural law, but rather as an empirical regularity reflecting broader labor market adjustment mechanisms.

Subsequent work has emphasized that the Okun coefficient should be understood as a reduced-form summary statistic. Prachowny (1993) demonstrates formally that the estimated Okun coefficient reflects the combined effects of changes in hours worked, labor force participation, and labor productivity, in addition to employment itself. When output rises, firms may initially adjust along the intensive margin by increasing hours per worker before expanding headcount. At the same time, improved labor market conditions can induce discouraged workers to re-enter the labor force, while higher capacity utilization may raise measured productivity. The unemployment rate therefore serves as a proxy for cyclical slack that captures multiple margins of adjustment simultaneously. The magnitude of the Okun coefficient depends critically on how strongly these margins respond to output fluctuations.

Empirically, the relationship has been estimated using two principal approaches: the "gap" model and the "first-difference" model. In developing-country contexts where potential output and natural rates of unemployment are difficult to estimate reliably due to structural change and data limitations, the first-difference specification is often preferred, as it avoids reliance on unobservable trend components. The gap specification relates deviations of actual output from potential output to deviations of unemployment from its natural rate. In contrast, the first-difference model directly links changes in unemployment to output growth, avoiding the need to estimate unobservable natural rates. Regardless of specification, a consistent finding in the literature is that the Okun coefficient varies across countries and over time.

Lee (2000), using OECD data, finds that although the negative relationship between output and unemployment is generally present, the magnitude of the coefficient differs substantially across economies and exhibits structural instability. Sögner and Stiassny (2002) examine the structural stability of Okun's Law using Bayesian methods and time-varying parameter (Kalman filter) methods. While they find no strong evidence of discrete structural breaks, their Kalman filter estimates reveal significant time variation in the Okun coefficient for several OECD countries, including France, Germany, and the Netherlands, whereas the relationship remains comparatively stable in countries such as the United States and Austria. Specifically, they show that much of this variation is driven by changes in the responsiveness of employment to output, which they link to differences in labor market protection and labor hoarding behavior. These findings underscore that Okun's Law is best interpreted as an empirical regularity rather than a structural constant, with its magnitude reflecting evolving labor market institutions and macro adjustment mechanisms over time.

The asymmetry of the relationship has also received considerable attention in recent literature. Crespo Cuaresma (2003) show that Okun's coefficient is significantly larger during recessions than during expansions, indicating that unemployment responds more strongly to negative output shocks than to positive ones. Their results suggest that labor market adjustment is state-dependent, with firms reducing employment more aggressively during downturns while relying more on intensive margin adjustments during recoveries. Such findings reinforce the view that labor market institutions and adjustment costs shape the strength and nonlinear behavior of the output–unemployment link.

Despite variation in magnitude and stability, Okun's Law has proven to be broadly robust across advanced economies. As Ball et al. (2017, p. 1439) note, "It is rare to call a macroeconomic relationship a 'law.' Yet we believe that Okun's Law has earned its name." Analyzing 20 advanced countries, they conclude that the relationship remains broadly stable within countries over time. They argue that apparent "jobless recoveries" following the Global Financial Crisis are largely consistent with the traditional Okun framework once the slow pace of output recovery is taken into account. At the same time, they document significant cross-country heterogeneity. For example, Spain exhibits a relatively large Okun coefficient, reflecting the responsiveness of unemployment to

output fluctuations, while Japan's coefficient is substantially smaller, consistent with institutional features such as labor hoarding and long-term employment relationships. These differences highlight the importance of labor market structure in mediating cyclical adjustment.

Recent literature has also emphasized that aggregate Okun coefficients may conceal important heterogeneity within economies. Zanin (2014), using harmonized OECD data, shows that the responsiveness of unemployment to output varies systematically across age groups. Youth unemployment (ages 15–24) tends to be substantially more cyclical than adult unemployment. This finding is consistent with earlier work by Bell and Blanchflower (2011), who document the pronounced volatility of youth unemployment across business cycles. O'Higgins (1997) attributes this pattern to the role of aggregate demand conditions in driving youth unemployment. Young workers are disproportionately affected during downturns because firms respond to recessions by freezing recruitment, which primarily impacts new labor market entrants. Bell and Blanchflower (2011) relate this to standard last-in-first-out (LIFO) employment adjustment practices. Zanin (2014) also reports gender differences, with male unemployment typically displaying greater cyclical sensitivity than female unemployment, partly reflecting sectoral segregation across industries with differing cyclical exposure.

Taken together, the evidence from advanced economies suggests that while Okun's Law remains a useful empirical framework, its magnitude depends critically on institutional features, labor market structure, and demographic composition. These insights motivate examination of whether similar patterns hold in structurally different regions.

2.2 Okun's Law in the MENA Region

Evidence from the Middle East and North Africa (MENA) regions presents a more ambiguous picture. Moosa (2008), examining Algeria, Egypt, Morocco, and Tunisia, finds that estimated Okun coefficients are generally statistically insignificant, suggesting that output growth does not systematically translate into reductions in unemployment. He offers three explanations for this weak relationship. First, unemployment in these countries appears largely structural and/or frictional rather than cyclical, meaning that

skill mismatches and job-search inefficiencies limit the responsiveness of unemployment to output growth. Second, labor markets are characterized by rigidity, partly due to the dominant role of the public sector as the main employer, which reduces hiring and firing flexibility and weakens cyclical adjustment. Third, economic structures in these countries are often concentrated in government-led or capital-intensive sectors, such as oil in Algeria. As a result, growth in these sectors does not generate substantial employment gains. Taken together, these factors imply that unemployment in the region cannot be explained primarily by insufficient aggregate demand. Kreishan (2011), focusing on Jordan using annual data from 1970 to 2008, reaches a similar conclusion: despite periods of sustained economic growth, unemployment remained persistently high and statistically unrelated to output growth.

This institutional rigidity is further elaborated in the broader MENA labor market literature. Assaad and Barsoum (2019) document that public employment has historically been central to the region's social contract, with governments expanding public sector hiring not purely in response to market conditions but often in response to political pressures and social stability concerns. They show, for example, that public employment increased following the Arab Spring in several countries, reflecting social and political considerations rather than output dynamics. Similarly, World Bank (2013) emphasizes the prevalence of public-sector wage premia and queuing behavior among educated youth, whereby individuals may remain unemployed while awaiting government jobs. Such institutional features imply that a significant segment of employment is weakly tied to short-run business cycle fluctuations, thereby reinforcing Moosa's argument that aggregate unemployment in MENA may exhibit limited cyclical responsiveness to GDP growth.

The sectoral composition of growth also matters. In several MENA economies, particularly oil-exporting countries, output fluctuations are driven largely by external commodity price movements and capital-intensive extraction sectors. As emphasized in the resource-curse literature (Sachs & Warner, 2001), natural resource abundance can crowd out tradable sectors and weaken export-led growth. In economies where output is concentrated in capital-intensive extraction industries, GDP growth may not translate into broad-based employment expansion. Distinguishing between oil-exporting and non-oil

MENA countries is therefore relevant when interpreting cross-country variation in Okun coefficients.

An additional factor that may attenuate the aggregate output–unemployment relationship is the prevalence of informality and labor market segmentation. La Porta and Shleifer (2014) documents that informal firms in developing economies operate alongside formal firms but differ substantially in productivity, scale, and managerial human capital. They emphasize that informal enterprises are typically small, low-productivity, and largely disconnected from the formal sector, with limited transition into formality. In labor markets where informal employment constitutes a significant share of total employment, cyclical adjustment may therefore occur through shifts between formal employment, informal activity, and inactivity rather than through movements into open unemployment.

Evidence from the World Bank (2013) indicates that informality remains widespread across MENA labor markets, with substantial shares of workers engaged in self-employment or informal wage work. If negative output shocks primarily trigger reallocation across employment categories instead of outright job destruction recorded as unemployment, the measured unemployment rate may respond weakly to GDP fluctuations. In such settings, the estimated Okun coefficient may understate the true cyclical sensitivity of labor demand, reflecting measurement and segmentation effects rather than the absence of output-driven adjustment.

Female labor force participation introduces a further complication. MENA countries consistently exhibit some of the lowest female participation rates globally, as documented by International Labour Organization (ILO) reports and the World Bank (2013). Participation decisions may be sensitive to economic conditions and household dynamics. Where participation is relatively elastic, women may withdraw from the labor force during downturns or enter during recoveries without immediate employment absorption. Such movements affect measured unemployment differently from contexts with more stable participation rates. As a result, aggregate unemployment statistics may not fully capture cyclical labor market slack, potentially contributing to weak Okun estimates at the national level.

