The Workplace Structure of Kuwait in 2011

Employment centres, establishment concentration, and spatial separation from residents

Abdulmalik Altararwa

August 2026

Completed descriptive research paper

Working paper Version: 13 August 2026 **Evidence class:** Descriptive spatial measurement

Abstract

Where firms locate and where people live are separate parts of an urban economy. This paper measures that separation in Kuwait using the public 2011 Census Geoportal of the Central Statistical Bureau. The source contains 773 lower-area polygons, 82,930 active establishments, and 715,968 workers attached to establishments. A transparent rule that requires at least 5,000 establishment workers per square kilometre in each polygon and at least 10,000 workers in each contiguous component identifies ten employment centres. Their 65 polygons contain 430,859 workers, or 60.2 percent of the source total. This classification is sensitive but not fragile: across nine pre-specified threshold combinations, centre shares range from 37.9 to 71.9 percent. Workers are more geographically concentrated than establishments. The top 50 polygons contain 60.2 percent of workers and 54.3 percent of establishments; inverse-Herfindahl effective polygon counts are 74.5 and 118.7. Finally, the spatial dissimilarity between establishment workers and residents classified by main occupation is 0.691 across lower areas, 0.570 across districts, and 0.172 across governorates. The decline shows how coarse aggregation cancels fine local differences. These results are historical classifications and distributional comparisons. They do not identify commuting flows, a current employment geography, or a causal effect of concentration.

Keywords: employment centres; workplace geography; spatial concentration; workplace-residence dissimilarity; Kuwait; census geography **JEL codes:** R12; R14; R23; R41

1. Introduction

A national employment total hides the spatial organization of work. The same number of workers can be spread across many mixed-use districts, concentrated in a central business district, or divided among commercial and industrial nodes. Those arrangements place different demands on roads, public transport, land, and infrastructure. They also create different environments for sharing suppliers, matching firms with workers, and learning from nearby activity. Before asking whether concentration raises productivity or congestion, however, we need to know what is concentrated, where it is concentrated, and how much the answer depends on the unit used to draw the map.

This paper measures Kuwait's workplace structure in the public 2011 Census Geoportal. Its object is the geography of active establishments and workers attached to establishments across 773 official lower-area polygons. The exercise has four parts. I first describe workplace intensity and worker composition. I then compare the spatial concentration of workers with the concentration of establishments. Third, I classify contiguous employment centres under a published density-and-total rule and repeat the calculation over a complete threshold grid. Finally, I compare the normalized spatial distribution of establishment workers with the distribution of residents, both at the original lower-area scale and after aggregation to districts and governorates.

This is not Kuwait's first employment map. Said (1986) mapped population, employment, and work trips using the 1980 census tract system. His analysis identified Kuwait City, Shuwaikh, Hawalli, Salmiya, the airport area,

Ahmadi, and Fahaheel as important employment locations. Adeel and Alfahad (2021) mapped jobs and residences by nationality in 2011 and studied transport accessibility. Da Cruz et al. (2024) used block-level 2011 employment density when examining residential segregation and urban governance. Those papers rule out a broad novelty claim. The contribution here is narrower and establishment-centric: an internally reproducible reconstruction of the complete local lower-area layer, a transparent centre classification with membership and sensitivity results retained in the research archive, a direct worker-versus-establishment concentration comparison, and a workplace-residence share-distance measure reported at three geographic scales.

The distinction between workers and establishments matters. A place may contain many workers because it contains many establishments, because its establishments are large on average, or both. The public polygon data cannot recover a firm-size distribution, but they can show whether the two spatial distributions coincide. They do not. The top ten polygons account for 24.5 percent of establishment workers and 19.4 percent of active establishments. The spatial Herfindahl index is 0.0134 for workers and 0.0084 for establishments. Expressed as inverse-Herfindahl effective polygon counts, the worker distribution resembles an equal distribution across 74.5 polygons, while the establishment distribution resembles one across 118.7 polygons.

The centre results make that concentration legible as a spatial system. Under the primary rule - 5,000 workers per square kilometre in each qualifying polygon, queen contiguity, and 10,000 workers in the resulting component - the algorithm returns ten centres. The three largest are Sharq-Mirqab-Qibla, Hawalli-Jabriya, and Shuwaikh Industrial-Rai. Together they contain 281,511 establishment workers. All ten centres contain 430,859 workers across 65 polygons, or 60.2 percent of the national source total. These clusters are measurement objects, not official planning districts or estimated labor markets. A nine-cell sensitivity grid makes their dependence on thresholds visible: the retained worker share ranges from 37.9 percent under the strictest combination to 71.9 percent under the most inclusive.

The residence-side comparison adds a different result. Across lower areas, the dissimilarity index between establishment workers and residents classified by main occupation is 0.691. It falls to 0.570 across 140 governorate-district combinations and 0.172 across six governorates. This decline is not a commuting estimate. The source contains no person-level origin-destination link, and the workplace and residence totals cover different source universes. It instead shows that large workplace-share surpluses and residence-share surpluses often occur within the same broad governorate. A governorate map consequently makes the distributions look much closer than a lower-area map does.

The paper is deliberately descriptive. Agglomeration economies, industrial land allocation, government office location, transport access, sector composition, and historical planning are all plausible explanations for the observed geography. A single historical cross-section cannot separate them. The estimands are counts, shares, classifications, and spatial distances - not treatment effects. There is no identifying variation, no causal estimator, and no sampling confidence interval.

The rest of the paper proceeds as follows. Section 2 places the measurement exercise in the Kuwait and employment-centre literatures. Section 3 describes the data, source universes, and validation. Section 4 states the descriptive estimands and algorithms. Section 5 presents the results. Section 6 discusses interpretation, policy use, and limitations. Section 7 concludes.

2. Literature and setting

2.1 Why workplace concentration is economically interesting

Urban concentration can lower the cost of sharing inputs and infrastructure, improve matching between firms and workers, and support learning across nearby organizations (Duranton and Puga 2003). These mechanisms make an employment centre more than a cartographic curiosity. The location and scale of centres can shape the number and direction of trips, the value of access, and the exposure of an economy to congestion or local disruption.

