

The Software Developer Index (SDI): A Multi-Dimensional Comparison of Software-Engineering Compensation, Affordability, Labor-Market Conditions, and Career Pay Dispersion Across 31 Countries

Kunal Dhanda

kunaliitkgp09@gmail.com

IIT Kharagpur: Indian Institute of Technology Kharagpur <https://orcid.org/0009-0000-3113-6902>

Research Article

Keywords: software engineering compensation, purchasing power parity, cost of living, composite index construction, double-counting bias, principal component analysis, labor-market tightness, cross-country comparison, data-quality confidence tiers, reproducible methodology

Posted Date: October 9th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11294861/v1>

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

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

The Software Developer Index (SDI): A Multi-Dimensional Comparison of Software-Engineering Compensation, Affordability, Labor-Market Conditions, and Career Pay Dispersion Across 31 Countries

Independent Research Report (v4, revised after three rounds of external methodological review)

October 2026

Abstract

We construct a composite *Software Developer Index* (SDI) ranking 31 countries on the economic conditions facing a working software engineer, combining two directly observed economic indicators – purchasing-power-parity (PPP)-adjusted compensation (levels.fyi, deflated by the World Bank Price Level Index) and a local-affordability ratio (nominal compensation divided by Numbeo’s city-level Cost of Living Plus Rent Index) – with two empirically grounded proxies for labor-market tightness (national unemployment rate) and career compensation headroom (the ratio of 90th-to-25th-percentile pay within the same local market). This report has now gone through three rounds of external review, each of which found a distinct statistical problem. The first round found that the original design rewarded “cheap country” twice (via both PPP-adjusted salary and a separate country-level cost-of-living pillar); fixing it reversed the headline ranking from India-first to United-States-first. The second round found that the fix introduced a different overlap: both remaining compensation-related pillars still contain the same nominal salary figure, so the Pearson correlation between them is $r = 0.90$ across the 31-country sample – compensation is not actually split into four independent dimensions, and the index should be read as a compensation-heavy composite. The third round asked us to confirm this with a model-based test rather than pairwise correlation alone, and to construct an index that explicitly collapses the two correlated pillars into one factor. A principal component analysis confirms it directly: the first three components explain 97.6% of total variance across the four pillars, with the fourth (the one that would be needed for all four pillars to be genuinely independent) contributing only 2.4%. A factor-adjusted index (“Model G”) that averages Compensation and Affordability into one term before combining it with labor-tightness and career-headroom still places the United States first and India third among reliable-sample countries. Combined with three further structural variants that isolate PPP-adjusted salary and the affordability ratio separately, and show the United States ranking first on *both* compensation-related pillars even when they are never combined into one score, this is evidence that the US’s top rank is not purely a mechanical artifact of the pillar overlap, though that overlap does mean the precise margin between the US and India is less tightly estimated than a naively equal-weighted four-“independent”-pillar average would suggest. Tiered data-confidence labels (very low / low / moderate / high, by sample size) are used throughout in place of a bootstrap, since levels.fyi exposes only percentile summary statistics, not individual-level microdata.

Keywords: software engineering compensation; purchasing power parity; cost of living; composite index construction; double-counting bias; principal component analysis; labor-market tightness; cross-country comparison; data-quality confidence tiers; reproducible methodology

1 Introduction

This is the fourth version of this report. We keep the original research question – where does a software engineer get the best overall economic outcome, once compensation, affordability, labor-market conditions, and career upside are all considered together – but the index construction has now been corrected three times, in response to three separate, substantive external reviews. Section 2 summarizes all three rounds of changes and why each was made; everywhere else in the paper describes the current (v4) methodology directly, with each prior version’s result kept visible rather than quietly absorbed.

2 What Changed, and Why

2.1 Round 1: the “cheap country” double-count (v1 → v2)

The central problem: double-counting. The original design computed Pillar 1 as PPP-adjusted compensation (nominal salary divided by the World Bank’s country-level Price Level Index) and, independently, Pillar 2 as an inverted country-level cost-of-living index. Both of these pillars are driven by essentially the same fact about a country – that its local prices are low relative to the United States – so a country with cheap local prices was rewarded for that fact twice in an otherwise equal-weighted four-pillar average, implicitly giving the “cheap country” signal roughly half the total weight of the index rather than one quarter. This was the single biggest driver of the original result (India ranking first, ahead of the United States), and it is a real methodological flaw, not merely a modeling choice.

The fix. In this revision, Pillar 1 (“Compensation”) is unchanged – PPP-adjusted salary via the World Bank Price Level Index, which answers *how much is this salary worth against a common international yardstick*. Pillar 2 is now redefined as “Affordability”: *nominal* (not PPP-adjusted) salary divided by *city-level* Numbeo Cost of Living Plus Rent, which answers *how far does this paycheck stretch in the specific city where the engineer actually lives and shops*. These two pillars now draw on different source data (a whole-economy World Bank price index vs. a city-level Numbeo consumer-price index) at different geographic granularity (country vs. city) answering genuinely different questions (international comparability vs. local purchasing power), rather than both re-deriving the same “is this country cheap” fact.

A second, related problem: geographic mismatch. The original cost-of-living figure was national, while the salary figure was drawn from a single benchmark *city* per country (e.g., Bengaluru for India, San Francisco for the United States). This conflated “the country is cheap” with “the specific tech hub this salary was measured in is cheap,” which can differ substantially (Bengaluru is not representative of India’s cost of living as a whole, any more than San Francisco is representative of the United States’). This revision replaces the national Numbeo figure with the city-level Numbeo figure for the same city used for the salary data, for all 31 countries (Table 1).

