INTRATEC SOLUTIONS

intratec

METHODOLOGY · JUNE 2026

Report Methodology

Industry Economics & Competitiveness

33	14	7
COUNTRIES	PILLARS	INDUSTRIES

CHAPTER 1

What the Report Measures

SUMMARY

- Industry Economics & Competitiveness scores manufacturing competitiveness across 33 countries and 7 commodity families on a consistent monthly basis.
- A 14-pillar architecture decomposes competitiveness into separately scored dimensions — 8 Base Pillars for the shared operating environment and 6 Industry-Specific pillars per commodity family.
- Three outputs serve different needs: pillar and composite scores (0–100, relative), ranks (ordinal position 1–33), and reference datasets (absolute costs and indexes).

Industry Economics & Competitiveness is a monthly intelligence program on manufacturing competitiveness for non-agricultural commodity industries. It covers 33 countries and 7 industry sectors, and each edition focuses on a single country, positioning it within a consistent, comparative framework that is regenerated every month from the latest available data — so the analysis tracks shifting cost and structural conditions rather than freezing into an annual snapshot. The report serves investors, policymakers, procurement executives, EPCs, cost estimators, and financial modelers.

The framework at a glance

33	7	14	Monthly
countries benchmarked	commodity families	competitiveness pillars (8 Base + 6 Industry-Specific)	updated each edition

Figure 1.1 Industry Economics & Competitiveness benchmarks 33 countries across 7 commodity families and 14 pillars, regenerated every month.

Three principles govern the design:

* **Objective measurement** — the analysis relies on verifiable, quantitative data: costs, volumes, rates, macroeconomic metrics, and policies.

* **Standardized assessment** — identical methods are applied across every country and every month, so cross-country comparisons stay valid.
* **Actionable synthesis** — multi-dimensional data is distilled into intuitive 0–100 scores and explicit rankings, bridging raw data and strategic insight.

What competitiveness means

Competitiveness here is the relative cost and structural advantage of producing a given commodity in one country versus the others in the benchmarking set — the fixed group of countries every assessment is measured against. Lower production cost and stronger structural conditions mean higher competitiveness.

The measure is comparative, not absolute. A country is always positioned against the full set of peers rather than against a fixed external standard, so a result answers "more or less competitive than the others," not "competitive in absolute terms." That comparative basis is what keeps month-to-month and country-to-country readings consistent.

What it covers

The report spans seven commodity families, each treated as a distinct industry. Together these non-agricultural, commodity-based sectors form the upstream and midstream nodes of almost all global value chains. Because the commodities are highly commoditized, their local pricing and margin dynamics act as sensitive barometers of a country's structural manufacturing costs. Sector competitiveness is evaluated on the dominant production route used in the focus country — for example, naphtha-based versus ethane-based cracking.

Table 1.1 The seven commodity families, each treated as a distinct industry.

INDUSTRY SECTOR	INDUSTRY SECTOR
Olefins & Derivatives	Polymers
Aromatics & Derivatives	Fertilizers
Alcohols & Organic Acids	Inorganic Chemicals
Metals	

The benchmarking universe

The report scores 33 countries individually, forming a fixed benchmarking universe. Each published edition focuses on one country and benchmarks it against the other 32. Holding the universe constant is deliberate: it keeps comparisons stable, so a change in a country's standing reflects a change in underlying conditions rather than a change in the comparison group.

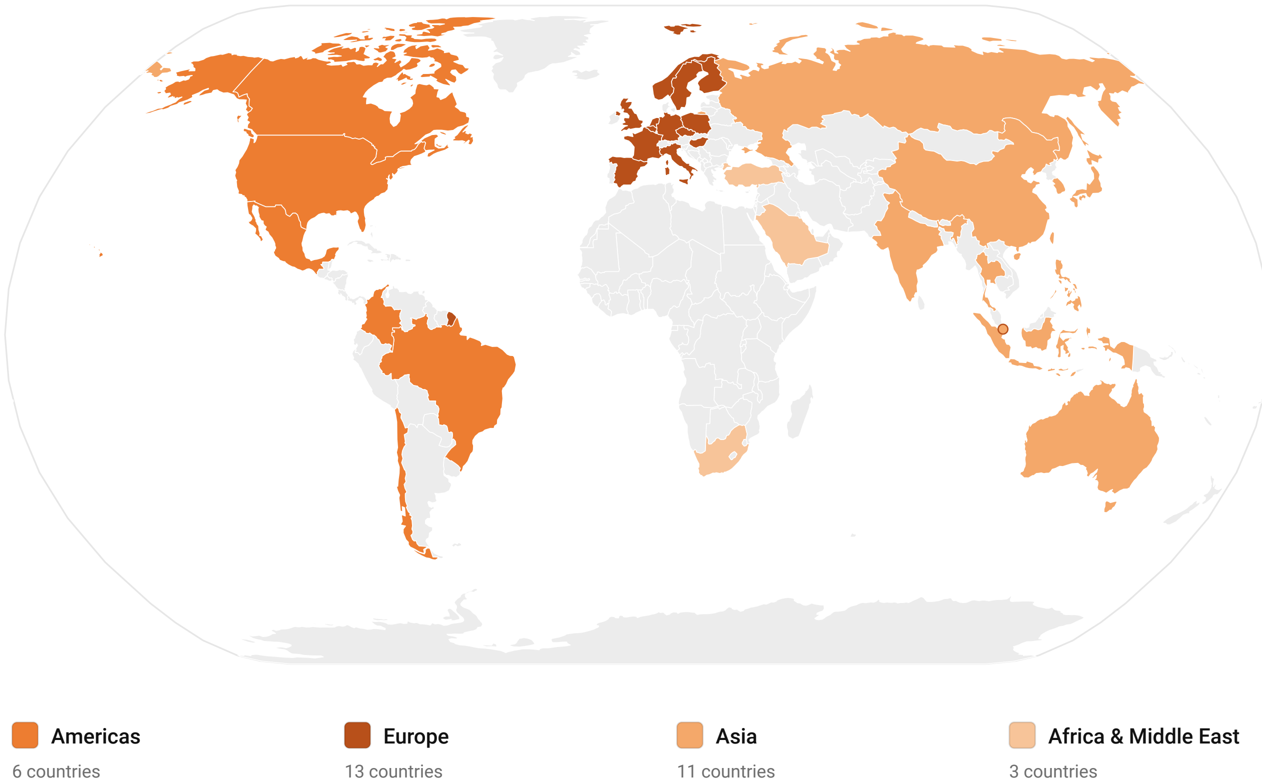

Figure 1.2 The 33 benchmarked countries, shaded by Intratec world region — Americas, Europe, Asia, and Africa & Middle East. The universe is held constant across editions, so a shift in a country's standing reflects underlying conditions rather than a changed comparison group.

The set is curated, not open-ended. Together the 33 countries account for more than 80% of global manufacturing output, with enough statistical depth and geographic diversity to make the 0–100 scores meaningful. They combine three complementary profiles: mature, high-income industrial hubs (such as the United States, Germany, and Japan); major resource-rich exporters (such as Saudi Arabia and Canada); and dynamic emerging supply-chain nodes (such as Mexico, India, and Brazil). The size of the set matters methodologically: 33 countries give a statistically significant baseline for what counts as "average," "leading," or "lagging" — the distribution against which every score is normalized.

Table 1.2 The fixed 33-country benchmarking universe, by region.

REGION	COUNTRIES
Americas	Brazil, Canada, Chile, Colombia, Mexico, United States
Europe	Belgium, Czech Republic, Finland, France, Germany, Hungary, Italy, Netherlands, Norway, Poland, Spain, Sweden, United Kingdom
Asia	Australia, China, India, Indonesia, Japan, Philippines, Russia, Singapore, South Korea, Taiwan, Thailand
Africa & Middle East	Saudi Arabia, South Africa, Turkey

Because the same 33 countries anchor every edition, a score is always read against a stable backdrop: the country in focus this month is itself one of the 32 comparators in every other country's edition, so a single consistent distribution underlies the entire program. That is what lets a procurement team weigh a Mexico edition against a Germany edition without correcting for a shifting peer group — the yardstick does not move from one edition to the next, only the country held up against it.

Three kinds of output

A reader meets three distinct outputs throughout the report. Distinguishing them early prevents the most common confusion.

* A **pillar** is one dimension of competitiveness — such as labor or energy — evaluated on a 0–100 basis. The framework uses 14 pillars: 8 Base Pillars, shared across all commodities, and 6 Industry-Specific pillars, defined per commodity family. The pillars are the subject of Chapters 3 through 5.
* A **score** is the 0–100 value summarizing performance on a pillar or overall, where 50 represents the global mean and higher values indicate greater competitiveness. Scoring is the subject of Chapter 6.
* A **reference dataset** is an absolute-value dataset — such as actual costs or indexes — provided alongside the scores but outside the 0–100 framework. Reference datasets are the subject of Chapter 7.

In short, scores answer the strategic, relative question — how competitive is a country versus its peers — while reference datasets answer the tactical, absolute question — what something actually costs. The rest of this document follows the journey of a number from raw input to finished score and dataset.