The same pattern can emerge without a productivity effect. Planning rules may place offices or industrial activity in designated zones. Public employment may be administratively concentrated. Road and port access may determine where land is usable for particular activities. Large sites may be separated from residential districts for environmental or logistical reasons. Historical investments can persist after the original reason for a location has changed. A workplace map records the equilibrium produced by these forces; it does not identify their individual contributions.

The empirical employment-centre literature offers several ways to describe that map. Giuliano and Small (1991) use a transparent density-plus-total logic: a subcentre must contain high-density units and enough total employment to be substantively important. McMillen (2001) develops a nonparametric alternative based on a smoothed employment-density surface. The second method can reduce reliance on a single threshold, but it introduces smoothing and model choices of its own. For Kuwait's one public cross-section of 773 irregular polygons, a deterministic density-and-total rule is easier to audit and reproduce. The important discipline is to publish the thresholds, refuse to call them natural boundaries, and show how the answer changes when they move.

Spatial scale creates another constraint. Fotheringham and Wong (1991) show that statistical patterns can change when observations are aggregated or rezoned, the modifiable-areal-unit problem. This paper cannot redraw the original census into a set of equally plausible alternative block systems. It therefore does not solve the problem. It runs two focused checks that are possible with the public data: a full grid of centre thresholds and a workplace-residence index recalculated at lower-area, district, and governorate scales.

2.2 What the Kuwait literature already establishes

The historical benchmark is Said's (1986) study of urban development and transport. Using the 1980 census, Said maps spatial distributions of population and employment and studies work-trip production, attraction, and interzonal linkages by nationality. Kuwait City is a major employment location in that account, while Shuwaikh, Hawalli, Salmiya, the airport, Ahmadi, and Fahaheel are also important. The study means that neither employment concentration nor a multi-node workplace structure is a new observation for Kuwait. It also has information this paper does not: observed journey-to-work linkages and travel-time-based spatial interaction models.

Adeel and Alfahad (2021) provide the closest recent map comparison. They combine 2011 spatial data with transport-accessibility measures and qualitative evidence, presenting the differentiation of jobs and residences by nationality. Their focus is equitable transport access, including access to bus stops and social constraints on mobility. This paper does not replicate that analysis. It uses the public CSB establishment layer at lower-area resolution, treats the workplace geography as the main measurement object, and documents a centre algorithm and concentration tables retained in the local research archive. Its dissimilarity index is a distance between distributions, not an accessibility measure.

Da Cruz et al. (2024) study spatial patterns, residential segregation, and urban governance in Kuwait. Their analysis explicitly uses 2011 block employment density as a correlate of residential patterns. That prior use rules out any claim that this is the first analysis of the Geoportal's employment density. The overlap also sharpens the contribution here: workers, active establishments, establishment-worker nationality, and residence-side occupation fields are considered jointly, while the output remains a descriptive workplace paper rather than a segregation model.

The incremental contribution therefore rests on transparency and joint measurement rather than discovery of a previously unseen node. The paper asks which lower areas qualify under a declared employment-centre rule; how workers and establishments differ in concentration; how robust the centre share is to thresholds; and how workplace shares differ from residence shares at multiple scales. This framing is deliberately modest, but it produces a reusable benchmark for later transport, land-use, or longitudinal work.

2.3 Kuwait's spatial structure and the interpretation of a 2011 baseline

Kuwait's urban and economic geography includes a historic commercial and administrative core, dense mixed-use districts, industrial and storage areas, ports, oil-related southern activity, and residential settlements extending away from the centre. The lower-area polygons do not map one-to-one onto economic functions, but their names make several recognizable clusters visible. The 2011 geography is still a historical snapshot. Later roads, towns, relocations, and boundary revisions lie outside the source.

This vintage discipline matters for comparison. The project retains a separately audited crosswalk to contemporary PACI geography, but 232 of 773 links are held out because the evidence is insufficient for a defensible match. None of the main analysis uses that crosswalk. Every result uses original 2011 polygons or aggregations built directly from their 2011 governorate and district identifiers. The paper can therefore say where the source counted work in 2011, not where Kuwait works today.

3. Data

3.1 Source and observation unit

The primary source is the Kuwait Central Statistical Bureau (CSB) 2011 Census Geoportal. The archived GeoJSON contains 773 lower-area polygons and 507 non-geometry attributes covering population, buildings, households, and establishments. Each record includes English and Arabic names, governorate and district identifiers, a published area, geometry, and the census fields used in this paper. The raw file is stored without modification; its SHA-256 hash and source-level provenance appear in the project manifest.

The main observation unit is the lower-area polygon. A composite CSB identifier is unique across all 773 records. For scale comparisons, I aggregate these same polygons into 140 unique governorate-district combinations and six governorates. I do not assign polygons to contemporary geography.

Table 1 separates the main source populations. There are 82,930 active establishments and 715,968 workers attached to establishments. Of those workers, 92,534 are recorded as Kuwaiti and 623,434 as non-Kuwaiti. The same file records 2,822,431 residents and 1,604,676 residents classified by main occupation.

Table 1. National source totals, 2011

Source measure	Count
Lower-area polygons	773
Active establishments	82,930
Workers attached to establishments	715,968
Kuwaiti establishment workers	92,534
Non-Kuwaiti establishment workers	623,434
Residents	2,822,431
Residents classified by main occupation	1,604,676

Note: Counts belong to different census fields and need not describe the same person universe. Source: Kuwait CSB 2011 Census Geoportal.

3.2 Workplace and residence-side universes

The file combines several universes in each polygon, and the distinction is central to the design. a_est_status_active measures active establishments. The Kuwaiti and non-Kuwaiti establishment-labor fields sum to the workplace worker count. Ten main-occupation population fields sum to the number of residents classified by occupation. Eighteen main-economic-activity population fields independently sum to exactly the same total in every polygon. Finally, resident nationality fields sum to total population.

These identities validate internal arithmetic; they do not turn the universes into person-level matches. The 715,968 establishment workers are not the same denominator as the 1,604,676 occupation-classified residents. Some resident categories need not appear in the public establishment total, and the source documentation available here does not support a complete reconciliation. For that reason, the main workplace-residence statistic first normalizes each side to sum to one. It compares where each source population is distributed, not how many workers commute between the two.