Renamed pillars. “Job security” is renamed **Labor-Market Tightness**: the underlying data (national unemployment rate) is a macro labor-market indicator, not a measure of occupation-specific employment protection, dismissal cost, or severance law, and the new name says so directly. “Promotion growth rate” is renamed **Career Compensation Headroom**: the underlying statistic (the ratio of 90th- to 25th-percentile pay within one cross-sectional snapshot of a market) measures

how far senior pay pulls above junior pay in that market, not how quickly a given individual is actually promoted, which no public dataset tracks at this granularity.

Confidence tiers instead of a single cutoff. The original draft flagged countries below 150 levels.fyi submissions as “low confidence” with no stated justification for that particular number. This revision reports four explicit tiers – very low ($n < 50$), low (50–149), moderate (150–499), high ($n \geq 500$) – so the reader can judge sensitivity to sample size directly rather than trusting a single binary cutoff. These tiers are researcher-defined, descriptive data-quality classifications based on submission count, not statistical confidence intervals, and we use the term “confidence tier” rather than “confidence interval” throughout to keep that distinction visible.

On bootstrapping. A genuine nonparametric bootstrap (resampling individual salary observations, recomputing the index, and reporting a confidence interval) was considered and is not possible with the data actually available: levels.fyi’s public interface exposes only four summary statistics per country (25th/50th/75th/90th percentile), not the underlying individual submissions, so there is nothing to resample. Reporting a simulated or distributionally-assumed confidence interval in place of a true bootstrap would itself be a fabrication this report is specifically trying to avoid; the tiered confidence labels above are reported instead as an honest, if less quantitatively precise, substitute.

Softer causal/evaluative language. The original abstract described all four pillars as “real, publicly sourced data pillars,” which is imprecise: two pillars (compensation, affordability) are directly observed from named public sources, while two (labor-market tightness, career headroom) are explicitly labeled proxies for a different, unmeasured quantity. The abstract and conclusion of this revision say so directly, and the conclusion no longer states that any country is simply “the best country for software engineers,” in favor of the more precise comparative statement the data actually support (Section 5).

What did not change in round 1. The underlying raw data collection (levels.fyi percentiles, World Bank series, Numbeo indices, FX rate) was unchanged from v1 to v2 and was independently re-verified against the original source pages; see Section 6 for everything about the data itself that remains a limitation.

2.2 Round 2: compensation is still counted twice (v2 $\rightarrow$ v3)

The problem this time. A second review pointed out, correctly, that fixing the “cheap country” double-count did not fully separate the two compensation-related pillars. Pillar 1 (Compensation) is $\text{Salary}_{\text{nominal}}/\text{PLI}$; Pillar 2 (Affordability) is $\text{Salary}_{\text{nominal}}/\text{City COL}$. Both are nominal salary divided by a price-level measure. Using different price-level data sources (a country-level World Bank index vs. a city-level Numbeo index) makes the two pillars *not identical*, but it does not make them *independent*: if a country’s nominal salary is far above or below the sample average, that single fact pushes both pillars in the same direction, since both have salary as a factor. Different source data reduces redundancy; it does not eliminate it. This is a real, distinct problem from the round-1 issue, not a restatement of it, and we did not catch it ourselves – the round-2 review did.

What we did about it, concretely. Rather than asserting the pillars are “different enough,” Section 5.2 reports the actual Pearson correlation between every pair of the four Model-C pillars

across the 31-country sample. We also add three structural variants (Section 5.4) that do *not* average the two compensation pillars together: Model D keeps only PPP-adjusted compensation (paired with labor-tightness and career-headroom); Model E keeps only the affordability ratio (same pairing); and Model F reports the compensation and affordability pillars as two entirely separate rankings, with no combination at all. If the Model C ordering of the US and India depended entirely on the shared-salary overlap, it should fall apart under at least one of these three variants. Section 5 reports what actually happens.

Two further corrections made in round 2, unrelated to the double-counting issue. The World Bank series used for Pillars 1 and 3 (Price Level Index and unemployment) are both 2024 data, the latest available at the time of writing, while the levels.fyi and Numbeo figures used for Pillars 2 and 4 were collected on 2026-10-07. This report is therefore more precisely described as *a 2026 compensation-and-cost-of-living snapshot combined with the latest available (2024) macroeconomic indicators*, not a single contemporaneous cross-section, and we state this explicitly rather than implying every input shares one reference date. Second, “Affordability” (Pillar 2) is a ratio of nominal salary to a Numbeo index, i.e. dollars of salary per Numbeo cost-of-living-index point – it is *not* disposable income, and does not account for taxes, the nonlinearity of housing costs at different income levels, differences in consumption baskets or savings rates, or health-care and transportation costs that are structured very differently across the 31 countries. A country with a high affordability ratio by this measure is not necessarily a country where an engineer can save more money in practice; it is a country where the reported salary buys more Numbeo-index-weighted local consumption than the same nominal salary would elsewhere.

What did not change in round 2. The round-1 fix (city-level, non-PPP affordability ratio in place of a country-level, double-counted cost-of-living pillar) remains in place; round 2 adds a correlation analysis and structural robustness checks on top of it, rather than reversing it.

2.3 Round 3: confirming the overlap with PCA, and a factor-adjusted index (v3 → v4)

The ask. A third review accepted that Models D, E, and F made the round-2 fix more credible, but pointed out that $r = 0.90$ between two pillars is a pairwise statistic, and that a cleaner statistical answer to “how many independent dimensions does this index actually have” is a principal component analysis on all four pillars together, not a single pairwise correlation. It also suggested a specific structural fix we had not tried: rather than testing the two correlated pillars in isolation (Models D and E) or not combining them at all (Model F), collapse them into a single “compensation factor” *before* averaging, so the index has three equally-weighted terms instead of four unequally-independent ones.