Finally, demographic structure plays an important role in shaping labor market dynamics in the region. The OECD (2016) documents that MENA countries are characterized by a large youth population and persistently high youth unemployment rates compared to OECD and global averages. The report highlights the difficulties young people face in transitioning from education to stable employment, including limited private-sector job creation and continued preference for public sector employment. In labor markets where youth are disproportionately concentrated among new entrants and in more precarious or transitional forms of employment, cyclical shocks may affect demographic groups unevenly. If some young individuals remain unemployed while queuing for public sector positions and others circulate between short-term private employment and inactivity, aggregate unemployment dynamics may mask heterogeneous cyclical responses across age groups.

2.3 Aggregation, Heterogeneity, and the Case for Demographic Disaggregation

The combined evidence suggests that the weak aggregate Okun coefficients reported for MENA economies may partly reflect the aggregation of heterogeneous demographic groups rather than a uniform absence of cyclical labor demand. In highly segmented labor markets characterized by public sector queuing, informality, and low female participation, different demographic groups may respond differently to output fluctuations. Aggregate unemployment rates, which combine these heterogeneous groups, may therefore understate cyclical sensitivity for specific cohorts. This possibility provides a direct empirical justification for estimating demographic-specific Okun coefficients.

The demographic disaggregation approach proposed by Zanin (2014) provides a framework for addressing this issue. By estimating separate Okun coefficients for different age and gender groups, it becomes possible to assess whether cyclical responsiveness varies systematically across segments of the labor force. If youth or prime-age men employed in the private sector exhibit stronger output sensitivity than other groups, aggregate insignificance may conceal meaningful within-group dynamics.

Despite the extensive literature on Okun's Law in advanced economies and the growing body of work on labor market structure in MENA, relatively little research has combined these strands by estimating demographic-specific Okun relationships within the region. Existing MENA studies have largely relied on aggregate unemployment rates.

This gap motivates the present study, which applies a first-difference Okun framework to selected MENA economies and disaggregates unemployment by age and gender. In doing so, the analysis seeks to determine whether cyclical labor market sensitivity is unevenly distributed across demographic groups and obscured by structural features of regional labor markets.

3. METHODOLOGY

The preceding discussion suggests that the weak Okun relationship observed in many MENA economies is unlikely to reflect an absence of cyclical labor demand. Rather, the literature points to the way labor markets in the region are structured. Public sector dominance, segmented employment systems, and persistently low female participation all imply that aggregate unemployment may not respond uniformly to changes in output. If adjustment differs across groups, aggregate estimates may understate the cyclical responsiveness of the more sensitive segments of the labor market, even where the aggregate response itself is modest. The empirical strategy adopted here is therefore built around a simple question: whether, and in what way, the demographic composition of the labor force shapes the apparent weakness of Okun's Law in MENA.

3.1 Country Sample and Data

The analysis focuses on six MENA economies: Algeria, Egypt, Jordan, Morocco, Saudi Arabia, and the United Arab Emirates. The selection is designed to capture variation along two dimensions that the preceding discussion identifies as central to labor market adjustment in the region: the degree to which output is driven by hydrocarbon and capital-intensive sectors, and the extent to which measured unemployment reflects queuing for public-sector or nationalized employment rather than cyclical private labor demand.

Saudi Arabia and the United Arab Emirates are included as Gulf Cooperation Council (GCC) economies characterized by large migrant workforces and extensive nationalization policies. Because the ILO Modelled Estimates (ILOEST) unemployment series is not nationality-specific, the estimates for these economies should be interpreted as aggregate resident labor-market indicators rather than as direct measures of national unemployment. In both countries, nationalization policies (Saudization and Emiratization) have shaped domestic labor demand in ways that are not straightforwardly

tied to aggregate output fluctuations. The two are not substitutes, however. The Emirati economy is considerably more diversified beyond hydrocarbons, particularly through Dubai's trade and services sectors, and female labor force participation has historically been higher than in Saudi Arabia. Including both introduces variation within the GCC in economic diversification and female participation; the common-slope specification does not separately identify these differences, but the country-specific estimates in Appendix Table XI provide descriptive evidence on them.

Egypt and Jordan represent non-GCC economies with strong public-sector employment legacies, in which the public sector has historically played a central role in employment absorption alongside a substantial private and informal sector. In both cases, the literature documents persistent queuing behavior among educated youth awaiting public employment, which may weaken the cyclical link between output and measured unemployment (Assaad & Barsoum, 2019).

Algeria and Morocco are included as Maghreb economies with broadly comparable demographic structures but different productive bases. Algeria's output is heavily driven by hydrocarbon rents that flow largely through state channels, while Morocco has a more diversified economy with stronger private-sector manufacturing and services linkages to European markets. Their contrast ensures the non-GCC part of the panel spans both hydrocarbon-intensive and more diversified productive structures, although the design does not formally test the effect of hydrocarbon dependence.

Unemployment data are drawn from the ILO Modelled Estimates (ILOEST, November 2025 release; International Labour Organization, 2025) and consist of annual unemployment rates disaggregated by sex and age group. The data cross-classify three sex categories (total, male, female) with three age bands (15+, 15–24, and 25+), yielding nine series per country, including the four mutually exclusive cells (male youth, male adult, female youth, female adult) used in the demographic analysis. The ILOEST series provide internationally comparable, harmonized estimates constructed from national labor force surveys where available and imputation procedures where direct observations are missing. This yields a balanced panel across the six countries and across all demographic cells, which is essential for comparing estimated Okun coefficients across groups. The use of modelled rather than directly observed data introduces some

measurement concerns, particularly in country-years where underlying national data are sparse; these are discussed further in Section 3.5. Real GDP growth rates are obtained from the World Development Indicators (World Bank, 2026; series NY.GDP.MKTP.KD.ZG).

The sample period runs from 2000 to 2024. This window captures substantial cyclical variation, including the 2008 global financial crisis, the 2011 Arab Spring and its macroeconomic aftermath, the 2014–2016 oil price collapse, and the COVID-19 shock. ILOEST projections for 2025 onward are excluded from estimation.

3.2 Empirical Model

The empirical framework follows the first-difference specification of Okun's Law. The focus is on the short-run relationship between output growth and changes in unemployment, rather than on deviations from unobservable trends. This is a deliberate choice. In the MENA context, estimating potential output or natural rates of unemployment would introduce additional noise without improving identification.

The baseline specification is

$$\Delta u_{it} = \alpha_i + \beta g_{it} + \varepsilon_{it}, \quad (1)$$

where Δu_{it} denotes the annual change in the unemployment rate and g_{it} is real GDP growth. The coefficient β captures the extent to which unemployment responds to changes in economic activity, on average across the panel.

A dynamic specification including lagged GDP growth is also estimated to allow for delayed labor market adjustment

$$\Delta u_{it} = \alpha_i + \beta_0 g_{it} + \beta_1 g_{i,t-1} + \varepsilon_{it}, \quad (2)$$

where $g_{i,t-1}$ is real GDP growth in the previous year. The coefficient β_0 captures the contemporaneous response of unemployment to output growth, while β_1 captures the delayed response operating with a one-year lag. The cumulative long-run Okun coefficient is given by ($\beta_0 + \beta_1$), which measures the total adjustment of unemployment to output growth once lagged effects are taken into account.

The central concern of this dissertation is whether the aggregate Okun relationship masks differences across demographic groups. This is addressed in two complementary

steps. First, the baseline specification is estimated separately for each demographic series; total, male, female, youth (15–24), and adult (25+). These are available for all six countries across the full sample period. The group-specific estimates provide a transparent first reading of the data: each coefficient can be interpreted directly as the Okun coefficient for a given segment of the labor force, and differences in magnitude across groups offer initial evidence of heterogeneity.

Group-specific regressions alone, however, cannot establish whether the observed differences are statistically meaningful, since estimating each series in isolation yields no estimate of the standard error of the difference between coefficients. To test heterogeneity formally, the demographic series are stacked and a pooled specification with an interaction term is estimated

$$\Delta u_{igt} = \alpha_{ig} + \beta g_{it} + \gamma(g_{it} \times D_g) + \varepsilon_{igt}, \quad (3)$$

where α_{ig} denotes country-by-group fixed effects, allowing each demographic group within each country to have its own average change in unemployment. D_g equals one for the comparison group and zero for the reference group. β is the Okun coefficient for the reference group, and γ measures the difference between the comparison and reference groups; the comparison group's coefficient is therefore β+γ.