The nationality comparison has a further limitation. The workplace side can be split into Kuwaiti and non-Kuwaiti workers, but the residence side cannot cross-tabulate resident workers by nationality. The available denominators are all Kuwaiti and all non-Kuwaiti residents. Those secondary indices may describe spatial separation between workplace counts and national resident populations, but differences between the two nationality indices can also reflect age structure, labor-force participation, and source coverage. They are not evidence that one nationality travels farther.

3.3 Constructed variables

Active-establishment density and worker density divide the relevant polygon counts by the CSB's published area in square kilometres. Workers per active establishment divide total establishment workers by active establishments. This last measure is a polygon aggregate, not an average constructed from establishment microdata; a high value cannot reveal whether a polygon has one very large establishment or several medium-sized ones.

The non-Kuwaiti worker share is

$$100 \times (W\_i^N) / (W\_i^K + W\_i^N),$$

where *W[i]^(K)* and *W[i]^(N)* are Kuwaiti and non-Kuwaiti establishment workers in polygon *i*. For public maps, the share is suppressed when the denominator is below 20. This is a presentation-stability rule, not an assertion about the CSB's disclosure regime.

The build also reconstructs five percentage fields from their published count fields: Kuwaiti worker share, non-Kuwaiti worker share, technology use for all activities, technology use for some activities, and no technology use. The maximum absolute discrepancy from the Geoportal's percentage values is 0.005 percentage points, consistent with rounding to two decimal places.

3.4 Geometry and validation

All 773 polygons have unique identifiers, valid non-empty geometry, positive published area, and non-negative required counts. National active-establishment, worker, resident, and resident-occupation totals reproduce the archived source. Kuwaiti and non-Kuwaiti establishment workers sum exactly to total establishment workers. Occupation and economic-activity population categories reconcile in every polygon.

As a geometry check, the national area derived after projection to UTM zone 39N differs from the sum of the published area fields by 0.189 percent. The analysis uses the published areas for density so that the denominator remains traceable to the Geoportal. All validation assertions run before any final table or figure is exported. The machine-readable validation file contains 14 passed checks.

4. Empirical measurement

4.1 Descriptive estimands

The paper estimates no causal parameter. Its four families of descriptive estimands are:

1. polygon workplace intensity: counts and densities of active establishments and establishment workers;
2. spatial concentration: top-polygon shares, a spatial Herfindahl index, and the inverse-Herfindahl effective polygon count for workers and establishments;
3. threshold-defined employment centres: centre membership, counts, and worker shares under a primary rule and eight alternatives; and
4. workplace-residence dissimilarity: the distance between normalized workplace and residence distributions at three spatial scales.

Because the source is a census layer rather than a sample, sampling error is not the relevant uncertainty. Classification uncertainty, coverage differences, source error, and spatial aggregation are. The threshold grid and scale comparison address two of those margins directly.

4.2 Worker and establishment concentration

For a measure *x[i]* - workers or establishments - I order polygons from largest to smallest and calculate the cumulative national share in the top 1, 10, 50, and 100. I also calculate

$$\mathrm{HHI}_x = \Sigma_{i=1}^{773}((x\_i) / (\Sigma_j x\_j))^2.$$

This is a spatial HHI: the units are places, not firms. It is not an index of product-market concentration or market power. Its inverse, *1/HHI[x]*, is the number of equally sized polygons that would generate the observed index. Calling it an effective polygon count makes the comparison intuitive while retaining its geographic meaning.

4.3 Employment-centre classification

The centre algorithm adapts the density-plus-total logic of Giuliano and Small (1991). Let *d[i]=W[i]/A[i]* denote establishment-worker density in polygon *i*, using published area *A[i]*. For a chosen density threshold *mean d*, the algorithm retains every polygon satisfying *d[i]* ≥ *mean d*. It then constructs a queen-contiguity graph: two retained polygons are neighbors if their geometries touch at an edge or a point. Each connected component *c* is retained as an employment centre when

$$\Sigma_{i \square c} W_i \geq \text{mean } W,$$

where *mean W* is a component worker floor.

The primary rule sets *mean d=5,000* establishment workers per square kilometre and *mean W=10,000* workers. These values are a transparent calibration for this source, not an estimate of a natural discontinuity. The full sensitivity grid crosses density floors of 2,500, 5,000, and 7,500 workers per square kilometre with component floors of 5,000, 10,000, and 15,000 workers. For each of the nine cells, I report the number of centres, number of member polygons, combined worker count, and national worker share.

Centre labels summarize the largest named areas inside each component. Membership, not the short label, defines the result. A centre may cross familiar district names, and two components with similar names remain separate when they are not connected through qualifying polygons. None is an official centre, legal zone, functional labor market, or observed commuting basin.

4.4 Workplace-residence dissimilarity

For workplace counts *W[i]* and residence-side counts *R[i]*, the dissimilarity index is

$$D(W,R)=(1)/(2)\Sigma_i|(W\_i)/(\Sigma_j W\_j)-(R\_i)/(\Sigma_j R\_j)|.$$

The index ranges from zero, when the two normalized spatial distributions are identical, to one, when their mass is in disjoint units. It can also be read as the total-variation distance between the distributions. That interpretation refers to shares. It is not the number or share of people who need to move, because the two totals and source universes differ.

The primary comparison sets *W[i]* equal to establishment workers and *R[i]* equal to residents classified by main occupation. Secondary comparisons use Kuwaiti establishment workers against all Kuwaiti residents and non-Kuwaiti establishment workers against all non-Kuwaiti residents. I calculate each index across 773 lower areas, 140 districts, and six governorates.

Aggregation can combine positive and negative local share differences inside the same broader unit. The triangle inequality therefore implies that dissimilarity cannot increase when a fixed fine partition is aggregated. The size of the decline is substantively informative about how much fine local separation the coarser map conceals, but it should not be described as a statistical robustness interval or as a solution to MAUP.

5. Results

5.1 Workplace geography and the leading polygons

Figure 1 shows active-establishment density beside the non-Kuwaiti share of establishment workers. Dense establishment geography appears around the metropolitan core, Hawalli and nearby areas, and several commercial or industrial nodes. The nationality map is broadly dark because non-Kuwaitis account for 623,434 of the 715,968 establishment workers, or 87.1 percent. There is still meaningful local variation, including a larger Kuwaiti share in some central polygons.