What we did about it. Section 5.3 reports a full PCA on the four Model-C pillars: the first three principal components explain 97.6% of total variance, and the fourth – the one that would be needed for genuine four-dimensional independence – explains only 2.4%. This confirms, with a model-based test rather than a single pairwise number, that the data support “roughly three effective dimensions,” exactly as the round-2 correlation analysis suggested. We also implement the suggested “Model G” (factor-adjusted index, Section 5.4) and report where it ranks the United States and India.

What we did not do. The same review’s suggestion of a parametric/distributional sensitivity analysis on the percentile data (Section 6) is noted as a worthwhile further extension but is not implemented in this version, for the same reason given in round 1 for not fabricating a bootstrap: four percentile points per country support a tiered data-quality classification, not a precisely quantified distributional uncertainty estimate, and presenting one would imply more precision than the underlying data can support.

3 Data Sources

3.1 Compensation: levels.fyi

For each of the 32 countries in our sample we recorded the median, 25th-percentile, 75th-percentile, and 90th-percentile total compensation for the role “Software Engineer” in one benchmark metro/country hub per country (Table 1), collected 2026-10-07 with levels.fyi’s currency toggle set to Indian Rupees, so all countries’ figures are directly comparable in one currency before any conversion. Submission counts (n) range from 12,930 (Bengaluru, India) to 20 (Shenzhen, China); see the confidence tiers in Table 1.

3.2 Pillar 1 input – PPP adjustment: World Bank Price Level Index

Nominal compensation is converted to USD at the 2026-10-07 market rate (96.47 INR/USD) and deflated by the World Bank’s *Price Level Index (GDP)* (PA.NUS.GDP.PLI, US = 100, 2024):

$$\text{Salary}_{\text{PPP-USD}} = \frac{\text{Salary}_{\text{nominal-USD}}}{\text{PLI}/100}. \quad (1)$$

3.3 Pillar 2 input – affordability: Numbeo city-level cost of living

We use Numbeo’s *Cost of Living Index by City, 2026 Mid-Year* (New York City = 100), specifically the city-level *Cost of Living Plus Rent Index*, matched to the exact city used for the levels.fyi salary figure wherever the hub was a named city, and to the country’s principal tech hub where the levels.fyi hub was itself country-wide (Nigeria → Lagos, Egypt → Cairo, Switzerland → Zurich, Japan → Tokyo; flagged explicitly in Table 1). The affordability ratio is:

$$\text{Affordability}_c = \frac{\text{Salary}_{\text{nominal-USD}}}{\text{City Cost of Living Plus Rent Index}_c}. \quad (2)$$

Unlike Pillar 1, this uses *nominal*, not PPP-adjusted, salary – the PPP adjustment is reserved for Pillar 1 only, which is what removes the double-count described in Section 2.

3.4 Pillar 3 input – labor-market tightness

The World Bank/ILO modeled *unemployment, total (% of total labor force)* series (SL.UEM.TOTL.ZS, 2024), inverted so lower unemployment contributes positively.

3.5 Pillar 4 input – career compensation headroom

The ratio of 90th-percentile to 25th-percentile levels.fyi total compensation within the same hub.

Table 1: Benchmark hub for salary (levels.fyi) and matched city for cost of living (Numbeo), per country. *hub mismatch: salary was country-wide, matched to the country’s principal tech hub for COL.

Country	Salary hub (levels.fyi)	COL city (Numbeo)	<i>n</i>	Confidence
India	Greater Bengaluru	Bangalore	12,930	high
United States	San Francisco Bay Area	San Francisco, CA	12,904	high
United Kingdom	Greater London Area	London	3,786	high
Canada	Greater Toronto Area	Toronto	2,909	high
Brazil	Greater Sao Paulo	Sao Paulo	1,572	high
Israel	Tel Aviv Metropolitan Area	Tel Aviv-Yafo	1,495	high
Singapore	Singapore (country)	Singapore	1,447	high
Taiwan	Greater Taipei Area	Taipei	1,139	high
Netherlands	Greater Amsterdam Area	Amsterdam	1,067	high
Ireland	Greater Dublin Area	Dublin	1,052	high
Australia	Greater Sydney	Sydney	990	high
France	Greater Paris Area	Paris	891	high
Switzerland*	Switzerland (country-wide)	Zurich [proxy: salary hub was country-wide]	682	high
Spain	Greater Barcelona Area	Barcelona	596	high
Egypt*	Egypt (country-wide)	Cairo [proxy: salary hub was country-wide]	516	high
Japan*	Japan (country-wide)	Tokyo [proxy: salary hub was country-wide]	436	moderate
Indonesia	Jakarta Metro Area	Jakarta	328	moderate
Sweden	Greater Stockholm	Stockholm	288	moderate
Vietnam	Ho Chi Minh City Metro Area	Ho Chi Minh City	260	moderate
Mexico	Greater Mexico City	Mexico City	258	moderate
Turkey	Istanbul Metropolitan Municipality	Istanbul	242	moderate
Hong Kong	Greater Hong Kong Area	Hong Kong	164	moderate
Germany	Stuttgart Metro Region	Stuttgart	157	moderate
Malaysia	Greater Kuala Lumpur	Kuala Lumpur	135	low
Argentina	Greater Buenos Aires Region	Buenos Aires	134	low
New Zealand	Auckland Region	Auckland	117	low
Austria	Vienna Metropolitan Area	Vienna	105	low
Poland	Tricity (Gdansk)	Gdansk	88	low
Norway	Greater Oslo Region	Oslo	88	low
South Africa	Greater Cape Town Region	Cape Town	47	very low
Nigeria*	Nigeria (country-wide)	Lagos [proxy: salary hub was country-wide]	43	very low
China	Greater Shenzhen Region	Shenzhen	20	very low