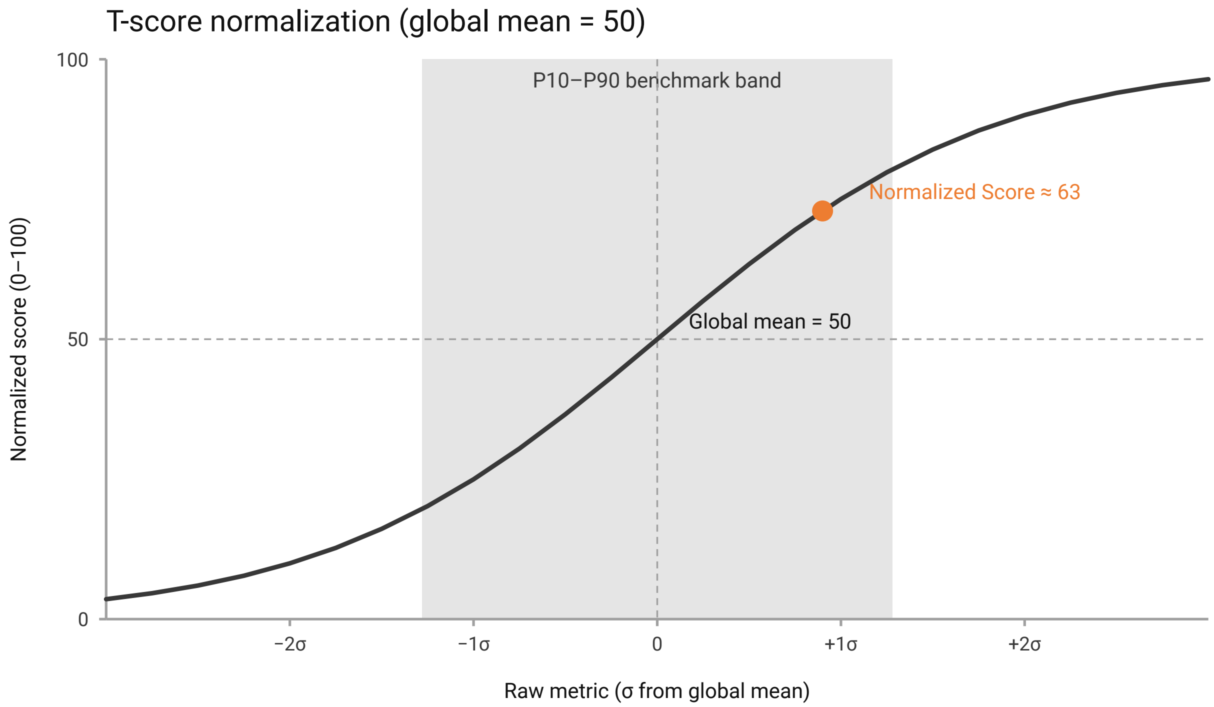

Figure 1.3 The 0–100 scale: 50 marks the global mean, about one standard deviation maps near 60 and 40, and the curve compresses toward the extremes.

Every published figure passes through the same two-layer validation applied across all Intratec assessments — automated checks on the full dataset, plus expert human review of a risk-guided sample. Chapter 2 describes that process as far as a reader needs it to trust the scores. The report is regenerated every month from the latest available data, so pillar scores, rankings, and reference datasets all update together; a single yardstick — 33 countries, 14 pillars, and one fixed 0–100 scale — applies to every edition.

NOTE

Shared conventions. The geographic regions, the currency anchor and exchange-rate references, and the unit system that underlie every figure are conventions applied across all Intratec assessments, not specifics of this report. They are documented on the Intratec methodology site.

CHAPTER 2

How the Data Reaches the Report

SUMMARY

- Every figure moves through a fixed monthly sequence: collection from 350+ sources, preparation and alignment to a common basis, and two-layer validation before publication.
- Official public data from national statistics agencies takes priority; private sources supplement rather than override it.
- A two-layer validation — automated cross-reference checks followed by risk-guided expert review — ensures each published value is a vetted figure, not a raw data pull.

Every score and reference figure in Industry Economics & Competitiveness begins as raw data gathered from outside sources, then moves through a fixed monthly sequence — collection, preparation, alignment to a common basis, and two-layer validation — before it is published. This chapter follows that path only as far as a reader needs it to trust the numbers; the full sourcing, preparation, and price-assessment detail lives on the Intratec methodology site.

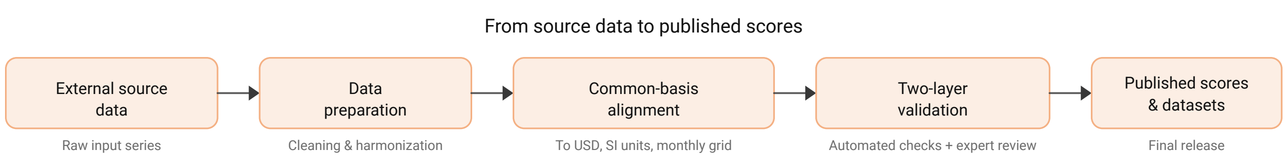

Figure 2.1 Every figure moves through a fixed monthly path — collection from outside sources, preparation, alignment to a common basis, and two-layer validation — before publication.

Where the data comes from

The analysis draws on six categories of source, with the vast majority of inputs coming from public, official records. The emphasis on official data is deliberate: figures published by national authorities are verifiable and free of any single party's commercial interest.

Table 2.1 The six categories of source behind every assessment.

SOURCE CATEGORY	WHAT IT PROVIDES
National statistics bureaus	Official trade, price, and economic statistics
Governmental agencies	Regulatory, fiscal, infrastructure, and labor data
International multilateral organizations	Cross-country economic indicators and technical references
Commodity exchanges	Quoted commodity and derivatives prices — settlement prices, spreads, freight quotations
Producers' data	Prices of commodities as sold and published by the producers themselves
Technology licensors	Published consumption figures, investment levels, and plant capacities, verified before use

Government-reported trade statistics form the foundation; private sources — private trading platforms, producers' data, and licensors' publications — are used only with permission from the rights holders, and supplement rather than override official records. The report does not rely on its own surveys except in very exceptional circumstances, which keeps subjective influence and self-reporting bias out of the assessments.

When more than one source covers the same data point, official and recognized institutions take priority for final figures, and **accuracy is favored over speed**: a slower official source outranks a faster private one. The collection portfolio counts **350+** sources and is expanded continuously — pertinent new sources are mapped and evaluated, and any discontinued source is replaced quickly so coverage stays uninterrupted.

> NOTE
>
> The official-first principle is disclosed openly. The internal parameters used to compare and select among specific sources are not published.

The monthly collection window

Data collection runs on a fixed schedule. Automated systems gather figures from official sources during the **last week of each monthly cycle**, and processing and assessment then take place at the start of the following month, before the edition is published. The same cutoff applies every month, so each edition rests on a consistent, repeatable collection point.

Official trade statistics are generally released **one to three months** after the period they describe — a normal feature of official data, since governments need time to compile, validate, and publish national figures. The reporting calendar is built around that rhythm rather than treating it as a shortcoming. When an official source lags, is temporarily unavailable, or is discontinued, internally developed mathematical models generate interim estimates so the published series never breaks. Where the available inputs are insufficient or unsuitable for a representative figure, analysts step in — estimating from a wide range of factual market information and proposing model adjustments — rather than leaving a gap. Every modeled value carries a status label — **Preliminary** or **Forecast** — so an interim estimate is always distinguishable from a settled official figure (see Chapter 8 for what those labels mean and how figures settle to Final).

Preparing the raw figures

Before any value enters an assessment, it passes through a preparation stage. Source data arrives in many file and number formats, each provider with its own conventions, so every incoming figure is cleaned, its format standardized to a single internal convention, validated against expected ranges, and screened for anomalies. Only figures that clear these steps move forward, so downstream calculations rest on a consistent, vetted foundation.

Official trade statistics also aggregate transactions that differ widely in volume, specification, and delivery conditions. Mixing them indiscriminately would distort the assessed level, so trades are filtered and clustered along four categories before use:

* **Minimum-volume thresholds** — trades too small to represent the market level are excluded.
* **Specification and grade** — only trades matching the assessed specification are kept together.
* **Location basis** — trades are grouped by the location basis they reflect.
* **Trade size and delivery terms** — clustering prevents distortions from mixing different sizes or delivery conditions.

The filter categories are disclosed; the specific numeric thresholds and the statistical techniques applied within them are proprietary.

> **Example:** a single unusually large cargo mixed into a month of smaller spot trades would pull the aggregate away from the mid-market level; clustering keeps such trades from skewing the assessment.

A common, comparable basis

Cleaned figures are then aligned so that any value can be compared against any other. Internally, every monetary figure is standardized to **US dollars** using a single set of reference exchange rates, and every physical quantity to the **International System of Units (SI)** using fixed conversion factors; series published less often than monthly — quarterly or annual official series — are interpolated onto the monthly grid, so all series share one timeline and any month can be compared across the full dataset. The currency anchor, the exchange-rate references, and the unit conventions themselves are shared across all Intratec assessments and are documented on the methodology site.