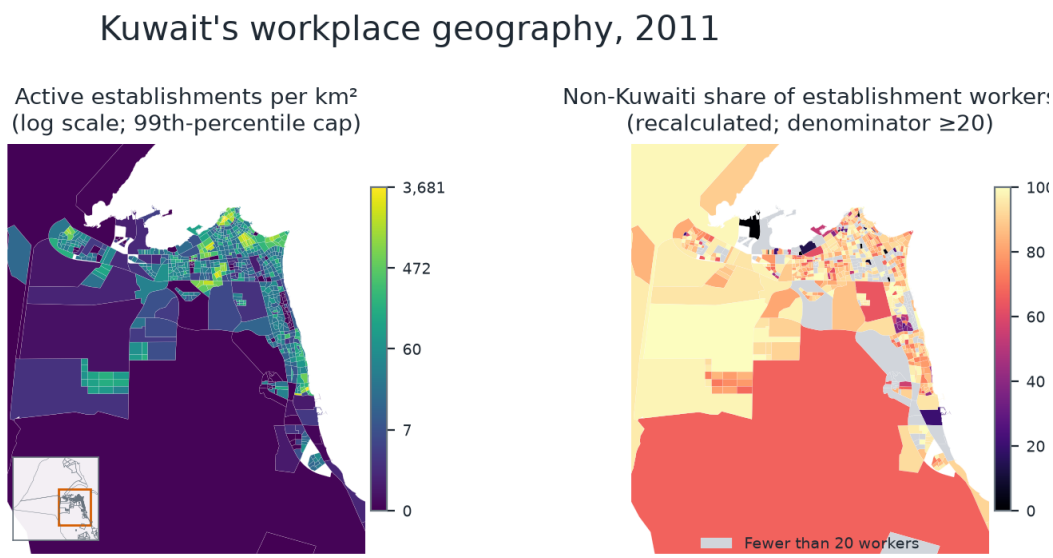

Source: Kuwait Central Statistical Bureau, 2011 Census Geoportal. Unit: 773 lower-area polygons. Establishment density divides active establishments by published polygon area. The worker share divides non-Kuwaiti establishment workers by Kuwaiti plus non-Kuwaiti establishment workers; polygons with fewer than 20 workers are grey. Main panels crop to the metropolitan/industrial core; the inset locates that crop within the full country. Pure 2011 geography is used; no contemporary boundary match is imposed.

Figure 1. Kuwait's workplace geography, 2011

The map is deliberately cropped to make the metropolitan and industrial pattern readable, with a full-country locator preserving the national extent. It uses counts attached to establishments. A dark polygon on the right does not say where non-Kuwaiti residents live, how they travel, or which industries employ them.

Figure 2 ranks the 15 polygons with the most establishment workers. Hawalli 4 is the largest at 52,319, or 7.3 percent of the national total. Jleeb Al-Shiyoukh 1 follows at 25,236. The next group includes Shuwaikh Industrial-2 block 1 at 15,810, Fahaheel 11 at 13,527, Shuwaikh Industrial-1 block 2 at 12,563, Ardhiya Makhzen 2 at 12,203, Sharq 5 at 11,637, and Subhan 11 at 11,426. The list mixes dense commercial areas with industrial and storage locations. Count and density should therefore be kept separate: a large worker count can occupy a compact mixed-use polygon or a more extensive industrial site.

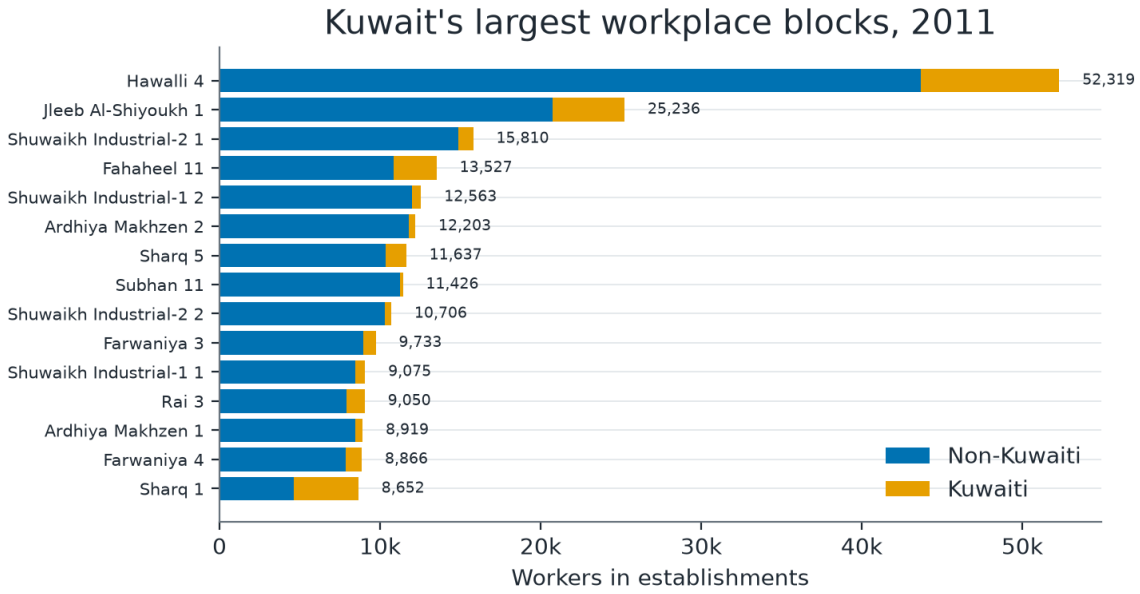

Source: Kuwait Central Statistical Bureau, 2011 Census Geoportal. Bars are workers attached to establishments in each lower-area polygon, split by nationality. They are workplace counts, not residents or commuter flows. The 15 polygons account for 31.7% of the national establishment-worker total.

Figure 2. Kuwait's largest workplace polygons, 2011

One local comparison foreshadows the workplace-residence result. Hawalli 4 alone contains 7.31 percent of establishment workers but only 0.23 percent of occupation-classified residents, a workplace-share surplus of 7.08 percentage points. In Jleeb Al-Shiyoukh, block 1 has a workplace-share surplus of 2.71 points, while blocks 2 and 3 have residence-share surpluses of 3.41 and 2.72 points. Large differences can therefore occur among neighboring polygons carrying the same broad place name. District aggregation necessarily smooths them.