4 Methodology

Each of the four pillars is standardized (z-scored) across the 31 countries with complete World Bank series (Taiwan is excluded from the composite for this reason, though its raw figures are retained in the dataset file). The revised Software Developer Index (“Model C”) is the equal-weighted mean:

$$\text{SDI}_c = \frac{1}{4} \left(z_{\text{compensation}} + z_{\text{affordability}} - z_{\text{unemployment}} + z_{\text{career-headroom}} \right)_c. \quad (3)$$

Whether this four-term average actually behaves as four independent dimensions, or fewer, is an empirical question we answer directly in Section 5.2 (pairwise correlation) and Section 5.3 (principal component analysis), rather than assuming either way. We also report the original (“Model A”, naive v1) index exactly as specified in the first draft – PPP salary, country-level cost of living, unemployment, and the same headroom ratio – so the effect of the round-1 fix is directly visible rather than asserted. Equal weighting is a modeling choice, not a derived optimum (there is no ground truth to fit against), so we additionally report a weight-sensitivity analysis across four schemes in Section 5.

To address the round-2 finding that Pillars 1 and 2 share nominal salary as a common factor, we additionally compute three structural variants that do not rely on averaging both compensation pillars together:

$$\text{Model D (PPP-only)} : \frac{1}{3} (z_{\text{compensation}} - z_{\text{unemployment}} + z_{\text{career-headroom}}) \quad (4)$$

$$\text{Model E (affordability-only)} : \frac{1}{3} (z_{\text{affordability}} - z_{\text{unemployment}} + z_{\text{career-headroom}}) \quad (5)$$

$$\text{Model F (separate, unaveraged)} : z_{\text{compensation}} \text{ and } z_{\text{affordability}} \text{ reported as two separate rankings} \quad (6)$$

$$\text{Model G (factor-adjusted)} : CF = \frac{1}{2} (z_{\text{compensation}} + z_{\text{affordability}}); \text{ SDI}_G = \frac{1}{3} (CF - z_{\text{unemployment}} + z_{\text{career-headroom}}) \quad (7)$$

Model G collapses the two correlated compensation measures into a single compensation factor *before* averaging, so that the index has exactly three equally-weighted terms instead of four unequally-independent ones – the most direct structural response to the round-2 finding, suggested by the round-2 review itself.

5 Results

Table 2 reports both models side by side. Two things stand out.

First, **the result reverses**: under Model A (the original, double-counted design), India ranks first (SDI 0.83) and the United States second (0.71); under Model C (the corrected design), **the United States ranks first (SDI 1.63) and India second (0.89)**. This is exactly the direction the external review anticipated, and confirms that the original India-first result was substantially an artifact of counting the same “cheap country” signal twice rather than a robust finding.

Second, the United States’ advantage in Model C is now driven by its affordability ratio, not just its PPP-adjusted salary: at the city level, a San Francisco Bay Area software engineer’s nominal salary relative to San Francisco’s own (very high) cost of living is still favorable once the double-counted country-wide comparison is removed, whereas India’s advantage in Model A was partly an artifact of comparing Bengaluru salaries against all-India (not Bengaluru-specific) living costs, which substantially understated Bengaluru’s own cost of living relative to the rest of India.

Table 2: Model C (revised) ranking, with Model A (naive v1) shown for comparison. [†]low confidence (50–149 submissions); [‡]very low confidence (<50); *hub mismatch (salary hub was country-wide, matched to the country’s principal tech hub for COL).

Rank	Country	SDI (C)	SDI (A, naive)	Comp. PPP (US\$)	Affordability ratio	Unemployment %	Car
1	United States	1.63	0.71	273,864	3,088	4.0	
2	India	0.89	0.83	158,186	2,421	4.2	
3	United Kingdom	0.42	0.15	158,700	1,721	4.4	
4	Israel	0.39	-0.13	156,796	1,879	3.6	
5	Netherlands	0.37	0.01	155,263	1,779	3.7	
6	Switzerland*	0.33	-0.30	148,945	1,631	4.3	
7	Poland [†]	0.33	0.24	141,749	1,871	2.8	
8	Nigeria [‡] *	0.33	1.21	33,950	156	3.0	
9	Ireland	0.32	-0.07	154,628	1,821	4.3	
10	Germany	0.16	-0.13	124,458	1,900	3.4	
11	Canada	0.14	-0.16	121,399	1,932	6.3	
12	Australia	0.11	-0.15	123,665	1,615	3.9	
13	Singapore	0.11	-0.27	143,050	1,101	2.7	
14	Mexico	0.08	0.22	99,945	1,381	2.7	
15	Hong Kong	0.02	-0.21	105,356	1,108	3.0	
16	China [‡]	-0.00	0.15	98,673	1,650	4.6	
17	Japan*	-0.07	0.10	86,080	1,298	2.5	
18	Norway [†]	-0.14	-0.39	113,693	1,425	4.0	
19	Austria [†]	-0.15	-0.28	104,076	1,479	5.2	
20	Brazil	-0.21	0.10	82,857	1,282	6.8	
21	New Zealand [†]	-0.23	-0.32	84,367	1,538	4.7	
22	Argentina [†]	-0.26	0.01	89,622	1,107	7.2	
23	Spain	-0.28	-0.28	120,302	1,503	11.4	
24	Vietnam	-0.32	0.34	43,522	575	1.6	
25	France	-0.37	-0.34	100,688	1,217	7.4	
26	Sweden	-0.45	-0.45	96,192	1,292	8.4	
27	Egypt*	-0.46	0.26	56,288	618	6.8	
28	Turkey	-0.50	-0.13	100,206	989	8.8	
29	Malaysia [†]	-0.53	-0.00	69,350	780	3.9	
30	Indonesia	-0.53	0.12	45,459	663	3.3	
31	South Africa [‡]	-1.03	-0.92	129,725	1,587	32.3	