The coefficient γ is therefore the central parameter of interest: it captures, in a single estimate with its own standard error, whether and by how much cyclical responsiveness differs across the demographic dimension. The null hypothesis of demographic homogeneity (H_0: $\gamma = 0$) is assessed using a Wald test.

This framework also allows the two dimensions to be compared directly. The interaction is estimated separately for the age dimension (youth versus adult) and the gender dimension (male versus female) and is extended to the four mutually exclusive demographic cells defined by the intersection of age and gender (male youth, male adult, female youth, female adult). Comparing the estimated interaction effects across these specifications makes it possible to assess which dimension (age or gender) is associated with the larger and more robust difference in cyclical responsiveness.

3.3 Empirical Strategy

The model is estimated with country fixed effects. Because the dependent variable is the first difference of the unemployment rate, time-invariant differences in unemployment level that reflect persistent structural features such as the size of the public sector or the extent of informal employment, are already removed by differencing. The country fixed effects allow each economy to have its own average change in unemployment over the sample, corresponding to a country-specific linear drift in the unemployment level that is unrelated to output growth.

The aggregate Okun relationship is estimated first, using total unemployment rates, to establish a benchmark. The demographic specifications described above are then estimated and compared against it. If the aggregate coefficient is weak while particular demographic coefficients are stronger, this would support the central argument that aggregation masks underlying cyclical responses.

To account for common shocks, the specification is augmented with time effects as a robustness check. This is particularly relevant in a region where output fluctuations are often influenced by global factors such as commodity price movements or external demand conditions. Several econometric issues arise in this setting.

First, unobserved heterogeneity across countries is addressed through fixed effects. This is not merely a technical correction but reflects the fact that labor market institutions, sectoral composition, and participation patterns differ systematically across the sample.

Second, standard errors are clustered by country to account for heteroskedasticity and within-country serial correlation. With only six clusters, however, cluster-robust inference is itself unreliable, since its validity rests on a large number of clusters. All inference therefore uses a t-distribution with $G - 1 = 5$ degrees of freedom; the wild-cluster-bootstrap p-values reported in Section 4.5 are obtained from the empirical bootstrap distribution. Significance levels are therefore interpreted with caution, and the analysis emphasizes the magnitude and consistency of estimated coefficients across specifications rather than relying solely on conventional significance thresholds.

Third, the estimated relationship is inherently a reduced form. Changes in unemployment reflect not only employment adjustment but also shifts in participation, hours worked, and movements into informal activity. This is particularly relevant in

MENA economies, where labor market adjustment does not always occur through changes in measured unemployment.

Finally, the issue of endogeneity must be addressed. In principle, unemployment may also influence output, raising concerns about simultaneity. In this context, short-run movements in GDP partly reflect factors external to contemporaneous labor-market conditions, including commodity price fluctuations, external demand, and global macroeconomic shocks. This is more plausible for the oil-exporting economies (Algeria, Saudi Arabia, and the United Arab Emirates), where output is heavily driven by hydrocarbon revenues, than for the more domestic-demand-driven economies of Egypt, Jordan, and Morocco.

The estimates should therefore be interpreted as reduced-form associations between output growth and changes in unemployment: although GDP growth may partly reflect external shocks, simultaneity and omitted variables cannot be fully ruled out. The robustness checks in Section 4.5 examine sensitivity to common shocks, influential periods, few-cluster inference, and individual countries; simultaneity and time-varying omitted variables remain limitations. All models are estimated in R (version 4.5.1) using the fixest package documented by Bergé et al. (2026).

3.4 Hypotheses

The empirical analysis tests three hypotheses. The first concerns the aggregate relationship:

$$H_1: \beta < 0.$$

That is, higher real GDP growth is associated with a fall in the unemployment rate, consistent with Okun's Law.

The second concerns age. Drawing on the "last-in, first-out" dynamics documented in the literature, in which younger workers on more precarious contracts bear a disproportionate share of cyclical adjustment, youth unemployment is expected to be more responsive to output than adult unemployment:

$$H_2: \gamma_{age} \neq 0 \text{ with the prior expectation } \gamma_{age} < 0.$$

The formal test is a two-sided Wald test of $H_0: \gamma_{age} = 0$. The directional prior is that youth unemployment is more responsive, implying a negative interaction coefficient.

The third concerns gender and is treated as exploratory. In some advanced economies men are more cyclically sensitive owing to their concentration in cyclical sectors such as manufacturing and construction. In the MENA context, however, low female labor-force participation and the segmentation of employment between the public and private sectors make the expected direction of any gender difference less clear. The gender comparison is therefore examined without committing to a directional prior:

$$H_3: \gamma_{gender} \neq 0.$$

Because the expected direction is ambiguous in the MENA context, the gender hypothesis is non-directional.

3.5 Limitations

First, although the ILOEST series extends back to 1991, the empirical sample begins in 2000; earlier country-years may rely more heavily on imputation, particularly where underlying labor-force survey coverage is incomplete, and the post-2000 window also coincides with the period of greatest cyclical variation in the sample. This yields twenty-five annual observations per country, or twenty-four after first-differencing. Second, the ILOEST series are modelled rather than directly observed for every country-year. Where underlying national labor force data are sparse, the ILO applies imputation procedures to ensure comparability, which strengthens cross-country consistency but may smooth short-run fluctuations in the underlying series. This could attenuate estimated cyclical responsiveness in affected country-years.

Third, unemployment rates may not fully capture labor market slack in economies where adjustment occurs through informality or withdrawal from the labor force, a concern that is particularly relevant for the non-GCC cases in the sample. Finally, while the selected countries capture meaningful variation within MENA along the dimensions identified in the literature (hydrocarbon dependence, public sector role, and labor market segmentation), the results should not be interpreted as representative of the region.

4. EMPIRICAL ANALYSIS

4.1 Descriptive Analysis

This section establishes the empirical basis before the estimation that follows. It proceeds from the cross-country picture of unemployment and growth, to the

demographic structure of the labor force, and finally to a preliminary visual reading of the output–unemployment relationship itself. The aim is descriptive: to set out the patterns that motivate the disaggregated analysis, and to identify the questions that the regressions in Sections 4.2 to 4.5 are designed to answer.

Table I: Summary statistics by country, 2000–2024

Country	Unemployment rate (%)				Change in U (pp)		GDP growth (%)		
	Mean U	SD U	Min U	Max U	Mean ΔU	SD ΔU	Mean g	SD g	N
Algeria	14.4	5.8	9.8	29.8	-0.75	1.79	3.1	2.2	25
Egypt	9.9	2.1	6.5	13.2	-0.09	1.06	4.3	1.6	25
Jordan	15.0	2.4	11.9	19.8	0.12	1.19	4.0	2.5	25
Morocco	10.0	1.3	8.9	13.6	-0.19	0.70	3.7	2.9	25
Saudi Arabia	5.6	0.8	3.5	7.7	-0.04	0.69	3.9	3.8	25
United Arab Emirates	2.6	0.5	1.6	4.3	0.00	0.55	3.9	4.4	25
Sample average	*9.6*	*2.2*	—	—	*-0.16*	*1.00*	*3.8*	*2.9*	*150*

Notes: Sources: ILO Modelled Estimates (Nov 2025); World Bank WDI. U is the aggregate unemployment rate (15+); ΔU its annual percentage-point change; g is real GDP growth. N is annual observations per country (25 for levels; 24 for first-differenced ΔU) or pooled observations (average row). Sample average is the unweighted mean of the six countries.

Table I summarizes unemployment and real GDP growth for each of the six economies over 2000–2024. Two features stand out. First, average unemployment varies widely across the sample, from 2.6% in the United Arab Emirates and 5.6% in Saudi Arabia to 14.4% in Algeria and 15.0% in Jordan. This gap reflects the structural differences set out in Chapter 3: the GCC economies record low unemployment rates, while the Maghreb and public-sector-dominant economies carry persistently higher rates. Second, the volatility of unemployment differs just as sharply. The standard deviation of the annual change in unemployment ranges from 0.55 percentage points in the UAE to 1.79 in Algeria, with the high-unemployment economies also exhibiting the most volatile year-to-year movements. Average GDP growth, by contrast, is more uniform across the sample, clustering between 3.1% and 4.3%, though its variability differs markedly across countries. Growth volatility is highest in Saudi Arabia and the UAE, where the standard deviation of growth reaches 3.8 and 4.4 respectively, while Algeria's measured growth volatility is more moderate at 2.2. Finally, the mean change in unemployment is small in most countries, consistent with the absence of a strong long-run trend and with the first-difference framing adopted here. Algeria is the exception, with an average annual decline of 0.75 percentage points, indicating a country-specific downward trend over the period.