5.2 Workers are more concentrated than establishments

Figure 3 compares the complete cumulative distributions. The worker curve rises more quickly than the establishment curve. The highest-ranked polygon contains 7.3 percent of workers but 3.0 percent of establishments. The top ten contain 24.5 percent of workers and 19.4 percent of establishments; the top 50 contain 60.2 and 54.3 percent; and the top 100 contain 80.9 and 74.5 percent.

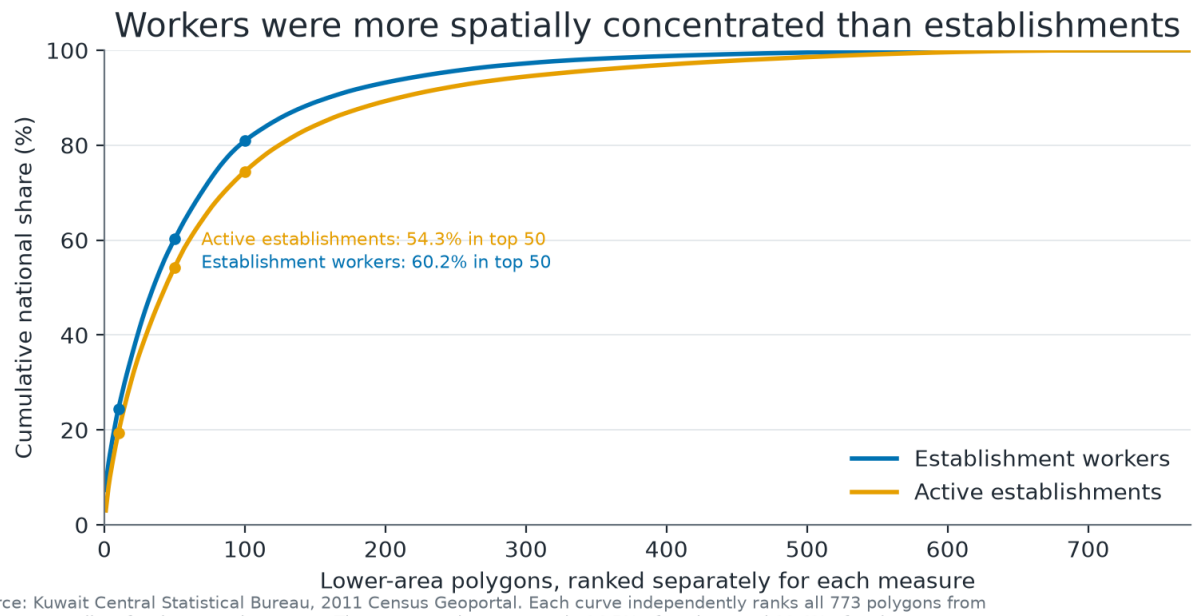

Source: Kuwait Central Statistical Bureau, 2011 Census Geoportal. Each curve independently ranks all 773 polygons from largest to smallest for the named measure. The top ten polygons contain 24.5% of workers and 19.4% of active establishments; the top 50 contain 60.2% and 54.3%. Polygon areas differ, so these are spatial concentration measures, not firm-size or product-market concentration.

Figure 3. Workers were more spatially concentrated than establishments, 2011

Table 2. Spatial concentration of workers and establishments

Measure	Top 1	Top 10	Top 50	Top 100	Spatial HHI	Effective polygons
Establishment workers	7.3%	24.5%	60.2%	80.9%	0.0134	74.5
Active establishments	3.0%	19.4%	54.3%	74.5%	0.0084	118.7

Note: Each row ranks polygons separately by the named measure. Effective polygons equal the inverse spatial HHI. These are place-concentration measures, not product-market HHIs.

The gap suggests that the spatial distribution of establishment size matters. It cannot show the underlying firm-size distribution, because the public record aggregates all establishments in a polygon. Even so, the result is economically useful. An infrastructure plan based only on establishment counts would spread attention more widely than one based on worker exposure. Conversely, the worker curve alone would miss places with many small establishments. Which measure belongs in an application depends on whether the object is workers, business units, trips, or service coverage.

5.3 Ten centres under the primary classification

Figure 4 maps the ten centres identified by the primary rule and lists them by worker count. Sixty-five qualifying polygons contain 430,859 workers, 60.2 percent of the national establishment-worker total, and 38,481 active establishments, 46.4 percent of the establishment total. The difference repeats the national concentration result: centre polygons hold a larger share of workers than establishments.

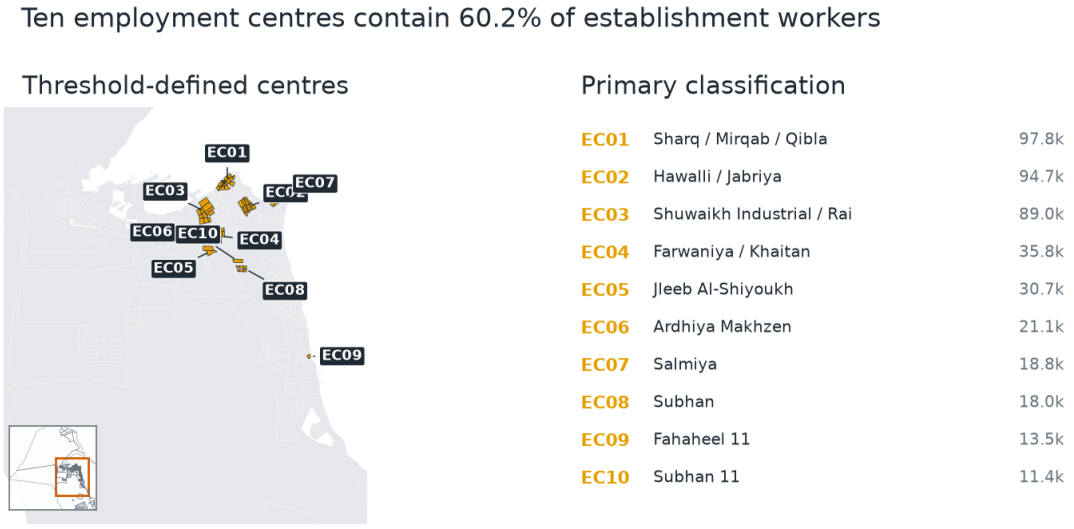

Source: Kuwait Central Statistical Bureau, 2011 Census Geoportal. A primary employment centre is a queen-contiguous set of lower-area polygons in which every polygon has at least 5,000 establishment workers per km² and the component has at least 10,000 workers. This rule identifies 10 centres, 65 polygons, and 430,859 workers. Labels correspond to employment_centres_2011.csv. Boundaries and thresholds are measurement choices, not natural or administrative centre definitions.