Table 3: Rank of each country under four alternative weighting schemes applied to Model C

Country	Equal (25/25/25/25)	Compensation-heavy (40/20/20/20)	Affordability-heavy (20/40/20/20)	Car
United States	1	1	1	1
India	2	2	2	2
United Kingdom	3	3	3	4
Israel	4	4	4	3
Netherlands	5	5	5	6
Switzerland	6	7	7	8
Poland	7	8	8	5
Nigeria	8	16	16	20
Ireland	9	6	6	7
Germany	10	10	10	9
Canada	11	11	11	10
Australia	12	12	12	11
Singapore	13	9	9	14
Mexico	14	13	13	13
Hong Kong	15	14	14	17
China	16	15	15	12
Japan	17	19	19	15
Norway	18	17	17	18
Austria	19	18	18	16
Brazil	20	21	21	22
New Zealand	21	22	22	19
Argentina	22	23	23	23
Spain	23	20	20	21
Vietnam	24	27	27	27
France	25	24	24	24
Sweden	26	25	25	25
Egypt	27	28	28	28
Turkey	28	26	26	26
Malaysia	29	29	29	29
Indonesia	30	30	30	30
South Africa	31	31	31	31

5.1 Weight-sensitivity analysis

Table 3 reports each country’s rank under four weighting schemes applied to Model C: equal weights, compensation-heavy, affordability-heavy, and career-heavy. The United States ranks first under every scheme tested. India ranks second or third in three of the four schemes, falling to third only under career-heavy weighting, where Nigeria’s extreme (and, per its “very low” confidence tier, highly unreliable) growth-headroom ratio temporarily pushes it to second place – a result we do not treat as substantive, given it rests on only 43 salary submissions. Excluding Nigeria, the United States/India ordering at the top is stable across all four weighting schemes, which is a meaningfully stronger claim than either model alone could support.

5.2 Pillar correlation: is compensation really split into two independent pillars?

Table 4: Pearson correlation among the four Model-C pillar z-scores (31 countries)

	Compensation	Affordability	Labor-tightness	Career headroom
Compensation	1.00	0.90	-0.04	-0.43
Affordability	0.90	1.00	-0.04	-0.53
Labor-tightness	-0.04	-0.04	1.00	0.04
Career headroom	-0.43	-0.53	0.04	1.00

Table 4 answers the round-2 review’s central question directly. Compensation and Affordability correlate at $r = 0.90$ ($r^2 = 0.81$, i.e. roughly 81% of the cross-country variance in one pillar is linearly shared with the other) – a very high correlation, and clear confirmation that the two pillars are not statistically independent, even though they draw on different source data at different geographic granularity. Labor-tightness is essentially uncorrelated with every other pillar ($|r| \leq 0.04$), so it is a genuinely independent fourth dimension. Career headroom is moderately *negatively* correlated with both compensation pillars ($r = -0.43$ and $r = -0.53$): higher-paid markets tend to have *smaller* relative spread between their 25th- and 90th-percentile pay, which is itself a notable finding we had not anticipated. The practical consequence of the $r = 0.90$ figure is that Model C should be read as having roughly three effective independent dimensions (a compensation factor, labor-tightness, and career headroom) rather than four, with compensation counted at close to double weight relative to the other two – exactly the concern the round-2 review raised, now quantified rather than argued about. Section 5.3 tests this “roughly three dimensions” claim directly with principal component analysis, rather than resting on pairwise correlations alone.

5.3 Principal component analysis

Table 5: Principal component analysis of the four Model-C pillar z-scores ($n = 31$ countries). Loadings are the eigenvector coefficients for each pillar on each component.

Component	Variance explained	Compensation	Affordability	Labor-tightness	Career headroom
PC1	56.6%	-0.61	-0.63	+0.05	+0.47
PC2	24.9%	-0.04	-0.04	-1.00	-0.01
PC3	16.1%	-0.41	-0.25	+0.03	-0.88
PC4	2.4%	-0.67	+0.73	-0.00	+0.11

Table 5 reports a PCA on the four standardized Model-C pillars across the 31 countries with complete data. PC1 explains 56.6% of total variance and loads almost equally on Compensation (-0.61) and Affordability (-0.63), confirming that these two pillars move together as a single underlying factor rather than as two separate ones. PC2 explains a further 24.9% and loads almost entirely on Labor-tightness (-1.00), confirming it as a genuinely separate dimension exactly as the correlation matrix suggested. PC3 explains 16.1% and loads primarily on Career headroom (-0.88). The first three components together explain 97.6% of total variance, with the fourth component – which would be needed to treat all four pillars as truly independent – contributing only 2.4%. This is an objective, model-based confirmation of the “roughly three effective dimensions” claim in Section 5.2: a fourth independent compensation-related dimension does not meaningfully exist in this data, which is exactly why Model G (Section 5.4) collapses the two compensation pillars into one factor rather than averaging them as two.