Figure 1: Unemployment and GDP growth by country, 2000–2024

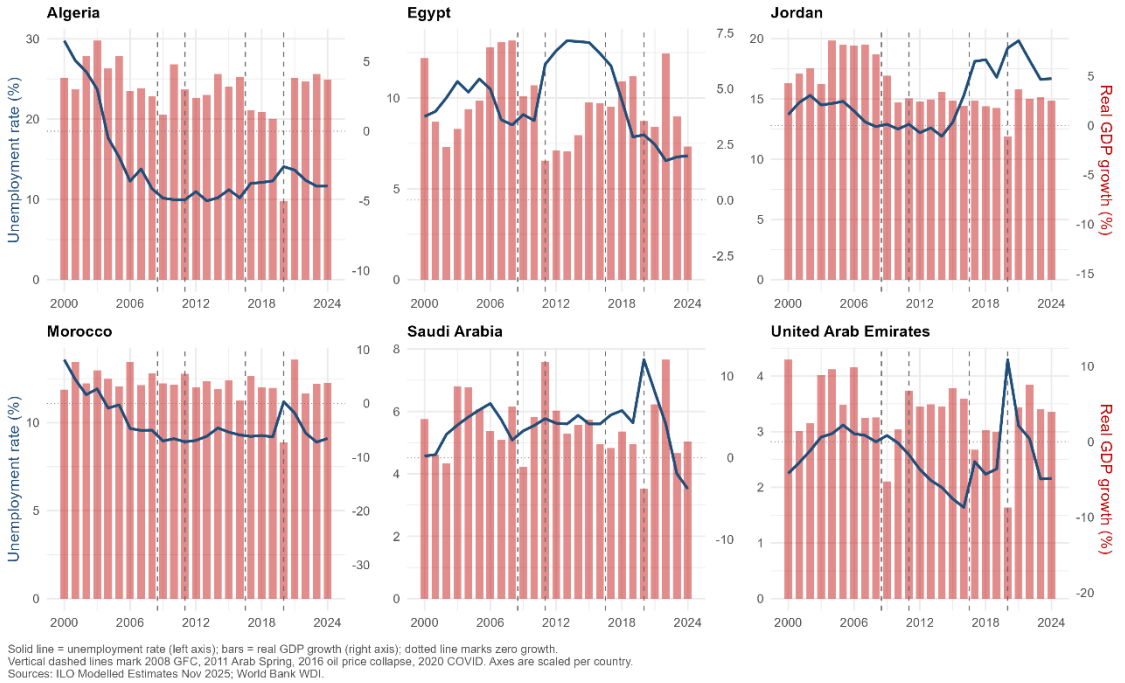

Figure 1 plots the two series together for each country, with the unemployment rate measured against the left axis and real GDP growth against the right. The figure makes the central puzzle of the dissertation visible. In several economies, sustained periods of positive growth coincide with unemployment rates that remain high or fall only modestly, while the contractions associated with the 2008 global financial crisis, the 2011 Arab Spring, the 2014–2016 oil price collapse, and the 2020 COVID-19 shock are not always matched by proportionate movements in unemployment. The per-country trajectories also confirm that the high-unemployment economies are those with the most pronounced swings, while the GCC economies display comparatively flat unemployment paths despite substantial variation in output. Because Figure 1 plots the level of unemployment, rather than its change, against growth, it is descriptive context rather than a direct depiction of the Okun relationship, which is examined in first differences below. It nevertheless makes the broader jobless-growth puzzle visible, with output and unemployment moving together only loosely at the aggregate level.

Table II: Unemployment dynamics by gender and age, pooled across countries

Group	Unemployment rate (%)		Change in U (pp)	
	Mean U	SD U	Mean ΔU	SD ΔU
By gender				
Male	8.06	5.29	-0.18	1.14
Female	16.61	7.16	-0.09	1.69
By age				
Youth (15–24)	23.92	10.16	-0.20	2.66
Adult (25+)	6.52	3.81	-0.07	0.91
Aggregate				
Total	*9.57*	*5.22*	*-0.16*	*1.10*

Notes: Sources: ILO Modelled Estimates (Nov 2025); World Bank WDI. Pooled across the six countries, 2000–2024 (N = 150 country-years per group; 144 for ΔU statistics).

Table II turns to the demographic structure of unemployment, pooling across countries to summarize the broad margins. The disparities are pronounced. Female unemployment averages 16.6%, more than double the male rate of 8.1%, and youth unemployment, at 23.9%, is almost four times the adult rate of 6.5%. The same groups are also more volatile: the standard deviation of the change in unemployment is 2.66 percentage points for youth against 0.91 for adults, and 1.69 for women against 1.14 for men. This pattern (higher and more variable unemployment among youth and women) is the first descriptive indication that the aggregate rate may average over groups with different cyclical behavior. Greater year-to-year volatility is consistent with, but does not by itself establish, stronger sensitivity to output fluctuations. These gaps are persistent features of the sample rather than features of particular years.

Table III: Unemployment dynamics and labor-force size, four-way demographic cells

Demographic cell	Unemployment rate (%)		Change in U (pp)		
	Mean U	SD U	Mean ΔU	SD ΔU	Share of LF (%)
Male, Adult (25+)	5.40	3.89	-0.10	0.91	68.3
Male, Youth (15–24)	20.96	9.92	-0.22	2.84	13.2
Female, Adult (25+)	11.96	5.28	0.04	1.60	14.9
Female, Youth (15–24)	35.44	16.09	-0.06	4.28	3.6

Notes: Sources: ILO Modelled Estimates (Nov 2025); World Bank WDI. Share of LF = mean share of the total 15+ labor force, pooled across the six countries, 2000–2024. The four cells are mutually exclusive and exhaustive; shares sum to ~100%.

These differences in the level and volatility of unemployment would matter little for the aggregate if the groups concerned were small. Table III therefore reports, alongside the unemployment statistics, each group's share of the total labor force, pooled across the six countries. The four mutually exclusive cells defined by the intersection of

age and gender are highly unequal in size: adult men account for roughly 68% of the labor force, adult women for around 15%, young men for 13%, and young women for under 4%. The implication is important for interpreting the aggregate Okun relationship. The groups with the most distinctive unemployment dynamics (youth and women) are precisely those with the smallest labor-force weights, so their behavior could be heavily diluted in the aggregate, which is dominated by adult men. This is the mechanism through which aggregation may conceal demographic group-specific cyclical responses, and it provides the descriptive foundation for the disaggregated estimation that follows.

Table IV: Labor-force shares by demographic cell and country

	Male		Female	
Country	**Adult**	**Youth**	**Adult**	**Youth**
Algeria	68.9	15.1	13.2	2.8
Egypt	64.3	13.7	16.7	5.3
Jordan	65.4	18.0	12.5	4.1
Morocco	61.3	14.1	19.8	4.8
Saudi Arabia	76.9	9.0	12.1	1.9
United Arab Emirates	72.9	9.2	15.2	2.6

Notes: Source: ILO Modelled Estimates (Nov 2025), economically active population. Each cell is the country's mean share of its total 15+ labor force over 2000–2024. Adult = 25+, Youth = 15–24. Rows sum to 100%.

This compositional structure is not uniform across the sample. Table IV reports the same four labor-force shares by country, revealing substantial heterogeneity along the lines of the typology developed in Chapter 3. The female share of the labor force is markedly lower in the GCC economies than in Egypt or Morocco, where female participation is higher: young women constitute under 2% of the labor force in Saudi Arabia and around 2.6% in the UAE. Adult men correspondingly dominate the GCC labor force to an even greater degree, accounting for 77% in Saudi Arabia. These differences indicate that the extent of any aggregation masking is itself country-specific, being most severe where the labor force is most heavily concentrated in a single demographic group.