Figure 4. Threshold-defined employment centres, 2011

Table 3. Primary employment-centre classification

Rank	Centre label	Polygons	Workers	Worker share	Active establishments	Non-Kuwaiti worker share
1	Sharq / Mirqab / Qibla	21	97,850	13.7%	11,104	84.4%
2	Hawalli / Jabriya	11	94,682	13.2%	7,609	84.0%
3	Shuwaikh Industrial / Rai	11	88,979	12.4%	7,207	93.9%
4	Farwaniya / Khaitan	6	35,778	5.0%	4,819	88.0%
5	Jleeb Al-Shiyoukh	2	30,683	4.3%	3,207	83.3%
6	Ardhiya Makhzen	2	21,122	3.0%	205	95.8%
7	Salmiya	3	18,816	2.6%	1,951	86.6%
8	Subhan	7	17,996	2.5%	226	96.7%
9	Fahaheel 11	1	13,527	1.9%	2,110	80.3%
10	Subhan 11	1	11,426	1.6%	43	98.4%
All centres		**65**	**430,859**	**60.2%**	**38,481**	**87.9%**

Note: Every member polygon has at least 5,000 establishment workers per square kilometre, and each queen-contiguous component has at least 10,000 workers. Centre labels abbreviate membership; the locally archived centre table contains the binding identifiers.

The three largest components each contain roughly 89,000 to 98,000 workers, but their establishment structures differ. Sharq-Mirqab-Qibla contains 11,104 active establishments and 8.8 workers per establishment. Hawalli-Jabriya contains 7,609 and 12.4 workers per establishment. Shuwaikh Industrial-Rai contains 7,207 and 12.3. At the more extreme end, Ardhiya Makhzen has 21,122 workers across 205 establishments, while Subhan 11

has 11,426 across 43. Those ratios identify industrial or storage polygons where workers are attached to relatively few establishments in aggregate. They do not reveal firm identities or within-polygon dispersion.

The names of several classified centres echo Said's (1986) historical employment geography: the central city, Shuwaikh, Hawalli, Salmiya, and Fahaheel appear in both. That overlap is suggestive of persistence, but it is not a longitudinal estimate. The 1980 and 2011 sources use different boundaries, counts, and methods. A valid persistence study would require harmonized spatial units and explicit treatment of boundary change.

5.4 Sensitivity to the centre rule

Figure 5 reports every pre-specified threshold combination. Centre counts range from five to twelve. Their combined share of national establishment workers ranges from 37.9 to 71.9 percent. The strictest rule - 7,500 workers per square kilometre and 15,000 workers per component - keeps five centres, 34 polygons, and 271,704 workers. The most inclusive rule - 2,500 per square kilometre and 5,000 per component - keeps twelve centres, 93 polygons, and 515,081 workers.

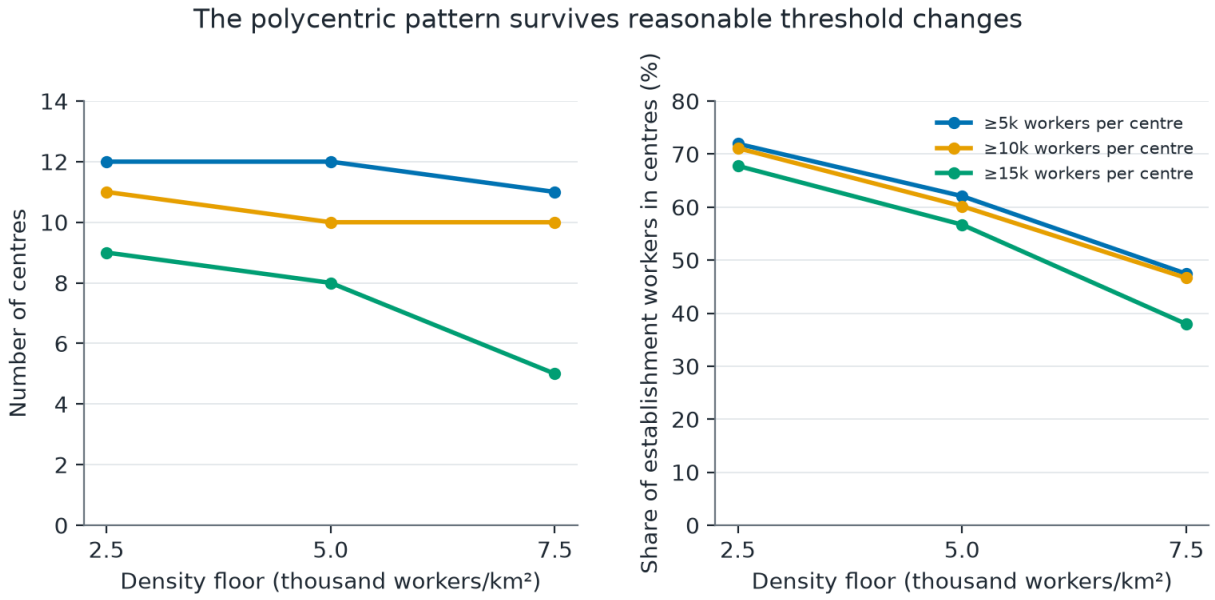

Source: Kuwait Central Statistical Bureau, 2011 Census Geoportal. Each line applies the same queen-contiguity algorithm while changing the per-polygon density floor (2,500, 5,000, or 7,500 workers/km²) and component-total floor (5,000, 10,000, or 15,000 workers). The paper's primary 5,000/10,000 rule returns 10 centres and 60.2% worker coverage. Sensitivity does not eliminate the modifiable-areal-unit problem.

Figure 5. Employment-centre sensitivity, 2011

Holding the density threshold at 5,000 and raising the component floor from 5,000 to 15,000 reduces the result from twelve centres and 62.1 percent of workers to eight centres and 56.7 percent. Holding the component floor at 10,000 and raising density from 2,500 to 7,500 reduces the included share from 71.1 to 46.6 percent, while the centre count moves only from eleven to ten. The nearly stable count conceals a large contraction in membership. Reporting only the number of centres would therefore give a misleading impression of robustness.