Two-layer validation

A figure is validated in two layers before publication, both run every month. First, an **automated stage** cross-references the same data point across multiple independent official records to confirm it agrees, and applies mathematical models to flag inconsistencies — applying the same consistent rules each month removes manual-handling error and participant bias. Second, **expert human review** acts as the final quality gate: market analysts with sector-specific knowledge review the model outputs and add the contextual judgment that automated rules alone cannot provide.

That human review is **risk-guided rather than random**. Any value the automated stage flags as anomalous is always routed to an expert; a sample of the remaining, unflagged values is also reviewed each month, an independent check that the automated rules are not silently missing problems. The combination concentrates expert attention where an issue is most likely while keeping the whole pipeline under human oversight — which is why a published score can be read as a vetted figure rather than a raw data pull.

CHAPTER 3

The 14-Pillar Architecture

SUMMARY

- Competitiveness is decomposed into 14 separately scored pillars — 8 Base Pillars shared across all sectors and 6 Industry-Specific pillars scored per commodity family.
- From the pillar level upward, all combination is at equal weight: Base Pillars average into a Base Pillars Score, Industry-Specific pillars into an Industry Composite, and the two combine into each family's Industry Score and the country's Overall Score.
- The combination logic above the pillar level is fully disclosed; the parameters that synthesize observed metrics into a pillar score are proprietary.

Industry Economics & Competitiveness does not reduce a country to a single number and stop there. It decomposes industrial competitiveness into 14 distinct, separately scored dimensions — its **pillars** — and only then synthesizes them upward into the composite scores a reader sees first. This chapter sets out that architecture: what a pillar is, why the framework is built from many pillars rather than one figure, the split between Base and Industry-Specific pillars, and the aggregation pyramid that carries observed data all the way to a country's Overall Score. Chapters 4 and 5 then describe the pillars one by one; Chapter 6 explains how to read the scores they produce.

What a pillar is

A **pillar** is one distinct dimension of industrial competitiveness — a thematic group of related metrics, such as labor costs, energy and utilities, or the tax environment, synthesized into a single 0–100 score. Each pillar isolates one driver of a country's standing so it can be measured, scored, and compared on its own terms. Together, the pillars form the structural backbone of every assessment: individual pillar scores roll up into industry-level and country-level composites, so the same building blocks support both a granular view of one driver and the headline ranking.

Why pillars instead of one number

The bottom line: a single overall figure conceals exactly the information a reader needs. Industrial competitiveness is shaped by many independent drivers — cost structures, infrastructure endowment, macroeconomic stability, market scale, and fiscal conditions — and collapsing them into one value hides where a country is genuinely strong or weak.

Pillars expose the structure behind a ranking. Each one scores a separate driver on the same 0–100 scale, with a higher score always indicating a more competitive position, so the reasons a country places where it does can be traced and individual drivers compared across countries. A single composite also cannot separate conditions that apply to every industry from economics specific to one sector. The two-layer pillar architecture keeps these apart: Base Pillars capture the country's structural operating environment, while Industry-Specific pillars capture sector economics such as input prices, conversion margins, and tariff exposure. Decomposing performance into separate, comparable drivers — rather than collapsing it into one value — is what makes the index transparent and actionable.

> **Example:** A country may score well above the global mean of 50 on Energy & Utilities Costs while scoring poorly on Logistics & Infrastructure. An overall figure would average these into an unremarkable mid-range number; the pillar view shows that the country is a cheap place to run an energy-intensive plant but an expensive place to move product to market.

Base and Industry-Specific pillars

The 14 pillars divide into two layers. The 8 **Base Pillars** score the country's structural operating environment once and are shared across all industries. The 6 **Industry-Specific pillars** capture sector economics and are scored separately for each of the 7 commodity families. The split reflects a design assumption that runs through the whole framework: structural conditions and sector-specific economics each carry weight, and neither dominates competitiveness on its own.

The Base Pillars capture conditions that apply equally to every industry — the operating environment a plant faces regardless of what it makes.

Table 3.1 The 8 Base Pillars — the country's structural operating environment, scored once and shared by all industries.

BASE PILLAR	WHAT IT CAPTURES
Manufacturing Labor Costs	The cost of the manufacturing workforce
Construction Labor Costs	The cost of the construction workforce
Capital & Construction Costs	The cost of building and equipping a plant
Energy & Utilities Costs	The cost of power, fuel, and industrial utilities
Logistics & Infrastructure	The endowment and cost of moving goods within the country
Freight Costs	The cost of international freight to and from the country
Macroeconomic Environment	Macroeconomic stability and conditions
Domestic Tax Environment	The fiscal and tax conditions facing producers

The Industry-Specific pillars capture economics that differ from one commodity family to the next — input prices, conversion margins, trade exposure, and the scale of in-country demand. Each is scored separately for each of the 7 commodity families.

Table 3.2 The 6 Industry-Specific pillars — sector economics, scored separately for each of the 7 commodity families.

INDUSTRY-SPECIFIC PILLAR	WHAT IT CAPTURES
Commodity Prices	Local price levels of the commodity
Feedstock-to-Product Margins	The margin between feedstock and product
Industrial Production	The scale and activity of domestic production
Global Trade Integration	How connected the sector is to international trade
Tariff Protection & Market Access	Tariff exposure and access to markets
Domestic Market Size	The scale of in-country demand for the commodity

The aggregation pyramid

Every pillar score is built through the same bottom-up hierarchy, and the same hierarchy then carries pillar scores up to the headline numbers. The structure is stated openly; only the internal parameters that synthesize the lowest levels are held back.

A score begins as **observed metrics** — real-world quantities in their own units, such as a wage rate or an energy price. Those raw values are normalized across the fixed 33-country universe into **0–100 indicators**, scaled so that 50 represents the global mean. Indicators are grouped into **components**, and components are synthesized into a pillar's final **0–100 score**. Up to this point the synthesis is proprietary; from the pillar level upward, the combination logic is simple and disclosed.

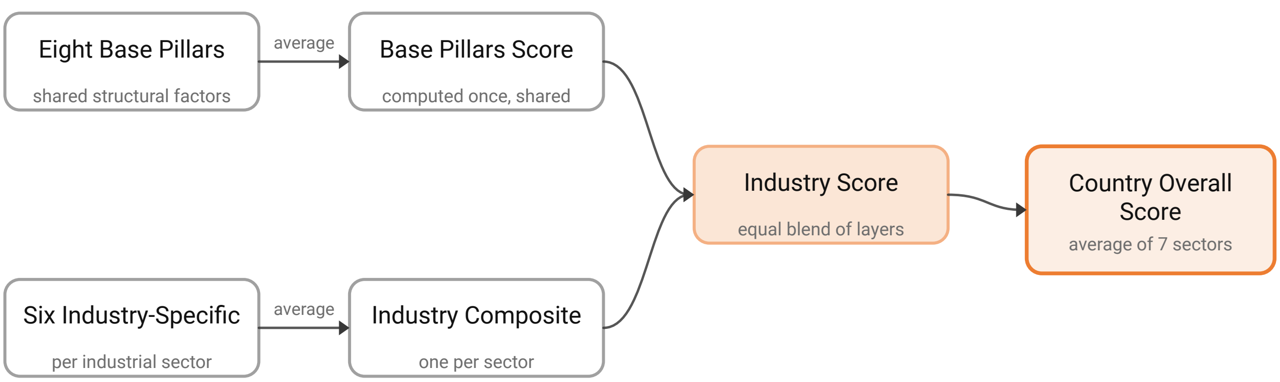

Figure 3.1 Pillar scores roll up in four equal-weight stages: to the Base Pillars Score and per-sector Industry Composites, then to each Industry Score, then to the Country Overall Score.

From the pillar level upward the combination proceeds in four steps, all at equal weight:

* The 8 Base Pillar scores average into a single **Base Pillars Score**, shared across all sectors.
* For each commodity family, the 6 Industry-Specific pillar scores average into an **Industry Composite**.
* For each commodity family, the Base Pillars Score and that sector's Industry Composite combine into the **Industry Score** — the central output of the index, one per sector.
* The Base Pillars Score combines with the **Industry-Specific Pillars Score** (the average of the seven Industry Composites) to give the country's single **Overall Score**.

Equal weighting at every step above the pillar level is the practical expression of the design assumption stated earlier: structural conditions and sector-specific economics are treated as equally important, so neither layer is allowed to dominate the result.

One further roll-up runs alongside this main chain. Each Industry-Specific pillar is also averaged at equal weight across the seven commodity families into a single cross-industry value for that pillar. This side roll-up is computed separately from the four steps above and does not feed the Industry Scores or the Overall Score; it exists to give a standalone, pillar-level read of one dimension across the whole commodity set, rather than within a single family. The Base Pillars need no such variant — they are already scored once per country and shared across every family, so a single value serves as both the family-level input and the cross-industry read.