Figure 2: Okun relationship by demographic cell, pooled across countries

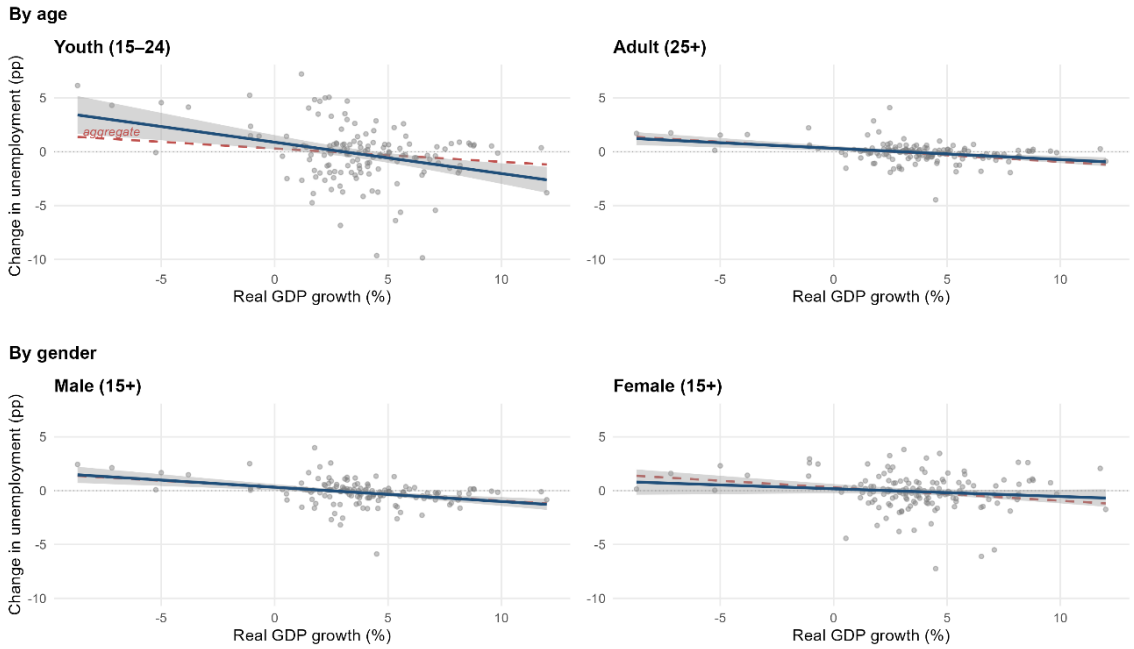

Each point is a country-year. Blue line = group's fitted Okun slope (95% CI shaded). Dashed line = aggregate (15+) slope, shown for reference.
Y-axes are shared across panels. Sources: ILO Modelled Estimates Nov 2025; World Bank WDI.

Figure 2 provides a visual reading of the Okun relationship across demographic groups, plotting the annual change in unemployment against real GDP growth for each cell, pooled across countries, with the aggregate slope shown for reference. The aggregate relationship is negative but shallow, consistent with the weak coefficients reported in the literature. The disaggregated panels, however, already suggest heterogeneity. The youth slope is visibly steeper than the adult slope, indicating greater cyclical sensitivity among younger workers, consistent with the "last-in, first-out" dynamic discussed in Chapter 2. By contrast, the adult and male slopes nearly coincide with the aggregate, which is expected given that adult men dominate the labor force and therefore drive the aggregate relationship. The female slope is flatter and weakly defined, in line with the more ambiguous cyclical behavior of female unemployment in a setting of low and variable participation. These visual differences are suggestive only; whether they are statistically meaningful is the question addressed formally in Sections 4.3 and 4.4.

Figure 3: Stability of the aggregate Okun relationship across candidate break points

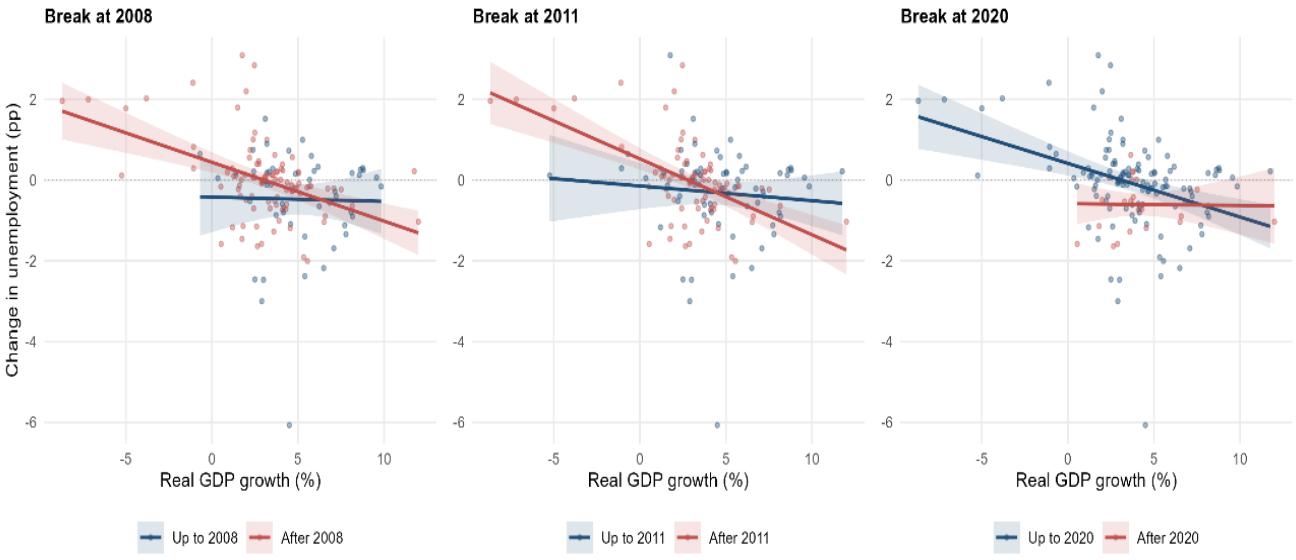

Pooled aggregate (15+) country-years, split before and after each candidate break. Lines are period-specific fitted Okun slopes (95% CI shaded). Axes shared across panels. The post-2020 slope is estimated on few observations and is correspondingly imprecise. Sources: ILO Modelled Estimates Nov 2025; World Bank WDI.

Finally, Figure 3 examines whether the aggregate relationship is stable over time, plotting the pooled output–unemployment scatter split before and after each of three candidate break points: the 2008 financial crisis, the 2011 Arab Spring, and the 2020 pandemic. The relationship does not appear time-invariant. For the 2008 and 2011 splits, the fitted slope is markedly steeper in the later period, suggesting that the aggregate Okun relationship strengthened over the sample rather than remaining constant; a pattern that is consistent with the time-variation documented by Lee (2000) and Sögner and Stiassny (2002). The 2020 split shows the opposite pattern, but the post-2020 slope rests on only four years of data and is correspondingly imprecise, so little weight should be placed on it. Because these comparisons are descriptive and pooled across countries, they cannot establish a formal structural break; they indicate only that the assumption of a constant coefficient warrants scrutiny, a possibility revisited in the robustness analysis of Section 4.5.

Taken together, the descriptive evidence sketches a coherent set of questions. The aggregate output–unemployment relationship is weak, as the regional literature reports, yet the labor force is composed of demographic groups that differ sharply in both the level and the volatility of their unemployment, and that carry very unequal weights in the aggregate. Whether these patterns reflect statistically meaningful differences in cyclical

responsiveness, and which dimension dominates, cannot be settled descriptively. The formal analysis begins with the aggregate Okun relationship in Section 4.2.

4.2 Aggregate Okun Estimates

Table V reports estimates of the aggregate Okun relationship for total unemployment (ages 15+), pooling the six countries with country fixed effects and standard errors clustered by country. Because the dependent variable is the annual change in unemployment, the estimation uses the 24 first-differenced observations per country over 2001–2024 (144 country-year observations in the baseline). The baseline specification yields an Okun coefficient of −0.135, statistically significant at the 5% level: a one-percentage-point increase in real GDP growth is associated with a fall in the unemployment rate of roughly 0.14 percentage points. The coefficient carries the expected negative sign, providing support for the first hypothesis (H_1) that the aggregate relationship is negative.

Table V: Aggregate Okun estimates

	Dependent variable: ΔU	
Variable	Baseline	Dynamic
GDP growth	-0.135*	-0.130*
	(0.041)	(0.037)
GDP growth (t−1)	—	-0.030
		(0.035)
Country fixed effects	Yes	Yes
Observations	144	144
Within R^2	0.146	0.153

Notes: Country fixed effects are included, and standard errors clustered by country are reported in parentheses The cumulative coefficient ($\beta_0 + \beta_1 = -0.159$, SE 0.064) is not significant at the 5% level (Wald test of $\beta_0 + \beta_1 = 0$, $p = 0.056$). * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Two features of this estimate are notable. First, in the baseline pooled specification a statistically significant negative association emerges. This contrasts with part of the regional literature, which reports aggregate Okun coefficients for MENA economies that are statistically insignificant (Moosa, 2008; Kreishan, 2011). Once the six countries are pooled and country fixed effects absorb country-specific average changes in unemployment, a discernible cyclical association emerges. Second, the magnitude is small. At −0.135, the coefficient is well below the advanced-economy average of -0.32 for the comparable annual first-difference specification, though not outside the range:

(Ball, Leigh, and Loungani, 2017) report country estimates ranging from -0.079 in Japan to -0.802 in Spain, with Austria at -0.135. One thing to note is that the pooled estimate also weights countries by the variance of their growth, so the two Gulf economies contribute roughly 60% of the identifying variation; the unweighted mean of the country-specific slopes in Appendix Table XI is -0.21. The aggregate relationship in the sample is therefore best characterized as present but weak. A small coefficient indicates weak transmission from output growth to measured open unemployment; it is not evidence that employment or labor demand is acyclical, since adjustment may occur through participation, informality, hours, or, in the case of GCC economies, the departure of expatriate workers. As the estimates are reduced-form, the coefficient is interpreted as a short-run association rather than a structural causal effect.