The sensitivity exercise supports a limited conclusion. Kuwait's 2011 workplace structure contains multiple large, contiguous high-density clusters under every rule in the grid. It does not support a claim that there were exactly ten true centres. Ten is the output of the declared primary classification; the share range is the more honest summary of how the rule matters.

5.5 Workplace-residence spatial dissimilarity

Figure 6 maps the signed difference between each lower area's share of establishment workers and its share of occupation-classified residents, then reports the dissimilarity indices. Blue areas contain a larger share of workplaces than of the residence-side population; orange areas contain the reverse. The map does not connect any blue area to an orange one.

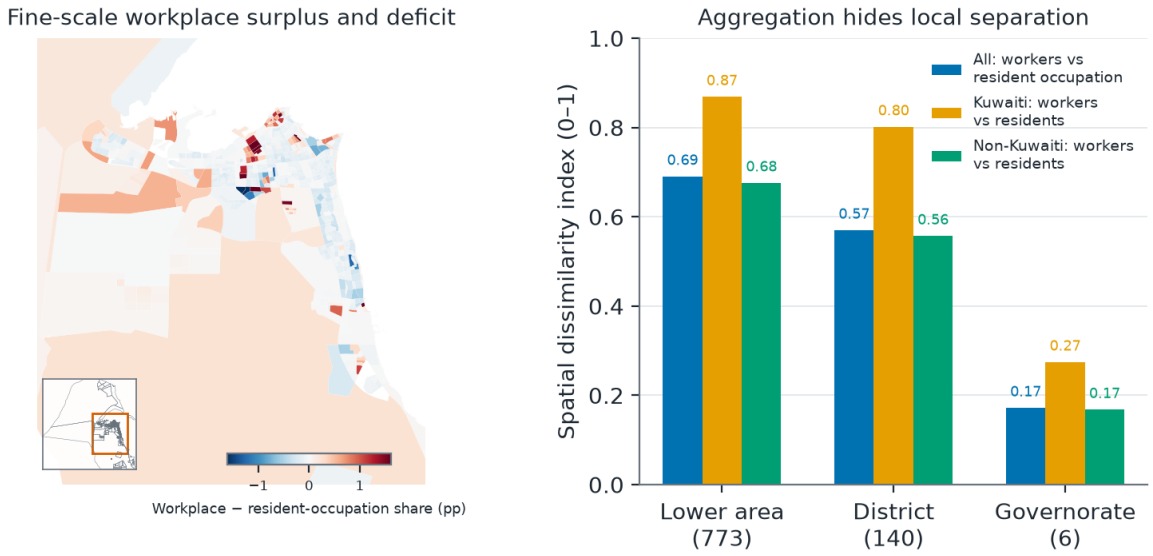

Source: Kuwait Central Statistical Bureau, 2011 Census Geoportal. Left: each polygon's normalized share of establishment workers minus its share of residents classified by main occupation; colors cap at the 99th percentile. Right: one-half the sum of absolute share differences after aggregation. The all-worker comparison uses resident occupation counts; nationality comparisons use total residents because the public file has no resident-worker-by-nationality cross-tab. These are spatial-distribution indices, not commuting flows.

Figure 6. Workplace and residence distributions differ most at fine geography

Table 4. Workplace-residence spatial dissimilarity

Comparison	Lower areas (773)	Districts (140)	Governorates (6)
Establishment workers vs. occupation-classified residents	0.691	0.570	0.172
Kuwaiti establishment workers vs. all Kuwaiti residents	0.869	0.802	0.275
Non-Kuwaiti establishment workers vs. all non-Kuwaiti residents	0.676	0.557	0.168

Note: The nationality rows use all residents by nationality because resident workers cannot be cross-tabulated by nationality in the public file. The index is a distance between normalized spatial shares, not an origin-destination or commuting measure.

The primary all-worker index is 0.691 at the lower-area scale. Put narrowly, the establishment-worker share distribution and occupation-classified resident share distribution are far from identical across the 773 polygons. The index falls to 0.570 when the same counts are aggregated to 140 districts and to 0.172 across six governorates. The large final decline means that workplace-share surpluses and residence-share surpluses frequently occur inside the same governorate and cancel after aggregation.

The centre summaries tell the same story in another form. Sharq-Mirqab-Qibla contains 97,850 establishment workers but 11,791 occupation-classified residents, a workplace-to-residence-side count ratio of 8.3. Shuwaikh Industrial-Rai has a ratio of 30.6. Ardhiya Makhzen exceeds 110, though its residence-side denominator is only 192. By contrast, Farwaniya-Khaitan has a ratio below one: 35,778 establishment workers and 74,648

occupation-classified residents. These are not jobs-to-employed-residents balance ratios in the conventional commuting sense, because the source populations differ. They are transparent summaries of which side of the public file is locally dominant.

The Kuwaiti comparison produces the highest dissimilarity at all scales: 0.869 across lower areas, 0.802 across districts, and 0.275 across governorates. The non-Kuwaiti values, 0.676, 0.557, and 0.168, sit close to the all-worker comparison. It would be tempting to interpret the difference as evidence that Kuwaiti workers commute farther. The data do not support that step. The residence denominators include children, retirees, and other nonworkers, and labor-force participation and source coverage differ by nationality. The nationality rows are therefore secondary spatial comparisons, not behavioral results.

6. Interpretation and limitations

6.1 What the results establish

Three descriptive findings survive the paper's measurement checks. First, employment was concentrated in a limited number of lower-area polygons, and workers were more concentrated than active establishments. Second, Kuwait's 2011 workplace system was multi-centred under every threshold combination examined, although the number and worker share of classified centres varied. Third, the workplace and residence-side distributions looked much more different at lower-area resolution than at governorate resolution.

These findings can discipline later research. A transport study can use the centre membership as candidate destinations, but it still needs network travel times and origins. A land-use study can ask whether particular planning designations predict centre membership, but it needs a credible comparison design. A longitudinal study can ask whether employment persisted or shifted, but it must harmonize boundaries and measurement universes. The present output is a baseline and an internally reproducible set of hypotheses, not a substitute for those designs.