> **NOTE**
>
> Structure, the 0–100 scale, and the combination logic from the pillar level upward are stated openly. The internal parameters that synthesize observed metrics, indicators, and components into a pillar score are not disclosed.

This pyramid is what later chapters build on. Chapters 4 and 5 take each of the 14 pillars in turn — what it measures, what data feeds it, how it enters the score, and a short example. Chapter 6 returns to the apex of the pyramid to explain how the resulting scores, ranks, and rankings are read.

CHAPTER 4

Base Pillars

SUMMARY

- The 8 Base Pillars score a country's structural operating environment once, shared across all 7 commodity families — from labor and capital costs through energy, logistics, macroeconomic stability, and taxation.
- Cost-based pillars invert the underlying value so a lower cost always yields a higher score; 50 marks the global mean on every pillar.
- Most pillars normalize against a twelve-month trailing average; the Domestic Tax Environment uses current-period values because tax rates change by policy step rather than continuously.

The eight **Base Pillars** describe the structural operating environment a process plant faces in a country, before any single commodity is considered. They are scored once per country and shared across all 7 commodity families, so they capture the cost-and-conditions backdrop that any industrial investor inherits regardless of what the plant makes. Chapter 3 placed these eight inside the wider 14-pillar architecture; this chapter takes each in turn.

Every Base Pillar is described with the same four beats: what it **measures**, the **data** categories that feed it, its **role in the score**, and a short **example**. Across all eight, each pillar resolves to a single 0–100 score and, for the cost pillars, inverts the underlying value so that "better for a manufacturer" always points upward; neither the exact normalization transform nor any below-pillar weight is disclosed. The one normalization exception — the Domestic Tax Environment pillar's use of current-period rather than trailing values — is explained in that pillar's section, alongside its Chapter 5 counterpart.

The eight pillars fall into three themes. **Input and build costs** group the four cost pillars a plant carries to staff and build — Manufacturing Labor Costs, Construction Labor Costs, Capital & Construction Costs, and Energy & Utilities Costs. **Moving goods** groups Logistics & Infrastructure and Freight Costs, the cost and reach of bringing inputs in and getting product to market. The **macro-fiscal setting** groups the Macroeconomic Environment and the Domestic Tax Environment, the country-level stability and fiscal burden every producer works within. All three themes combine at equal weight into the single Base Pillars Score.

Manufacturing Labor Costs

Measures. This pillar measures the cost of the manufacturing workforce — the wage rates and total employer labor cost for manufacturing roles, weighed against how much each worker produces.

Data. The pillar draws on a reference dataset of absolute labor costs and worker output across the 33 countries. Labor cost builds in four steps:

* **Direct Pay** — base wages plus overtime, bonuses, and allowances: the cash components paid directly to the worker.
* **Indirect Pay** — voluntary and customary benefits (health insurance, training, housing) and legally required benefits and employer taxes: paid leave, retirement and social-security contributions, unemployment insurance, workers' compensation, and payroll taxes.
* **Total Employer Cost per hour** — Direct Pay plus Indirect Pay, the pillar's cost input.
* **Adjusted Labor Cost per effective hour** — Total Employer Cost scaled by a **productivity factor** (manufacturing value added per employee, anchored at 1.0 for the most productive position, the United States).

Source categories are national statistics bureaus and government labor agencies; figures are stated on a common monetary basis.

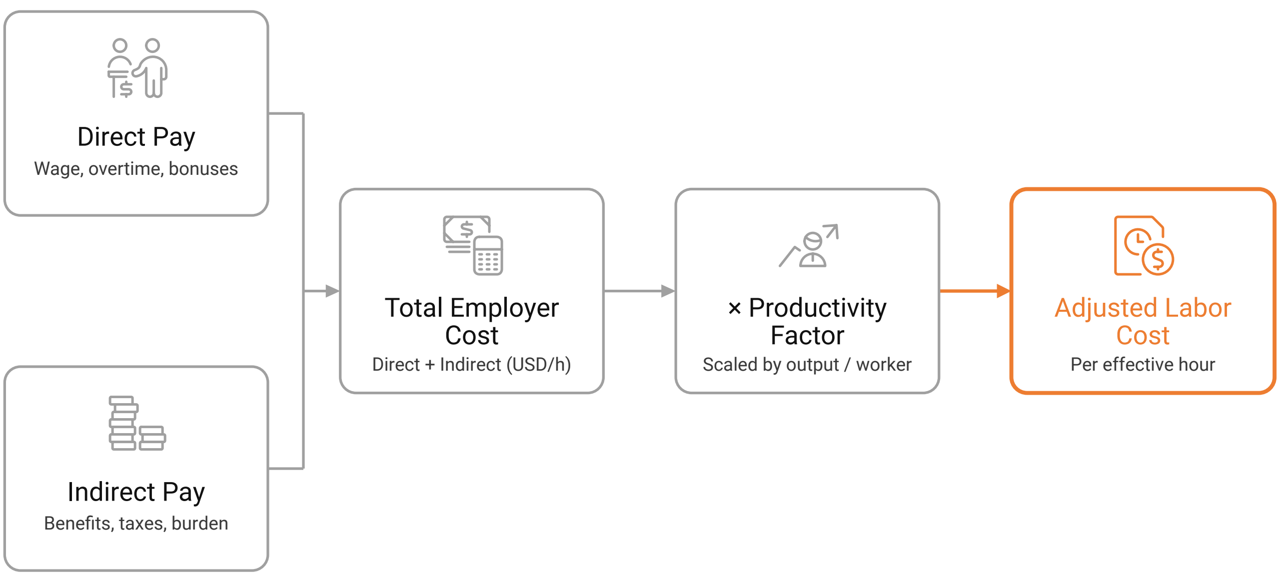

Figure 4.1 Labor cost builds from directly paid wages through benefits and other employment costs to total employer cost, then a productivity-adjusted cost.

Role in the score. Labor scores are **productivity-adjusted**: the Adjusted Labor Cost per effective hour is normalized across the 33-country set into a 0–100 score where 50 is the global mean, with cost inverted so that a lower adjusted cost produces a higher score. The productivity factor anchors at **1.0×** for the most productive position and rises **linearly** to **3×** at the lower bound of the distribution. Countries at or above the reference anchor all receive the same 1.0× factor. The anchor points are calibrated against a broader set of economies than the 33 scored countries, providing a more stable statistical foundation for the scale.

Example. A country with high headline wages but strong output per worker can still score above the global mean, because the productivity adjustment credits what each labor dollar actually yields rather than the raw wage rate alone.

Construction Labor Costs

Measures. This pillar measures the labor cost of the trades that build a process plant — distinct from the workforce that later operates it.

Data. The same labor reference dataset isolates a Construction Labor role alongside the manufacturing roles. The cost build-up is identical in structure:

* **Direct Pay** — base wages, overtime, bonuses, and allowances.
* **Indirect Pay** — voluntary and customary benefits together with legally required contributions and employer taxes.
* **Total Employer Cost per hour** — Direct Pay plus Indirect Pay, scoped to the construction sector.
* **Adjusted Construction Labor Cost per effective hour** — Total Employer Cost scaled by a **construction productivity factor** anchored at 1.0 for the United States.

Source categories are national statistics bureaus and government labor agencies covering construction-trade compensation; figures are stated on a common monetary basis.

Role in the score. Construction labor cost is productivity-adjusted and normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean, with cost inverted so that a country where skilled construction trades are cheaper for the output delivered scores higher. The productivity factor anchors at **1.0×** for the most productive position and rises linearly to **5×** at the lower bound of the distribution — a wider range than the 3× ceiling for manufacturing, because output-per-worker variation is greater in construction trades.

Example. Two countries with similar manufacturing labor scores can diverge on construction labor when one has a tight, high-cost market for construction trades, which lowers its Construction Labor Costs score without touching the manufacturing figure.

Capital & Construction Costs

Measures. This pillar measures the cost of building and equipping a process plant — the equipment, materials, and construction execution required to bring a plant on line.

Data. The pillar is grounded in the construction-cost machinery the report maintains for its reference datasets. The build-up runs in two stages.

Two primary components form the base construction-cost factor:

* **Material** — domestically sourced inputs such as steel and cement plus imported machinery, each priced at ex-factory level adjusted for freight, insurance, and import duties.
* **Construction Labor** — local construction crews and the premium for importing skilled trades where local execution capacity is thin.

Two near-neutral adjustments are then applied to that base factor:

* **Business-environment adjustment** — regulatory quality, government effectiveness, and corruption exposure.
* **Logistics-and-infrastructure adjustment** — the cost and reliability of moving goods within the country.

The result is an **Adjusted Construction Cost Factor** expressed with the United States at 1.00. Source categories span national statistics bureaus, governmental agencies, and market exchanges for the underlying material and price series.