The dynamic specification, which adds lagged GDP growth, is reported in the second column. The contemporaneous coefficient is essentially unchanged at −0.130 and remains significant, while the lagged coefficient is small (−0.030) and statistically insignificant. The cumulative long-run coefficient, $\beta_0 + \beta_1 = -0.159$, is not statistically significant at the 5% level. The estimates therefore provide little evidence of an additional one-year response beyond the contemporaneous relationship.

These aggregate estimates establish the benchmark against which the demographic results are assessed. In the baseline specification the relationship is negative and statistically significant, though weak in magnitude; Section 4.5 shows this significance is not robust. Crucially, a weak aggregate coefficient need not imply that no group responds strongly to output. Because the aggregate is a labor-force-weighted composite, it is dominated by the largest group (adult men) whose own cyclical response may be modest. The central empirical questions that follow are therefore twofold: whether cyclical responsiveness differs across demographic groups (Section 4.3), and, if so, how much those group-level differences actually contribute to the aggregate relationship given the very unequal labor-force weights documented in Section 4.1 (the decomposition in Section 4.4).

4.3 Demographic Okun Estimates

Having established a weak aggregate relationship that is statistically significant in the baseline specification, this section estimates the Okun coefficient separately for each

demographic group. Each series is estimated using the same first-difference specification with country fixed effects and country-clustered standard errors, in both the baseline and dynamic forms. These group-specific estimates provide a transparent first reading of demographic heterogeneity; whether the differences between groups are statistically significant is tested formally in Section 4.4. Results are reported in Table VI.

Table VI: Demographic Okun estimates, by age and gender

	Baseline	Dynamic		
Group	**β**	**β_0**	**β_1**	**Cumulative**
By gender				
Male	-0.146* (0.043)	-0.142* (0.039)	-0.026 (0.038)	-0.168 (0.066)
Female	-0.075 (0.048)	-0.071 (0.047)	-0.018 (0.036)	-0.089 (0.074)
By age				
Youth	-0.305** (0.063)	-0.311** (0.055)	0.017 (0.068)	-0.294* (0.111)
Adult	-0.113* (0.034)	-0.111* (0.031)	-0.019 (0.029)	-0.130* (0.050)
Four-way cells				
Male, adult	-0.118* (0.033)	-0.116* (0.031)	-0.012 (0.030)	-0.129* (0.050)
Male, youth	-0.357** (0.081)	-0.357** (0.073)	-0.010 (0.079)	-0.368* (0.133)
Female, adult	-0.080 (0.047)	-0.075 (0.046)	-0.023 (0.023)	-0.098 (0.063)
Female, youth	-0.120 (0.104)	-0.144 (0.105)	0.121 (0.123)	-0.023 (0.139)
Aggregate *Aggregate (15+)*	*-0.135** *(0.041)*	*-0.130** *(0.037)*	*-0.030* *(0.035)*	*-0.159* *(0.064)*

Notes: Each group estimated separately by first-difference OLS with country fixed effects; SEs clustered by country in parentheses. β is the baseline coefficient; β_0, β_1 the contemporaneous and lagged coefficients; cumulative = $\beta_0 + \beta_1$. Baseline N = 144; dynamic N = 144. * p<0.05, ** p<0.01, *** p<0.001.

The most pronounced difference is along the age dimension. Youth unemployment has a baseline Okun coefficient of −0.305, almost three times the adult coefficient of −0.113, and the youth estimate is significant at the 1% level. A one-percentage-point increase in growth is associated with a fall in youth unemployment of roughly 0.31 percentage points, against only 0.11 for adults. This is consistent with the "last-in, first-out" dynamic discussed in Chapter 2, in which younger workers, concentrated in more precarious and private-sector employment, bear a disproportionate share of cyclical adjustment. The gap is large enough to suggest that the age dimension is the main source of demographic heterogeneity in the sample, a proposition tested directly in the next section.

The gender dimension is weaker and less pronounced. Male unemployment has a baseline coefficient of −0.146, significant at the 5% level, while the female coefficient is −0.075 and not statistically significant. Men therefore appear more cyclically responsive than women, but the female estimate is too imprecise to support a firm conclusion. This pattern is consistent with the more ambiguous cyclical behavior of female unemployment in a regional context of low and variable female labor-force participation, and it is why the gender hypothesis (H_3) was framed without a directional prior. Whether the male–female difference is itself statistically significant is examined in Section 4.4.

The four-way cells, which interact age and gender, reinforce the importance of the age dimension, but the pattern is driven primarily by male youth. The male-youth coefficient is the largest in the table at −0.357 and is significant at the 1% level, whereas the female-youth coefficient, at −0.120, is close to the male-adult estimate of −0.118 and is imprecisely estimated. The cells therefore do not show a uniformly strong youth response across both genders. Neither female coefficient is statistically significant, and the female-youth estimate is unstable across specifications; this imprecision could reflect the volatility of the series itself: the standard deviation of the annual change in female-youth unemployment is 4.28 percentage points (Table III), the largest of the four cells and nearly five times that of adult men. Measurement error in the modelled series may contribute, but the regressions cannot separate the two. The dominant pattern within the cells is thus an age distinction concentrated among men rather than a general youth or gender effect.

Two broader features of the table are central to the argument of this dissertation. First, the dynamic specification adds very little. Across most groups, the lagged growth coefficient is small and statistically insignificant, and the cumulative coefficient ($\beta_0 + \beta_1$) is close to the baseline estimate. As at the aggregate level, the adjustment of unemployment to output is essentially contemporaneous, and the demographic results are robust to the inclusion of lagged growth. Second, and most importantly, the estimate for adult men who constitute roughly 68% of the labor force is −0.118, almost identical to the aggregate coefficient of −0.135. This is a direct illustration of the mechanism set out in Section 4.2: because the aggregate is dominated by adult men, whose own cyclical response is modest, the aggregate coefficient closely tracks theirs. The strongly responsive male-youth group, despite its large coefficient, carries insufficient weight to

pull the aggregate far from the adult-male value. The weak aggregate Okun relationship is therefore not simply a matter of heterogeneity being averaged away; it reflects the modest responsiveness of the numerically dominant group. Section 4.4 tests the statistical significance of these demographic differences formally, and the decomposition that follows quantifies how much each group contributes to the aggregate coefficient.

4.4 Heterogeneity Tests

The demographic estimates in Section 4.3 suggest that cyclical responsiveness differs across groups, but comparing coefficients estimated on separate samples does not establish that the differences are statistically significant. This section tests them formally. For each dimension, the two relevant groups are stacked into a single dataset, and a common model is estimated in which output growth is interacted with a group indicator. The coefficient difference is estimated using Equation (3).

The coefficient γ measures the difference in the Okun coefficient between the two groups directly, and a Wald test of H_0: $\gamma = 0$ asks whether that difference is statistically distinguishable from zero. Country-by-group fixed effects are included, so that each group within each country carries its own intercept. Results are reported in Table VII.

Table VII: Tests of demographic heterogeneity in Okun's coefficient

	Difference in Okun coefficient (γ)	**p-value**
Youth vs. adult	-0.191**	0.003
	(0.035)	
Female vs. male	0.071	0.09
	(0.034)	

Notes: Difference (γ) in the Okun coefficient between each group and its reference (adult; male), from stacked first-difference models with country-by-group fixed effects and country-clustered standard errors (in parentheses). A negative γ indicates the group is more cyclically responsive than its reference. p-values are from two-tailed Wald tests of H_0: $\gamma = 0$ (no difference). N = 288 per test. * p<0.05, ** p<0.01, *** p<0.001.