5.4 Structural robustness: does the US ranking survive un-averaging compensation?

Table 6: Top 15 countries under Model C (both compensation pillars, 1/4 each) vs. Model D (PPP-only, 3 equal pillars), Model E (affordability-only, 3 equal pillars), and Model G (compensation factor = mean of Pillars 1–2, then averaged with the other two pillars, 3 equal weights)

Country	SDI (C)	SDI (D)	SDI (E)	SDI (G)
United States	1.63	1.19	1.02	1.10
India	0.89	0.60	0.87	0.74
United Kingdom	0.42	0.38	0.23	0.30
Israel	0.39	0.25	0.20	0.22
Netherlands	0.37	0.28	0.19	0.24
Switzerland	0.33	0.32	0.19	0.26
Poland	0.33	0.18	0.24	0.21
Nigeria	0.33	1.19	1.01	1.10
Ireland	0.32	0.20	0.14	0.17
Germany	0.16	-0.06	0.14	0.04
Canada	0.14	-0.11	0.13	0.01
Australia	0.11	0.03	0.07	0.05
Singapore	0.11	0.34	-0.07	0.14
Mexico	0.08	0.14	0.21	0.17
Hong Kong	0.02	0.21	0.08	0.15

Table 6 reports Models D, E, and G alongside Model C for the top 15 countries. Four findings stand out.

First, **PPP-adjusted compensation and local affordability are not interchangeable measures**, and the difference between them is itself one of this report’s more interesting findings. In Model E (affordability-only), the United States ranks first (1.02), Nigeria a close second (1.01), India third (0.87). In Model D (PPP-compensation-only), Nigeria and the United States are effectively tied for first (1.19 vs. 1.19) and **India drops all the way to third (0.60)**. The gap between India’s Model D and Model E scores is the largest swing any country shows between the two single-pillar variants, which illustrates concretely why PPP adjustment (an economy-wide, country-level construct) and a city-level affordability ratio can tell quite different stories about the same country: PPP adjustment, by construction, mechanically favors low-price-level economies more strongly

Table 7: Model F: Compensation rank and Affordability rank reported separately (not averaged into one score)

Country	Compensation rank	Affordability rank	Rank difference
United States	1	1	0
United Kingdom	2	9	7
India	3	2	1
Israel	4	5	1
Netherlands	5	8	3
Ireland	6	7	1
Switzerland	7	11	4
Singapore	8	25	17
Poland	9	6	3
South Africa	10	13	3
Germany	11	4	7
Australia	12	12	0
Canada	13	3	10
Spain	14	15	1
Norway	15	17	2
Hong Kong	16	23	7
Austria	17	16	1
France	18	22	4
Turkey	19	26	7
Mexico	20	18	2
China	21	10	11
Sweden	22	20	2
Argentina	23	24	1
Japan	24	19	5
New Zealand	25	14	11
Brazil	26	21	5
Malaysia	27	27	0
Egypt	28	29	1
Indonesia	29	28	1
Vietnam	30	30	0
Nigeria	31	31	0

than a city-matched affordability ratio does, because the latter is anchored to one specific, often atypically expensive, tech hub rather than the whole economy.

Second, removing the affordability pillar in Model D lets the PPP-inflation artifact described in Section 6 (World Bank PPP mechanically favoring the lowest-price-level countries) reassert itself, with Nigeria’s extremely small sample ($n = 43$) producing a near-tie with the United States that we do not treat as a genuine finding.

Third, Model G – which collapses Compensation and Affordability into one factor *before* averaging, rather than averaging four terms of which two are highly correlated – puts the United States first (1.10) and India third (0.74), with Nigeria again a close second (1.10) for the same small-sample reason as in Model D. Excluding Nigeria on data-quality grounds, Model G’s US/India gap (1.10 vs. 0.74) is comparable in direction and magnitude to Model C’s (1.63 vs. 0.89), so collapsing the correlated pair into one factor – the single most direct structural response to the round-2 critique – does not change the ordering.

Fourth, and most directly responsive to the round-2 concern, Table 7 reports compensation rank and affordability rank *without ever averaging them*: the United States ranks **first on both** (rank 1 and rank 1, a rank difference of 0), which is evidence that the US’s strong position in Model C is not purely a mechanical consequence of counting one shared salary figure twice – a country whose top rank depended entirely on that overlap would not necessarily top both un-averaged rankings simultaneously. India ranks third on compensation and second on affordability (difference of 1), also consistent across both un-averaged views. We take this as evidence that the US/India ordering is more than an artifact of the remaining pillar correlation, while still agreeing with the round-2 review that the $r = 0.90$ correlation means the *margin* between the US and India in Model C’s single combined score is less precisely estimated than an equal-weighted four-independent-pillar average would normally imply.

5.5 Interpretation

Combining Sections 5.2 and 5.4, the most defensible summary of the data is: *under the equal-weight Model C, the United States ranks first and India second, and this ordering is stable across four tested weighting schemes and three structural variants that do not average the two compensation pillars together. However, the two compensation-related pillars are mechanically correlated ($r = 0.90$) because both are driven by the same nominal salary figure, so Model C’s single combined score should be interpreted as a compensation-heavy composite with roughly three effective independent dimensions, not as evidence that the US dominates on four statistically independent economic dimensions.* This is deliberately more hedged than either of the two single-sentence claims the data might seem to support at first glance (“India is the best country for software engineers” in v1; “the United States is the best country for software engineers” in v2) and is the direct result of two successive external reviews, each of which found a real problem the previous version had not.