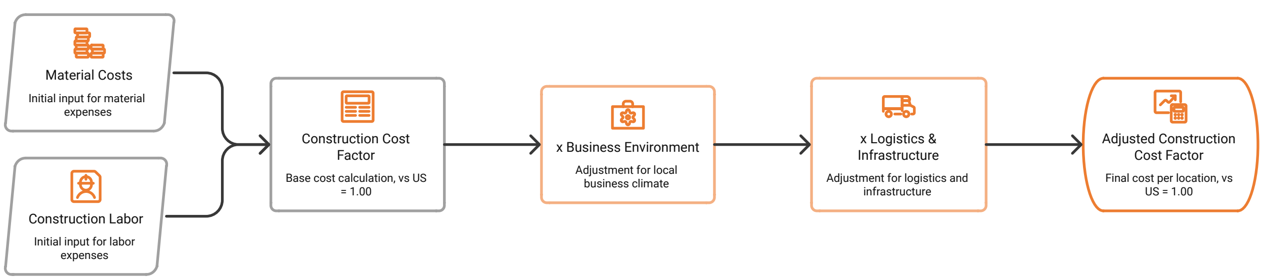

Figure 4.2 Capital & Construction cost builds from Material and Construction Labor into a base factor, then two near-neutral adjustments — business environment and logistics — yield the Adjusted Construction Cost Factor (US = 1.00).

Role in the score. The Adjusted Construction Cost Factor is normalized across the 33-country set into a 0–100 score scaled so that 50 is the global mean, with cost inverted so that a cheaper place to build scores higher.

Example. A country whose steel and equipment inputs are inexpensive but whose construction execution is hampered by weak local infrastructure may land near the global mean, the cheap materials offsetting the costlier execution.

Energy & Utilities Costs

Measures. This pillar measures the cost of the energy and industrial utilities a plant consumes day to day — the recurring inputs that drive operating expense for most processes, including electricity, fuels, steam, water, and industrial gases.

Data. Two components feed the pillar:

* **Average Energy Price** (the larger component) — a consumption-weighted blend of electricity, natural gas, and coal standardized to USD per MMBtu. Energy pricing carries greater weight in the pillar.
* **Cost of ten core process utilities** — covering four utility families: steam; water utilities (cooling water, chilled water, demineralized water, and process water); atmospheric gases (compressed air, nitrogen, and oxygen); and process gases (hydrogen and carbon monoxide). The relative weight of each utility within the ten-utility group is qualitative only.

Each utility's cost is **modeled**, not observed. Every country runs through the same cost-estimating model, driven by three country-level inputs: **labor cost**, **construction-cost inflation**, and **energy cost**. These three inputs adapt the model to local conditions, so the resulting figures are comparable across the 33-country set. For each utility, a **representative facility size** is assumed — one that carries defined purity and pressure specifications for gases and steam — and a larger facility carries a lower unit cost. Each utility can be reported on up to three cost bases: an **on-site cash cost** (a consumer self-generating the utility for its own needs), an **off-site cash cost** (a large-scale external supplier serving multiple nearby plants), and a **contract price** (what a customer pays that off-site supplier, adding depreciation, corporate overhead, and return on capital to the cash cost). Cash-cost bases exclude those capital recovery items.

Figure 4.3 The two components of the Energy & Utilities Costs pillar: Average Energy Price (the larger component) and the cost of 10 core process utilities across four families — steam, water utilities, atmospheric gases, and process gases.

Source categories include market exchanges and government energy agencies behind the energy price series.

Role in the score. The combined energy and utility cost level is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean, with cost inverted so that a country able to supply energy and utilities cheaply scores higher.

Example. A country with abundant low-cost natural gas and power can score well above the global mean on this pillar, signaling a competitive home for an energy-intensive process even before sector economics are considered.

Logistics & Infrastructure

Measures. This pillar measures the endowment and cost of moving goods within a country — the quality and cost of the transport and infrastructure available to a plant for getting inputs in and product out.

Data. Five components feed the pillar:

* **Inland Transportation** (the most heavily weighted component) — road and rail freight throughput and network density.
* **Ports & Trade Logistics** — container-port throughput and integration into global shipping lanes.
* **Logistics System Efficiency** — economic output per unit of fuel consumed, a proxy for how productively the logistics network is operated.
* **Utility Infrastructure** — the reach of electricity, water, and gas networks.
* **Infrastructure Investment** — gross fixed capital formation as a share of the economy, capturing the pace at which the stock is being renewed.

Source categories are national statistics bureaus, governmental agencies, and multilateral organizations that compile infrastructure and logistics indicators.

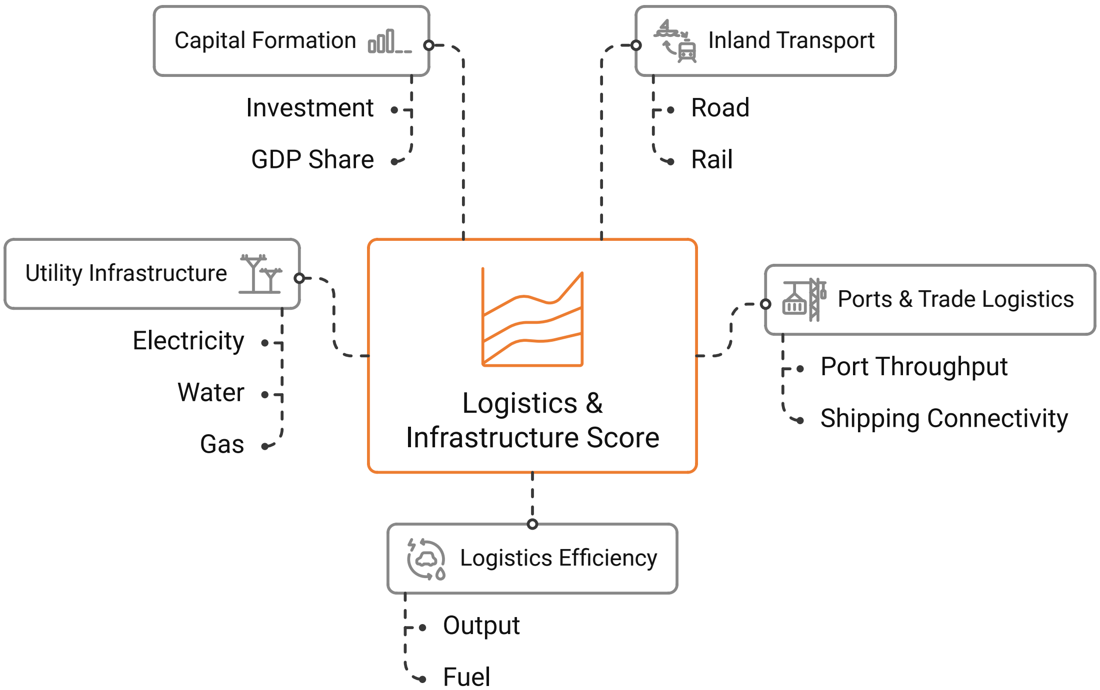

Figure 4.4 The five components of Logistics & Infrastructure: Inland Transportation (most heavily weighted), Ports & Trade Logistics, Logistics System Efficiency, Utility Infrastructure, and Infrastructure Investment.

Role in the score. The combined logistics-and-infrastructure measure is normalized across the 33-country set into a 0–100 score scaled so that 50 is the global mean. Because stronger infrastructure and lower internal transport cost both favor a manufacturer, a better-endowed, lower-cost logistics environment produces a higher score. For landlocked countries, the ports component is not applicable; its contribution is redistributed proportionally across the remaining components. Qualitative weights only; inland transportation carries the greatest weight but no percentage is disclosed.

Example. A country with dense, well-maintained road and rail and efficient ports scores high here even if its international freight rates are unremarkable, because the two movement pillars are scored separately.

Freight Costs

Measures. This pillar measures the cost of international freight to and from a country — the shipping cost on both inbound routes for feedstock and outbound routes for product.

Data. Two components feed the pillar:

* **Maritime Freight** (the greater-weight component) — ocean shipping rates across four cargo categories: containerized goods, dry bulk, liquefied gases, and bulk light liquids, measured in both import and export directions and anchored to each country's main traded commodity and principal partner.
* **Inland Freight** — road and rail rates blended by each mode's actual share of freight volumes.

Ocean freight carries the greater weight in the pillar, reflecting its outsized role in commodity trade; the relative weight is qualitative only. Source categories include market exchanges that publish freight rates for inbound and outbound routes.

Role in the score. Inbound and outbound freight cost is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean, with cost inverted so that a country with cheaper international freight scores higher. Pairing this pillar with Logistics & Infrastructure separates the two distinct movement problems a plant faces — reaching the rest of the world versus moving goods at home — so a country strong on one need not be strong on the other.

Example. A coastal country on busy, competitively priced shipping lanes can score high on Freight Costs while its internal Logistics & Infrastructure score lags, reflecting cheap ocean freight but congested domestic transport.

Macroeconomic Environment

Measures. This pillar measures the broad economic conditions a plant operates within — six dimensions that together describe the stability and depth of the macroeconomic environment a long-horizon investor inherits.

Data. Six components feed the pillar:

* **Inflation** — scored on a stability band; both deflation and high inflation are penalized, reflecting that both extremes raise uncertainty for long-term capital allocation.
* **Exchange-rate volatility** against the US dollar.
* **Short-term interest rate** — the cost of working capital.
* **Sovereign credit standing** — mapped onto a numeric scale drawn from published sovereign credit ratings.
* **Corporate borrowing cost** — proxied by the investment-grade corporate bond yield.
* **Gross national savings** as a share of the economy.