The age difference is large and highly significant. Youth are more cyclically responsive than adults by $\gamma = -0.191$, with a p-value of 0.003: the null hypothesis of no difference is rejected. This provides strong full-sample evidence of an age difference and supports H_2 under the baseline specification. The gender difference is smaller and much

weaker. Women are estimated to be less cyclically responsive than men ($\gamma = +0.071$), but this difference is not statistically significant ($p = 0.090$), and it is considerably less robust than the age result. There is thus at most weak evidence for a gender difference, and H_3 receives little support.

Two points clarify the interpretation. First, the age and gender findings are not symmetric. In the separate comparisons, the age difference is roughly three times the size of the gender difference and is significant at a far more demanding level. The joint model, in which both dimensions are entered simultaneously over the four age-by-gender cells, confirms the ordering but not the magnitude. There the two adjusted interactions are of similar size (youth $\gamma = -0.140$; gender $\gamma = +0.137$), however it is their precision that separates them: only the age interaction is distinguishable from zero at conventional levels ($p = 0.019$ against $p = 0.115$). Age is therefore the dimension on which the evidence for heterogeneity is statistically stronger, rather than the dimension with the larger adjusted point estimate. Second, the weak gender difference found here should be read alongside the result of Section 4.3 that the female coefficient is not itself statistically distinguishable from zero. The two are consistent: women's unemployment is neither clearly cyclical in absolute terms nor robustly distinguishable from men's. The age difference, by contrast, is large and precisely estimated. Taken together, the full-sample tests provide materially stronger evidence of heterogeneity by age than by gender. The stability of the age result across periods is assessed separately in Section 4.5.

Table VIII: Decomposition of the aggregate Okun coefficient by demographic cell

Demographic cell	LF share (%)	Okun β	Contribution	% of pooled-share total
Male, adult	68.3	−0.118	−0.081	56.0
Male, youth	13.2	−0.357	−0.047	32.7
Female, adult	14.9	−0.080	−0.012	8.3
Female, youth	3.6	−0.120	−0.004	3.0
Pooled-share implied aggregate	**100.0**	**—**	**−0.144**	**100.0**

Notes: Contributions equal pooled mean labor-force shares multiplied by baseline Okun coefficients (Table VI); percentages refer to the pooled-share total. Shares: ILOEST, pooled across the six countries, 2000–2024. Direct aggregate estimate: −0.135 (Table V).

The heterogeneity documented above raises a final question: if youth are so much more cyclically responsive than adults, why is the aggregate Okun coefficient nonetheless weak? The explanation examined here is the composition of the labor force. Because aggregate unemployment is the labor-force-weighted sum of the group-specific unemployment rates, each cell's contribution depends not only on its responsiveness but

also on its labor-force share. Table VIII summarizes this weighting using pooled mean labor-force shares; its implied total of −0.144 differs slightly from the direct estimate of −0.135 because one pooled share vector is applied across countries with different demographic compositions. Using actual country-year shares reproduces aggregate unemployment within 0.001 percentage points, confirming the accounting identity. In the corresponding first-difference decomposition, the share-shift component is −0.0004 (SE 0.0042), showing that changes in labor-force shares contribute negligibly to the estimated Okun coefficient.

The decomposition illustrates the proposed compositional explanation. Adult men, at around 68% of the labor force, account for 56% of the pooled-share total on their own, whilst having the second-shallowest cyclical response of the four cells. Their large weight pulls the aggregate toward their own modest coefficient of −0.118, which is why the directly estimated aggregate (−0.135) lies so close to it. Male youth, by contrast, are strongly cyclical but demographically outweighed: male youth have by far the steepest coefficient (−0.357), but their 13% labor-force share prevents them from dominating the aggregate. Because the accounting identity holds at the country-year level and the share-shift component is negligible, the aggregate Okun response is generated almost entirely by weighted changes in group-specific unemployment rates. The decomposition therefore shows that the comparatively modest response of the numerically dominant adult-male group has the greatest influence on the aggregate, while the stronger male-youth response remains important but receives a substantially smaller weight.

4.5 Robustness Checks

The inference reported so far relies on cluster-robust standard errors, which are known to be unreliable when the number of clusters is small. With only six countries, conventional clustered standard errors tend to overstate significance, so the p-values in the preceding sections should be regarded as suggestive rather than definitive. This section subjects the two central results, the aggregate Okun coefficient and the youth–adult interaction, to four checks: alternative fixed-effects specifications, exclusion of the pandemic period, a wild-cluster bootstrap, and a leave-one-country-out exercise. Results are collected in Table IX.

Table IX: Robustness of the aggregate and age results across specifications

Specification	Aggregate β	Aggregate p	Age γ	Age p
Baseline	−0.135*	0.022	−0.191**	0.003
Add year fixed effects	−0.064	0.399	−0.191**	0.003
Exclude 2020	−0.087	0.129	−0.114	0.100
Exclude 2020–2021	−0.085	0.155	−0.082	0.199
Pre-pandemic (through 2019)	−0.103	0.139	−0.078	0.261
Wild-cluster bootstrap (p-value)	−0.135	0.013	−0.191	0.010
Leave-one-country-out (range)	−0.11 to −0.16	—	−0.17 to −0.21	—

Notes: The aggregate column reports the Okun coefficient β from the total-unemployment model; the age column reports the youth–adult interaction γ from the stacked model. Conventional p-values use a t(5) reference distribution. Wild-cluster-bootstrap p-values are calculated separately from the bootstrap distribution. Leave-one-country-out p-values, based on five clusters, use a t(4) reference distribution. The leave-one-out range reports the smallest and largest coefficient when each country is dropped in turn. * p<0.05, ** p<0.01, *** p<0.001.

The two results behave very differently, and the contrast is itself informative. The aggregate coefficient is fragile. Although negative and significant in the baseline specification, it loses significance once year fixed effects are added ($\beta = -0.064$, $p = 0.40$), and it is not significant when the pandemic year is excluded ($\beta = -0.087$, $p = 0.13$) or in the pre-2020 sample ($\beta = -0.103$, $p = 0.14$). This is not surprising: with only six countries, common shocks like the pandemic account for much of the shared variation in output growth, so year effects absorb most of the variation that identifies the aggregate coefficient. The aggregate association is therefore present in the baseline pooled data but rests heavily on common time variation, and on the pandemic observation in particular.

The age interaction behaves differently. It is essentially unchanged by the inclusion of year fixed effects ($\gamma = -0.191$, $p = 0.003$). This contrast with the aggregate is informative, but its force should not be overstated. Year effects vary only over time, whereas the youth–adult differential is identified from the difference in responsiveness between the two groups within each country-year. A shock common to both groups therefore shifts them equally and cannot, by construction, move γ. The check confirms that the age result does not rest on the common time variation that carries the aggregate;

it does not, and cannot, rule out shocks whose effects differ by age. That distinction matters here, because the pandemic is precisely such a shock. The wild-cluster bootstrap provides a separate check. Following Webb (2023), the p-value is re-estimated using a six-point weight distribution with 9,999 replications, imposing the null and resampling residuals by country. Webb shows that the conventional two-point weights admit an ambiguous set of valid p-values when clusters are few, whereas the six-point distribution yields finer resolution and more reliable finite-sample inference. Under this alternative few-cluster procedure the age difference remains significant (bootstrap $p = 0.010$), and it stays between -0.17 and -0.21 across all leave-one-country-out samples, significant in every case under t(4) inference, so it is not an artifact of any single country. Its sensitivity is to the pandemic period, and that sensitivity is substantial: excluding 2020 roughly halves the coefficient to -0.114 ($p = 0.10$), and in the pre-2020 sample it is smaller and not significant ($\gamma = -0.078$, $p = 0.26$). The age heterogeneity is therefore robust to country deletion and to few-cluster inference, but its magnitude and significance are concentrated in the later, crisis-affected part of the sample, and its stability over normal periods cannot be established from these data.

The gender difference, examined in the preceding section, is weaker still: it does not reach significance under the bootstrap ($p = 0.058$) and is not pursued further here. Taken together, the checks support the following reading: The aggregate Okun relationship in these economies is weak and fragile, sensitive to common shocks much like the pandemic. The more durable finding is the demographic contrast: youth unemployment is substantially more cyclically responsive than adult unemployment, a difference that survives the year-effect and few-cluster checks, though it too weakens materially when the pandemic period is removed. This pattern is consistent with the central argument of the dissertation; that the aggregate coefficient is uninformative about the cyclical experience of the groups within it, while being honest about the limits that six countries and a single large shock impose on statistical inference.