6 Limitations

1. **No public, occupation-specific employment-protection dataset exists at this country coverage.** We attempted to retrieve the OECD’s Employment Protection Legislation index, which would be a better-targeted measure of job security than aggregate unemployment, but could not obtain a complete, machine-readable table across our full country sample (OECD EPL coverage also excludes several countries in our sample, e.g. Nigeria, Vietnam, Egypt). The national unemployment rate is reported instead, and is explicitly named “Labor-Market Tightness” rather than “job security” to reflect what it actually measures.

2. **No public, longitudinal, country-level promotion-tracking dataset exists.** “Career Compensation Headroom” (the p90/p25 ratio) is a cross-sectional dispersion measure, not a measure of how quickly an individual is promoted. A market can have wide pay dispersion because senior roles are rare, not because promotion is fast; the two are conceptually distinct and this report does not have the data to separate them. Nigeria’s data illustrate exactly how unstable this pillar can be: with only $n = 43$ submissions, its p90/p25 ratio (9.24) is roughly double the next- highest country’s, and under career-heavy weighting (Section 5) this single noisy ratio is enough to push Nigeria into second place ahead of India. The same small-sample pattern recurs in Models D and G (Section 5.4), where Nigeria comes within statistical noise of the United States under every specification that gives compensation or career headroom enough individual weight. Across all of these, Nigeria is an interesting statistical signal, but currently a data-quality artifact that needs substantially more salary data before it could support any country-level conclusion – not a finding about Nigeria’s labor market or compensation structure, and not a result we treat as competing with the US/India comparison this report is actually about.
3. **Economy-wide PPP baskets do not match a software engineer’s actual consumption basket.** The World Bank Price Level Index used for Pillar 1 is built from an economy-wide basket that includes many non-tradable, locally produced goods whose prices are structurally low in poorer countries for reasons unrelated to what a software engineer actually buys (often a more internationally-tradable-goods-heavy basket resembling a global tech hub’s). This still mechanically favors the lowest-PLI countries in the sample (Nigeria, Egypt, India, Vietnam, Indonesia) in Pillar 1, even after the Pillar 2 fix described in Section 2; it is a property of the PPP data source itself, not of this report’s index construction, and we do not attempt a post hoc correction, since any correction would itself require an assumption about a software-engineer- specific consumption basket for which no public data exists.
4. **One benchmark hub, or in four cases one whole country, stands in for each country.** Most countries have large internal cost-of-living and salary dispersion; the chosen hub is the one with the deepest levels.fyi coverage, which maximizes salary-data reliability at the cost of treating that hub (rather than the national distribution) as the unit of analysis. For Nigeria, Egypt, Switzerland, and Japan, the levels.fyi salary figure itself was already country-wide rather than city-specific (Table 1), and we matched it to that country’s principal tech hub for the Numbeo affordability figure as the closest available proxy – an approximation we flag rather than hide.
5. **Self-selection in levels.fyi submissions.** Submissions are voluntary and plausibly skew toward large, well-known, often US-headquartered multinational employers, particularly outside the United States, which would bias non-US compensation figures upward relative to the true national distribution of software-engineering pay. A reported PPP-adjusted salary of \$158,186 for India should be read as *the PPP-adjusted median of levels.fyi’s self-selected Bengaluru submitter pool*, not as a claim about the typical Indian software engineer’s purchasing power.
6. **Taiwan is excluded from the composite index** (not a World Bank reporting member); its raw levels.fyi and Numbeo figures remain in the accompanying dataset file.
7. **Equal weighting is a transparent default, not a derived optimum** – addressed directly via the four-scheme sensitivity analysis in Section 5 rather than asserted without testing.

8. **The two compensation pillars are correlated at $r = 0.90$** (Section 5.2) because both are driven by the same nominal salary figure. This means Model C’s four pillars are not four independent dimensions; see Section 5.4 for the structural checks (Models D, E, F) run specifically to test how much this affects the US/India ordering.
9. **Temporal mismatch across data sources.** Compensation and cost- of-living data were collected 2026-10-07; the World Bank Price Level Index and unemployment series are both 2024, the latest available at the time of writing. This report combines a 2026 compensation/cost-of-living snapshot with 2024 macroeconomic indicators, not a single contemporaneous cross-section.
10. **The affordability ratio is not disposable income.** It is nominal salary divided by a Numbeo cost-of-living index, and does not account for taxes, the nonlinearity of housing costs across income levels, differing consumption baskets, savings rates, or health-care and transportation cost structures, all of which vary substantially across the 31 countries.
11. **A parametric/distributional sensitivity analysis was considered but not implemented.** Given only four percentile points per country, one could fit a parametric distribution (e.g. lognormal) to each country’s reported percentiles and examine how sensitive the career- headroom ratio is to the fitting assumption. We flag this as a natural further extension rather than implementing it here, to avoid presenting a distributional assumption with more apparent precision than the underlying four-point data actually support.
12. **Single point-in-time snapshot** (2026-10-07 for compensation and cost of living; 2024 for the two World Bank series, as noted above). Compensation, cost of living, and unemployment all move over time; this is a snapshot, not a trend.