Price stability and short-term financing cost carry the greatest weight in the pillar; the weights are qualitative only. Source categories are national statistics bureaus, governmental agencies, and multilateral organizations that publish macroeconomic data.

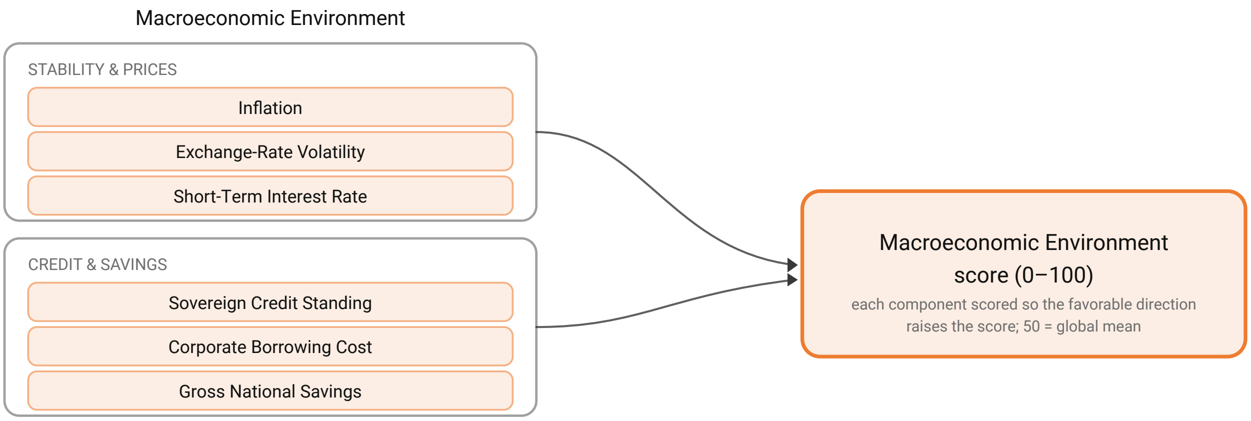

Figure 4.5 The Macroeconomic Environment pillar combines six components — inflation, exchange-rate volatility, short-term interest rate, sovereign credit standing, corporate borrowing cost, and gross national savings — each oriented so the favorable direction raises the score.

Role in the score. The six dimensions are each scored so the favorable condition raises the score: a moderate, stable inflation band scores highest; a steadier currency scores higher; lower short-term rates score higher; stronger sovereign credit scores higher; lower corporate bond yields score higher; and greater national savings depth scores higher. Together the six are normalized across the 33-country set into a 0–100 score scaled so that 50 is the global mean. Because this pillar scores conditions rather than a single price, there is no single cost inversion — the favorable direction is built into each dimension's orientation.

Example. A country with low, stable inflation and a steady currency scores above the global mean here, signaling a more predictable basis for long-horizon capital decisions, while a country with volatile prices or a swinging exchange rate scores below it.

Domestic Tax Environment

Measures. This pillar measures the fiscal burden manufacturing bears in a country — the tax and duty environment facing producers.

Data. The pillar combines three fiscal charges with import duties on machinery:

Fiscal charges:

* **Combined corporate income tax rate.**
* **Indirect tax on goods** — VAT or GST-equivalent.
* **Employer payroll contributions.**

Plus import duties on **capital machinery** — the tariffs levied on imported process equipment.

Source categories are governmental agencies and multilateral organizations that compile comparable tax statistics.

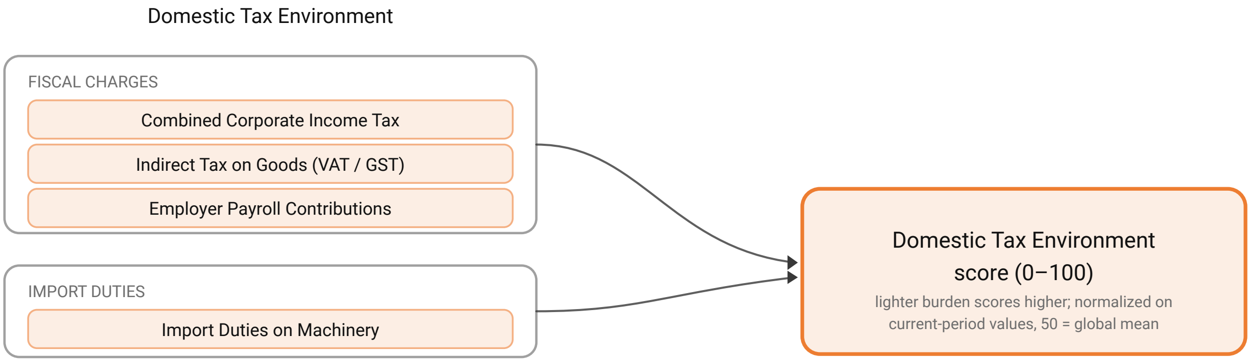

Figure 4.6 The Domestic Tax Environment pillar combines three fiscal charges and import duties on machinery, normalized on current-period values so a lighter burden scores higher.

Role in the score. The combined fiscal burden is normalized against the **current period's value** (not a trailing twelve-month average) and then scaled across the 33-country universe into a 0–100 score where 50 is the global mean, with burden inverted so that a lighter tax and duty load produces a higher score. Current-period normalization is used here — as with the Tariff Protection & Market Access pillar in Chapter 5 — because tax rates and tariff schedules change by policy decision rather than varying continuously.

Example. A country that pairs a moderate corporate rate with low import duties on machinery scores above the global mean on this pillar, even where its labor or energy costs are only average.

The eight Base Pillars score a country's shared operating environment once, for every commodity family; the economics that vary by commodity sector are scored separately in the six Industry-Specific pillars of Chapter 5.

CHAPTER 5

Industry-Specific Pillars

SUMMARY

- Six Industry-Specific pillars score sector economics — commodity prices, feedstock margins, production scale, trade integration, tariff exposure, and home-market demand — separately for each of the 7 commodity families.
- All six are equal-weight averaged into the Industry Composite for each family, which combines with the shared Base Pillars Score to produce that family's Industry Score.
- Because each family is scored independently, a country can lead one industry and lag another on exactly these dimensions.

The six **Industry-Specific pillars** capture the economics that change from one commodity industry to the next. These six are scored separately for each of the 7 commodity families: Olefins & Derivatives, Aromatics & Derivatives, Alcohols & Organic Acids, Polymers, Fertilizers, Inorganic Chemicals, and Metals. The same country therefore carries a different Industry-Specific profile for polymers than for fertilizers, because input prices, conversion margins, production scale, trade exposure, and home-market demand all differ by industry. Chapter 3 placed these six inside the wider 14-pillar architecture; this chapter takes each in turn.

The seven commodity families

The 7 commodity families define the industry scope. Each family groups related commodities that share feedstock chains, production technologies, and market dynamics; the representative commodities listed below anchor the data collection and scoring for that family.

Table 5.1 The seven commodity families and their representative commodities.

COMMODITY FAMILY	REPRESENTATIVE COMMODITIES
Olefins & Derivatives	Ethylene, propylene, butadiene
Aromatics & Derivatives	Benzene, toluene, paraxylene, styrene
Alcohols & Organic Acids	Methanol, ethylene glycol, acetic acid
Polymers	HDPE, polypropylene, PVC, polystyrene
Fertilizers	Ammonia, urea, DAP, potassium chloride
Inorganic Chemicals	Caustic soda, soda ash, titanium dioxide, aluminum oxide
Metals	Carbon steel, stainless steel, aluminum

For each family, commodity scores are normalized to 0–100 against the 33-country universe and equal-weight averaged within the family; the seven family scores are then equal-weight averaged into the pillar score. All scoring uses a twelve-month average of the underlying indicators, with one exception: the Tariff Protection & Market Access pillar uses current-month values because tariff schedules change by policy step.

Every Industry-Specific pillar is described with the same four beats: what it **measures**, the **data** categories that feed it, its **role in the score**, and a short **example**. Two conventions hold across all six. First, each pillar resolves to a single 0–100 score scaled so that 50 is the global mean across the 33-country universe — above 50 is stronger than average, below 50 weaker. Second, the direction matches the Base Pillars: a higher score always signals a more competitive position, so where a pillar rests on a cost or price the scoring inverts the underlying value, and where it rests on scale or openness a larger or more connected position scores higher. Neither the exact normalization transform nor any below-pillar weight is disclosed.

Because all six are evaluated per family, a country can lead one industry and lag in another on exactly these dimensions — abundant low-cost feedstock for one chemical lifts that industry's margins while a thin home market holds another back. The six together form the Industry Composite for each family, which combines with the shared Base Pillars Score to produce that family's Industry Score; Chapter 6 covers how those numbers are read.