5. CONCLUSION

This dissertation asked whether Okun's Law holds in a sample of six MENA economies over 2000–2024, and whether the responsiveness of unemployment to output growth differs across demographic groups defined by age and gender. The motivation

was the well-documented "jobless growth" of the region, and the suspicion that a weak or insignificant aggregate Okun relationship might conceal meaningful variation across the groups that make up the labor force.

The findings can be stated simply. At the aggregate level, Okun's Law is weak and fragile: a one-percentage-point increase in output growth is associated with a fall in unemployment of about 0.14 percentage points in the baseline specification. The estimate is negative but far shallower than the responses typically estimated for advanced economies, and it is not robust to common year effects or the exclusion of the pandemic period. This places six selected countries between the two positions in existing literature, neither exhibiting a textbook Okun relationship nor the complete disconnect sometimes reported for MENA labor markets.

The more important results are demographic. Youth unemployment is estimated to be far more cyclically responsive than adult unemployment, with a youth–adult coefficient difference of approximately −0.19. This difference is robust to wild-cluster bootstrap inference designed for few clusters, to the exclusion of any individual country, and to the inclusion of common year effects that dissolve the aggregate relationship. It is, however, sensitive to the pandemic period: without 2020, the differential falls by roughly half and loses statistical significance, so its magnitude and significance are concentrated in the later, crisis-affected years, and its stability over normal, non-crisis periods remains uncertain. The gender estimates point toward weaker responsiveness among women than men, but the difference is marginal and does not withstand the same scrutiny, so it should be regarded as suggestive rather than established. The evidence for heterogeneity is stronger along the age dimension than along the gender dimension, although the stability of the age differential outside crisis periods remains uncertain.

The decomposition ties these results together and constitutes the dissertation's central contribution. The apparent contrast is that youth unemployment is more responsive while the aggregate relationship remains weak; the decomposition shows how unequal labor-force weights reconcile these results. The aggregate coefficient is approximately a labor-force-weighted average of the group-specific coefficients, and adult men, who make up roughly two-thirds of the labor force, account for more than half of it on their own, despite having one of the more modest cyclical responses. Their sheer weight pulls the aggregate

toward their own modest coefficient, while the strongly cyclical male-youth, being a small share of the labor force, contribute comparatively less. The weak aggregate Okun relationship is therefore not evidence that demographic heterogeneity is absent; it is consistent with the labor force being dominated by its numerically largest group, whose response is comparatively modest. The aggregate coefficient understates the estimated responsiveness of youth overall, and male youth in particular, precisely because these groups are numerically small.

These conclusions carry the usual caveats. The sample contains only six countries, so cluster-robust inference is inherently fragile; this was addressed with a wild-cluster bootstrap, but the small number of clusters remains a fundamental limitation, and the aggregate result in particular is fragile, which is sensitive to common year effects, to the pandemic observation, and, at the margin, to the exclusion of individual countries. The unemployment series are ILO modelled estimates rather than direct national measurements, which ensures cross-country comparability but may smooth short-run movements and create correlated measurement error across demographic series. The analysis uses unemployment responsiveness as an indicator of jobless growth but does not directly estimate employment creation, since employment can rise without unemployment falling when labor-force participation also changes. The analysis is also reduced-form: the Okun coefficient summarizes a bundle of underlying margins (participation, hours, and productivity) without separating them, and it should not be read as a structural or causal parameter. Finally, the demographic analysis treats age and gender as separate, additive influences on cyclical responsiveness. This specification can measure the youth effect and the gender effect individually, but it cannot test whether the two interact. For instance, whether the heightened responsiveness of youth is concentrated among young men. Such an interaction can be estimated, but with only six country clusters and a female-youth series whose year-to-year variation is the largest of the four cells, it is too imprecisely estimated to support reliable inference.

Future research could examine a broader country panel, estimate demographic-specific responses to common shocks, and use employment and participation data to separate the channels underlying unemployment adjustment. Within the present sample, the most defensible conclusion is that the aggregate Okun relationship is weak and fragile, while full-sample unemployment responsiveness differs more clearly by age than by

gender. The extent to which this age difference represents a stable normal-period relationship, rather than heightened responsiveness during major crises, remains an open question. What this dissertation establishes is narrower but, it is hoped, secure: in these economies, the weak aggregate relationship coexists with stronger estimated responsiveness among some demographic groups, whose comparatively small labor-force weights limit their influence on the aggregate estimate.

REFERENCES

Assaad, R. and Barsoum, G. (2019). Public employment in the Middle East and North Africa. *IZA World of Labor* 463.

Ball, L., Leigh, D. and Loungani, P. (2017). Okun's Law: Fit at 50? *Journal of Money, Credit and Banking*, *49*(7), 1413–1441.

Bell, D. N. F. and Blanchflower, D. G. (2011). Youth Unemployment in Europe and the United States. *IZA Discussion Paper* No. 5673. Bonn: Institute for the Study of Labor (IZA).

Bergé, L.R., Butts, K. and McDermott, G. (2026). Fast and user-friendly econometrics estimations: The R package fixest. arXiv preprint arXiv:2601.21749.

Crespo Cuaresma, J. (2003). Okun's Law Revisited. *Oxford Bulletin of Economics and Statistics* 65 (4), 439-451.

International Labour Organization. (2025). *ILO modelled estimates database, ILOSTAT* [database]. Available from https://ilostat.ilo.org/methods/concepts-and-definitions/ilo-modelled-estimates/.

Kreishan, F. M. (2011). Economic Growth and Unemployment: An Empirical Analysis. *Journal of Social Sciences* 7 (2), 228-231.

La Porta, R. and Shleifer, A. (2014). Informality and Development. *Journal of Economic Perspectives* 28 (3), 109-126.

Lee, J. (2000). The Robustness of Okun's Law: Evidence from OECD Countries. *Journal of Macroeconomics* 22 (2), 331-356.

Moosa, I. A. (2008). Economic Growth and Unemployment in Arab Countries: Is Okun's Law Valid? *Journal of Development and Economic Policies* 10 (2), 5-24.

OECD (2016). *Youth in the MENA Region: How to Bring Them In*. Paris: OECD Publishing.

O'Higgins, N. (1997). The challenge of youth unemployment. *International Social Security Review* 50 (4), 63-93.

Okun, A. M. (1962). Potential GNP: Its Measurement and Significance. In: *Proceedings of the Business and Economic Statistics Section, American Statistical Association*. Washington, DC: American Statistical Association.

Prachowny, M. F. J. (1993). Okun's Law: Theoretical Foundations and Revised Estimates. *The Review of Economics and Statistics* 75 (2), 331-336.

Sachs, J. D. and Warner, A. M. (2001). The curse of natural resources. *European Economic Review* 45 , 827-838.

Sögner, L. and Stiassny, A. (2002). An analysis on the structural stability of Okun's law—a cross-country study. *Applied Economics* 34, 1775-1787.

Webb, M.D. (2023), Reworking wild bootstrap-based inference for clustered errors. *Canadian Journal of Economics/Revue canadienne d'économique*, 56: 839-858

World Bank. (2026). *GDP growth (annual %) [Data set]*. World Development Indicators. https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG

World Bank (2013). *Jobs for Shared Prosperity: Time for Action in the Middle East and North Africa*. Washington, DC: World Bank.

Zanin, L. (2014). On Okun's law in OECD countries: An analysis by age cohorts. *Economics Letters* 125 (2), 243-248

APPENDIX

Table X: Leave-one-country-out estimates of the age interaction

Excluded country	Age interaction (γ)	Clustered SE	t(4) p-value
Algeria	−0.174	0.030	0.004
Egypt	−0.170	0.025	0.002
Jordan	−0.198	0.042	0.009
Morocco	−0.204	0.043	0.009
Saudi Arabia	−0.197	0.050	0.018
United Arab Emirates	−0.213	0.051	0.014

Notes: Each row re-estimates the stacked age model (Equation 3) excluding one country. With five clusters remaining, p-values use a t(4) reference distribution. The full-sample estimate is γ = −0.191 (SE 0.035, p = 0.003).

Table XI: Country-specific Okun slopes

Country	Okun slope (β)	Std. error	N
Algeria	−0.376	(0.154)	24
Egypt	−0.446	(0.104)	24
Jordan	−0.133	(0.096)	24
Morocco	−0.160	(0.036)	24
Saudi Arabia	−0.061	(0.035)	24
United Arab Emirates	−0.079	(0.022)	24

Notes: Each row is a separate first-difference OLS regression of the change in total unemployment on GDP growth for a single country (24 annual observations each). Standard errors are conventional OLS standard errors.