7 Conclusion

Three rounds of external review each found a real, distinct statistical problem in this index, and each fix changed either the result or our confidence in it: removing the “cheap country” double-count (round 1) flipped the ranking from India-first to US-first; quantifying the remaining compensation correlation and testing it structurally (round 2) did not flip the ranking again, but established that the margin behind it rests on roughly three effective independent dimensions, not four; confirming that with a principal component analysis and a factor-adjusted index (round 3) did not flip the ranking a third time either, and gave the “three dimensions” claim a model-based footing rather than a single pairwise correlation. Under the current (v4) methodology: the United States ranks first and India second, this ordering is stable across four weighting schemes, three structural variants that avoid averaging the two compensation pillars together, a factor-adjusted index that collapses them into one term, and separate un-averaged rankings on each individual pillar (Table 7). We regard the three successive corrections, and the fact that each was caught by someone asking where a pillar’s information was actually coming from rather than by any statistical test we ran ourselves at the time, as a more durable finding than either country’s rank: a composite index with four plausible-sounding, differently-sourced pillars can still have two of them sharing most of their variance, and that is not something equal weighting, sensible- looking source diversity, or even one round of sensitivity analysis will necessarily reveal on its own – it took three.

Data and Reproducibility

The complete dataset is in `sdi_dataset.json`, generated by `build_index.py`, with two top-level keys: `countries` (one entry per country, documented field-by-field in Table 8) and `pillar_correlation_matrix` (the full 4×4 Pearson correlation matrix underlying Table 4).

References

- [1] levels.fyi. *Software Engineer Salary by Location*. Retrieved 2026-10-07 from <https://www.levels.fyi/locations>.
- [2] Numbeo. *Cost of Living Index by Country, 2026 Mid-Year*. Retrieved 2026-10-07 from https://www.numbeo.com/cost-of-living/rankings_by_country.jsp.
- [3] Numbeo. *Cost of Living Index by City, 2026 Mid-Year*. Retrieved 2026-10-07 from <https://www.numbeo.com/cost-of-living/rankings.jsp?title=2026-mid>.
- [4] World Bank. *Price level index (GDP)*, indicator PA.NUS.GDP.PLI. Retrieved 2026-10-07 from <https://api.worldbank.org/v2/country/all/indicator/PA.NUS.GDP.PLI>.
- [5] World Bank/ILO. *Unemployment, total (% of total labor force) (modeled ILO estimate)*, indicator SL.UEM.TOTL.ZS. Retrieved 2026-10-07 from <https://api.worldbank.org/v2/country/all/indicator/SL.UEM.TOTL.ZS>.
- [6] OECD. *Indicators of Employment Protection*. Coverage and access limitations discussed in Section 6.
- [7] United Nations Development Programme. *Human Development Index: Technical Notes*. Methodological precedent for equal-weighted composite index construction.

Table 8: Data dictionary for `sdi_dataset.json`

Field	Description
<code>country</code>	Country name (key)
<code>hub</code>	levels.fyi benchmark metro/country hub used for salary data
<code>numbeo_city</code>	Numbeo city used for cost-of-living data (matched to hub where possible)
<code>hub_mismatch</code>	true if the salary hub was country-wide and COL city is a proxy (Nigeria, Egypt, Switzerland, Japan)
<code>n_submissions</code>	levels.fyi salary submission count for the hub
<code>confidence_tier</code>	descriptive data-quality tier (not a statistical CI): high ($n \geq 500$) / moderate (150–499) / low (50–149) / very low (<50)
<code>median_salary_inr</code>	levels.fyi median total comp, INR (site currency toggle), collected 2026-10-07
<code>salary_nominal_usd</code>	<code>median_salary_inr</code> / 96.47 (2026-10-07 FX rate)
<code>price_level_index</code>	World Bank PA.NUS.GDP.PLI, country-level, US=100, 2024 (latest available)
<code>salary_ppp_usd</code>	<code>salary_nominal_usd</code> / (<code>price_level_index</code> /100) – Pillar 1 (Compensation)
<code>city_cost_of_living_plus_rent</code>	Numbeo city-level Cost of Living + Rent Index, NYC=100
<code>affordability_ratio</code>	<code>salary_nominal_usd</code> / <code>city_cost_of_living_plus_rent</code> – Pillar 2 (Affordability)
<code>unemployment_rate</code>	World Bank/ILO SL.UEM.TOTL.ZS, % of labor force, 2024 – Pillar 3 input (inverted)
<code>growth_ceiling_p90_p25</code>	levels.fyi p90/p25 total-comp ratio within the hub – Pillar 4 input
<code>sdi_model_a_naive</code>	v1 index: mean $z(\text{PPP salary, } -\text{country COL, } -\text{unemployment, growth ratio})$
<code>sdi_model_c_revised</code>	v2 index: mean $z(\text{PPP salary, affordability ratio, } -\text{unemployment, growth ratio})$
<code>sdi_model_d_ppp_only</code>	v3 variant: mean $z(\text{PPP salary, } -\text{unemployment, growth ratio})$ – no affordability pillar
<code>sdi_model_e_affordability_only</code>	v3 variant: mean $z(\text{affordability, } -\text{unemployment, growth ratio})$ – no PPP pillar
<code>model_f_compensation_z / _affordability_z</code>	Pillars 1 and 2 reported as separate, un-averaged z-scores
<code>sdi_model_g_factor_adjusted</code>	v4 variant: mean($z(\text{PPP}), z(\text{affordability})$) collapsed to one compensation factor, then mean(factor, labor, career)
<code>weight_sensitivity_rank</code>	rank of this country under each of 4 alternative weighting schemes (Model C)
<code>pillar_correlation_matrix</code>	top-level key (not per-country): Pearson r among the 4 Model-C pillar z-scores
<code>pca</code>	top-level key (not per-country): PCA explained variance % and loadings for the 4 Model-C pillars