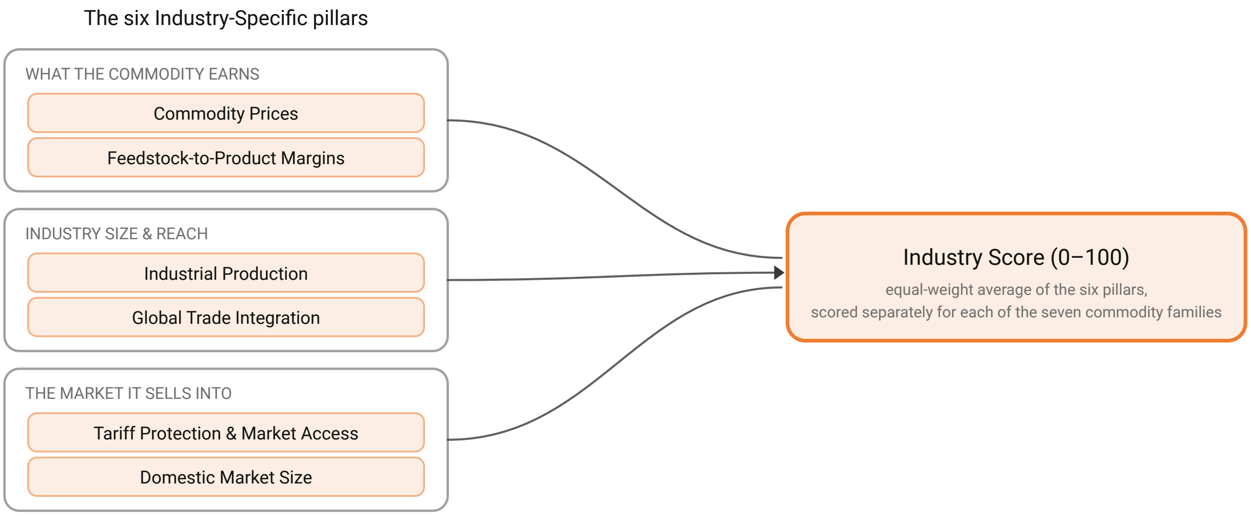

Figure 5.1 The six Industry-Specific pillars group into three themes and combine at equal weight into the Industry Score, scored separately for each of the seven commodity families.

Commodity Prices

Measures. This pillar measures the price level of a commodity in a country relative to its peers — how the local selling price of the family's products compares across the 33-country set, a direct read on relative cost competitiveness. Because the covered commodities are highly commoditized, their local pricing is a sensitive barometer of a country's underlying manufacturing economics.

Data. Prices are assessed as domestic landed prices, following a tiered determination logic:

* Where a direct domestic market quote exists, it is used.
* For net-exporting countries, the export price on a free-on-board basis serves as the domestic reference.
* For net-importing countries, the cost-insurance-freight import price adjusted for duties and inland handling gives the landed cost.

Regional price benchmarks bridge any gaps in direct observations. Settlement prices and spreads from commodity exchanges, alongside prices of commodities as sold and published by the producers, supply the underlying observations. Government-reported trade statistics anchor these price observations to each country. All figures are aligned to a common monetary basis so countries stay comparable.

Role in the score. The relative price level is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean. As a price-based pillar the scoring inverts the underlying value, so a country where the commodity is cheaper to obtain scores higher — "better for a manufacturer" points upward, the same convention used for the Base Pillars' cost dimensions.

Example. Two countries producing the same polymer can diverge sharply on this pillar when one sits in a region of structurally low local prices and the other in a high-price market, lifting the first country's Commodity Prices score above the global mean while the second falls below it.

Feedstock-to-Product Margins

Measures. This pillar measures the spread between feedstock cost and product price for the country's dominant production route — the margin earned by converting a raw input into the finished commodity, where the route is the main process or technology used in that country. It reflects the production economics that determine whether making the commodity locally is profitable.

Data. The pillar combines feedstock and product price observations for the prevailing production technology in each country: quoted prices and spreads from commodity exchanges and producers' published prices for both the input and the output, with the relevant route identified per country (for example, naphtha- versus ethane-based cracking for olefins). Technology licensors contribute published consumption figures that link a unit of product to its feedstock requirement. The series are stated on a common monetary basis so margins are comparable across the set.

Role in the score. The conversion margin for the country's main route is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean. Here a wider margin is the favorable condition, so a country whose dominant route earns a larger feedstock-to-product spread scores higher, with no cost inversion needed — the pillar already measures a margin in the competitive direction.

Example. A country with access to cheap ethane feeding an ethane-based route can show a far wider olefins margin than a peer reliant on costlier naphtha, scoring well above the global mean on this pillar even where the two are similar on most Base Pillars.

Industrial Production

Measures. This pillar measures the scale and output of the specific commodity industry within the country — how large and established local production capacity is for that family. Greater industrial depth signals an entrenched, efficient sector and contributes to a higher score.

Data. The pillar draws on production-volume and capacity indicators for the family in each country: government-reported trade and output statistics from national statistics bureaus and governmental agencies, supplemented by international multilateral organizations for cross-country context, and by technology licensors' published plant capacities verified before use. Volumes are placed on a common basis so the scale of one country's industry can be compared with another's.

Role in the score. The scale of the industry is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean. This pillar rests on scale rather than a cost, so it is not inverted:

a larger, more established production base produces a higher score, reflecting the depth and maturity of the sector in that country.

Example. A country hosting several large, long-running fertilizer complexes scores high on Industrial Production for that family, while a country with only a single small plant scores below the global mean — even if its costs are otherwise attractive.

Global Trade Integration

Measures. This pillar measures how connected a country's industry is to international trade flows for the commodity family — the export and import activity and overall openness that tie local production into global value chains. A more integrated industry is better positioned to source inputs and place output across borders.

Data. The pillar draws on trade-flow data for the family in each country across five dimensions:

* **Export value** — the scale of outward flows.
* **Import value** — the scale of inward flows.
* **Trade openness** — combined flows relative to domestic output.
* **Export-partner diversification** — how widely export partners are spread.
* **Import-partner diversification** — how widely import partners are spread.

The foundation is government-reported trade statistics from national statistics bureaus and governmental agencies, with international multilateral organizations supplying cross-country openness indicators. The flows are expressed on a common basis so connectedness is comparable across the 33 countries.

Role in the score. The degree of trade integration across the five dimensions is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean. Greater scale and broader diversification both favor a manufacturer, so the pillar is not inverted: a country whose industry trades more actively and more widely scores higher.

Example. A coastal country whose aromatics sector both imports feedstock and exports product at scale scores high on Global Trade Integration, while a largely inward-facing producer serving only its home market scores below the global mean on the same family.

Tariff Protection & Market Access

Measures. This pillar measures the tariff levels and trade barriers affecting the industry — how easily the family's goods cross the country's borders, gauging the ease of market access for both inputs and output. Export market access carries greater weight in this pillar than domestic import protection. For the export-exposed commodity industries the report covers, the ability to reach global markets drives long-run plant economics more than tariff shelter at home.

Data. The pillar draws on published tariff schedules and trade-barrier parameters for the family in each country: fiscal and regulatory data from governmental agencies, with international multilateral organizations supplying comparable cross-country tariff and trade-policy references. The parameters are placed on a common basis so the barrier facing one country's industry can be compared with another's. Because tariff schedules change by policy decision rather than varying continuously, this pillar normalizes against the current period's value rather than a trailing twelve-month average.

Role in the score. The tariff and barrier level is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean, with the barrier inverted so that lighter tariff protection and easier market access produce a higher score. As with the cost-based pillars, the more favorable condition for a manufacturer — freer trade — points upward.

Example. A country inside a large free-trade bloc, where the family's goods move across borders with little tariff friction, scores above the global mean on this pillar, while a country shielding the same industry behind high import duties scores below it.

Domestic Market Size

Measures. This pillar measures the size of in-country demand for the commodity — the scale of the local market the family's producers can sell into. A larger domestic market supports local production and consumption and contributes to a higher score. This is the dimension Chapter 4 set aside from the shared Base Pillars: in-country demand is a property of each commodity sector, so it lives here in the Industry-Specific layer.

Data. The pillar draws on consumption and demand indicators for the family in each country: government-reported trade and consumption statistics from national statistics bureaus and governmental agencies, supplemented by international multilateral organizations for cross-country demand context. Demand is expressed on a common basis so one country's market can be sized against another's.

Role in the score. The scale of in-country demand is normalized across the 33-country universe into a 0–100 score scaled so that 50 is the global mean. This pillar rests on market scale rather than a cost, so it is not inverted: a larger home market for the commodity produces a higher score, reflecting the demand a local producer can serve close to home.

Example. A populous country with heavy domestic consumption of a polymer scores high on Domestic Market Size for that family, giving local producers a sizable home market, while a small economy with thin local demand scores below the global mean even when its production costs are competitive.

The six Industry-Specific pillars are scored separately for each of the 7 commodity families and average at equal weight into that family's Industry Composite. Combined with the shared Base Pillars Score, each composite yields the family's Industry Score; the average of the seven composites is the Industry-Specific Pillars Score, which pairs with the Base Pillars Score to form the Overall Score. How these scores are read and ranked is covered in Chapter 6.