6.2 Why the paper is not causal

There is no treatment and no counterfactual. Density, land use, road access, public employment, sector composition, and residence patterns are jointly determined. Even a precise correlation between employment density and another spatial outcome would be exposed to reverse causality, omitted amenities and infrastructure, sorting, and spatial spillovers. The paper therefore does not estimate effects of concentration on productivity, congestion, land value, segregation, or welfare.

The absence of sampling intervals should not be confused with certainty. The arithmetic uses the entire public census layer, but the classifications depend on thresholds and polygons. Source definitions may omit or classify some work differently. Historical counts can be recorded with error. Those are measurement and design problems rather than sampling variation.

6.3 The public-data ceiling

Six limits define what can be learned from the locally archived file.

First, the reference year is 2011. The paper cannot describe current workplace geography. A current atlas requires a newer establishment source and reviewed cross-vintage boundaries.

Second, the observation is a polygon, not an establishment or worker. The data cannot recover an establishment-size distribution, firm ownership, detailed sector specialization, wages, productivity, or worker attributes within workplaces. Workers per establishment is only a ratio of polygon totals.

Third, the source has no person-level origin-destination links. Workplace-residence dissimilarity cannot identify who commutes, where they go, how far they travel, which mode they use, or how long the trip takes. Calling the statistic a “commuting mismatch” would overstate it.

Fourth, the workplace and residence-side universes differ. Normalizing each distribution makes their geography comparable but does not reconcile the underlying populations. The nationality-specific residence denominators are especially broad.

Fifth, the modifiable-areal-unit problem remains. The threshold grid addresses classification sensitivity, and the scale comparison shows aggregation sensitivity, but neither tests alternative zoning systems at the same scale. Queen contiguity also treats point-touching polygons as connected; a rook-contiguity or distance-based definition could split some components.

Sixth, centre thresholds are not institutional boundaries. A polygon just below 5,000 workers per square kilometre is excluded even if it is economically linked to a nearby centre, while a polygon just above the threshold is included. The complete membership and nine-cell grid allow readers to see the consequence, but no threshold can replace functional-flow data.

6.4 Policy use without overreach

The centre map is most useful as a screening device. It identifies a small set of workplace clusters where questions about access, infrastructure load, and exposure are likely to be quantitatively important in the 2011 baseline. The worker-versus-establishment comparison also shows why policy maps need to match the policy object. Worker safety, transit demand, and daytime exposure may call for worker counts. Business outreach or licensing services may call for establishment counts. Neither distribution is a general measure of economic importance.

The residence comparison adds a scale warning. Governorate averages can make workplace and residence distributions appear relatively aligned even when sharp surpluses sit beside each other locally. A project that allocates transport or services at governorate level could therefore miss the geography relevant for access. That is a research implication, not a policy-effect estimate. Current decisions still require current data.

7. Conclusion

Kuwait's public 2011 Census Geoportal supports a stronger workplace paper than a simple atlas, provided the claims remain within the data. Across 773 lower-area polygons, 715,968 workers attached to 82,930 active establishments were spatially concentrated. Workers were more concentrated than establishments: the top 50 polygons held 60.2 percent of workers and 54.3 percent of establishments, and their inverse-HHI effective polygon counts were 74.5 and 118.7.

A transparent density-and-total rule identifies ten contiguous employment centres containing 430,859 workers, or 60.2 percent of the establishment-worker total. The classification is informative but not unique. Across nine threshold combinations, five to twelve centres contain 37.9 to 71.9 percent of workers. Publishing that range is part of the result, not a robustness footnote.

The workplace distribution also differs sharply from the residence-side distribution at fine scale. Dissimilarity between establishment workers and residents classified by main occupation is 0.691 across lower areas, 0.570 across districts, and 0.172 across governorates. Coarse geography hides much of the local separation. It does not reveal commuting.

The paper's novelty is therefore incremental and internally reproducible from the local archive. Earlier research already mapped employment, jobs and residences, and employment density in Kuwait. This paper contributes an establishment-centric classification, a worker-versus-establishment concentration comparison, a complete threshold sensitivity, and a multi-scale spatial-distance measure. It establishes a defensible historical baseline. Explaining how that structure formed, how it changed, and what it did to productivity or access requires new data and a different empirical design.

Data availability and replication

The official source is the Kuwait Central Statistical Bureau 2011 Census Geoportal. The project preserves the raw GeoJSON byte-for-byte and records its hash, access route, rights status, and local provenance in data/source_manifest_fragment.csv. code/acquire.py checkpoints the documented source files without replacing conflicting bytes. code/build.py validates the source and reproduces all derived data, tables, and six figures. The centre and sensitivity tables are retained in the local research archive. Derived outputs are not represented as freely licensed: sharing remains subject to CSB permission/terms and attribution, which still require confirmation.

References

Adeel, Muhammad, and Reem Alfahad. 2021. *Towards an Equitable Transport System in Kuwait: Understanding the Social and Cultural Context of Transport Accessibility*. LSE Middle East Centre Kuwait Programme Paper Series 14.

Da Cruz, Nuno F., Dhari Alrasheed, Muneerah Alrabe, and Abdullah al-Khonaini. 2024. *Spatial Patterns and Urban Governance in Kuwait: Exploring the Links between the Physical, the Socio-Economic and the Political*. LSE Middle East Centre Kuwait Programme Paper Series 25.

Duranton, Gilles, and Diego Puga. 2003. "Micro-Foundations of Urban Agglomeration Economies." NBER Working Paper 9931.

Fotheringham, A. Stewart, and David W. S. Wong. 1991. "The Modifiable Areal Unit Problem in Multivariate Statistical Analysis." *Environment and Planning A* 23 (7): 1025-1044.

Giuliano, Genevieve, and Kenneth A. Small. 1991. "Subcenters in the Los Angeles Region." *Regional Science and Urban Economics* 21 (2): 163-182.

Kuwait Central Statistical Bureau. 2011. *Census Geoportal: Population, Building, Household, and Establishment Indicators*. Official public geographic data portal; archived locally with source hash and provenance.

McMillen, Daniel P. 2001. "Nonparametric Employment Subcenter Identification." *Journal of Urban Economics* 50 (3): 448-473.

Said, Galal M. 1986. "Urban Development and Transport in Kuwait." Paper presented at the World Conference on Transport Research, Vancouver.