CHAPTER 6

Reading Scores and Rankings

SUMMARY

- A score places a country on a 0–100 scale where 50 is the global mean of the fixed 33-country set.
- Score, rank, and ranking are distinct: a value on the scale, a position, and the ordered list of positions.
- Seven industry rankings sit beneath one overall ranking, and a change in standing is always relative to the set.

Chapters 3 through 5 built the 14-pillar architecture and the aggregation pyramid that carries observed data up to a country's composite scores. This chapter returns to the apex of that pyramid and explains how the finished numbers are read — what a score on the 0–100 scale means, how a score differs from a rank and a ranking, how the seven industry rankings relate to the single overall ranking, and what a shift in standing does and does not signal. The structure itself is Chapter 3′s; this chapter is about interpretation.

Reading a 0–100 score

The bottom line: a score places a country on a 0–100 scale where **50 is the global mean** — the average across the fixed 33-country benchmarking universe. A score above 50 marks a position stronger than that average; a score below 50, a weaker one. The ends of the scale mark the most and least competitive positions in the set. Every pillar resolves to a score on this scale, and so does every composite above it, so the same 0–100 reading applies whether the number describes one driver or a country overall.

Because the scale is normalized across the 33-country universe, a score is always a **relative** position, not an absolute measurement. It answers "more or less competitive than the other 32 countries," not "competitive against a fixed external threshold." The benchmark anchors that define the distribution — the 10th percentile (P10) and the 90th percentile (P90) — likewise describe where a country sits within the spread of the set, not a fixed verdict of good versus bad.

The scale is anchored by the distribution of the 33-country universe on the same twelve-month average that feeds each country's score — so the benchmarking population and the scoring input are always aligned. The underlying normalization uses a fixed slope in the normal case; when an extreme outlier would push a score outside the 0–100 range, the slope is automatically reduced so the worst-placed country lands exactly at the boundary rather than breaking the scale.

Within each pillar, component sub-scores are placed on the 0–100 scale individually before being combined. The pillar score is then the weighted average of those already-normalized sub-scores, with no further rescaling applied at any level of the hierarchy — the 0–100 property holds all the way from individual components up through every composite, because the inputs are already normalized rather than because any clamping is applied after the fact.

Why a lower cost produces a higher score

A higher score always signals stronger competitiveness, whatever the pillar measures. That single convention holds the framework together: across every pillar, “better for a manufacturer” points upward. For the cost-based pillars — energy cost, labor cost, and the like — the scoring **inverts** the underlying value, so a lower cost yields a higher score. A country with cheaper industrial power therefore scores high on Energy & Utilities Costs, even though the raw figure behind it is small. Pillars that rest on scale or openness rather than a cost are not inverted: there, a larger or more connected position is already the favorable one and scores higher directly.

Productivity-adjusted labor scores

Labor scores carry one further adjustment. Labor cost is weighed against **output per worker** rather than taken in isolation, so a higher-wage country that is also more productive is not penalized for the wage alone. The productivity factor is expressed as a ratio with the most efficient case set to **1.0×**, and the productivity-adjusted cost is then normalized on the same 0–100 scale as every other pillar.

> **Example:** Two countries can post similar manufacturing wage rates, yet the one whose workforce produces more per worker carries a lower productivity-adjusted labor cost — and therefore a higher labor score — than the one whose output per worker is lower.

Score, rank, and ranking

Three outputs sit close together and are easy to conflate; separating them is the key to reading the report.

* A **score** is the 0–100 value on a single dimension or composite, with 50 the global mean — a measure of *how far* a country stands from the average.
* A **rank** is a country’s ordinal position from **1 to 33**, where rank 1 is the most competitive and rank 33 the least — a measure of *where in line* it stands.
* A **ranking** is the full ordered list of all 33 countries on a given basis — the table a rank is read from.

A score and a rank can move independently: a country can hold its rank while its score drifts as the underlying data shifts, or change rank without a large change in its own score because a neighbor’s score moved. Reading both together is what distinguishes a genuine change in conditions from a reshuffling among close peers.

To make a ranking quick to read, the 33 positions are divided into **five tier bands — Top, Upper, Mid, Lower, and Bottom** — running from the most competitive countries down to the least. The bands give an immediate read of whether a country sits among the leaders, the upper-middle, the middle, the lower-middle, or the laggards of the set.

Table 6.1 The five tier bands across the 33-country ranking.

TIER BAND	POSITION IN THE RANKING
Top	The most competitive countries
Upper	Above the middle of the set
Mid	The middle of the set
Lower	Below the middle of the set
Bottom	The least competitive countries

Industry rankings and the overall ranking

The report publishes **seven industry rankings — one per commodity family — and one overall ranking**. An **industry ranking** measures a country's competitiveness for a single commodity family, such as Polymers, by combining the eight Base Pillars, shared across all manufacturing, with the six Industry-Specific pillars evaluated for that family. The **overall ranking** is the consolidated cross-industry view: it spans all seven families, summarizing a country's competitiveness across the whole set rather than within one sector.

A country can therefore lead one industry and lag in another. The Base Pillars are shared across every family, but the Industry-Specific pillars — commodity prices, margins, production, trade, tariffs, and market size — differ by family. A country with abundant low-cost feedstock for one chemical can rank near the top of that industry yet sit well down another family's ranking where its home market is small or its trade position weak. Those family-level differences are what drive the divergence across the seven industry rankings.

How the rankings are assembled

Each ranking is read off the composite scores built in Chapter 3's pyramid, all combined at equal weight from the pillar level upward. For each commodity family, a country's Base Pillars Score and that family's Industry Composite combine into an **Industry Score**, and the seven Industry Scores produce the seven industry rankings. Those per-family results then consolidate into the **Overall Score** — the Base Pillars Score paired with the Industry-Specific Pillars Score, the average of the seven Industry Composites — from which the single overall ranking is read. The layered structure is stated openly; the exact weighting applied to synthesize the levels below the pillar is not disclosed. Every position, industry and overall alike, is a rank within the fixed set of 33 countries.

Cross-industry pillar scores

Each Industry-Specific pillar also produces a cross-industry view: its equal-average across the seven commodity families, computed separately from the main four-stage aggregation chain and available as a standalone pillar-level score. This cross-industry score answers a different question than a family-level result — not “how competitive is a country for Polymers on this dimension?” but “how does the country perform on this dimension when no single family is singled out?” It is a side rollup rather than a step in the chain that leads to the Overall Score, so it does not affect any country’s ranking.

Reading a score in context

A score is a relative position within the 33-country set, and reading it as anything more invites error. It is **not an absolute investment threshold**, **not a recommendation**, and **not a forecast of profitability** — it states where a country stands among its peers on a competitiveness dimension, nothing further.

Scores also shift over time. A country’s score can change as its own underlying data is updated or as the peer set’s data moves, since the scale is normalized across the whole universe. Holding the **33-country universe fixed** is what makes that movement meaningful: because the comparison group does not change from edition to edition, a change in a country’s standing reflects a change in underlying conditions rather than a change in who it is being measured against.

> **NOTE**
>
> Scores, ranks, and tier bands are relative positions within the fixed 33-country universe — not absolute thresholds, recommendations, or profitability forecasts. Structure, the 0–100 scale, and the combination logic from the pillar level upward are stated openly; the parameters that synthesize the levels below the pillar are not disclosed.

CHAPTER 7

Reference Datasets

SUMMARY

- Four reference datasets report absolute costs and index levels, outside the 0–100 score framework.
- They span Plant Construction Cost Indexes, Plant Location Factors, Labor Costs and Productivity, and Industrial Utility Costs across the 33 countries.
- Built largely from modeled cost inputs, they stay current and several carry six-month forecasts.

The competitiveness scores answer a strategic, relative question — how competitive a country is against its peers. The reference datasets answer a tactical, absolute question — what something actually costs. They are the absolute-value figures published alongside the scores: real dollar costs and index levels, sitting outside the 0–100 framework and usable directly in a reader's own cost work.

An **absolute value** is an actual figure — a dollar cost or an index level — rather than a normalized score on the 0–100 scale used to rank competitiveness (Chapters 3 through 6). The two views are complementary: the scores show where a location stands, while the reference datasets supply the underlying magnitudes that cost estimating, escalation, and modeling require. A reference dataset is what gets used whenever concrete cost work is needed — escalating a capital estimate over time, adjusting an estimate from one country to another, or pricing utilities for a manufacturing-cost model.

The report publishes four reference datasets, detailed in turn below: Plant Construction Cost Indexes, Plant Location Factors, Labor Costs and Productivity, and Industrial Utility Costs. Together they span the principal cost components of commodity manufacturing across the 33 covered countries. Two properties hold across all four. They are built largely from modeled cost inputs rather than from lagged official trade statistics, so they stay current rather than trailing the reporting period, and selected series include six-month forecasts — keeping the figures usable for forward-looking work. And each is documented on its own with its sources, coverage, and units. The four-month analytical lag carried by the scored analysis does not apply to them; Chapter 8 covers that timing.

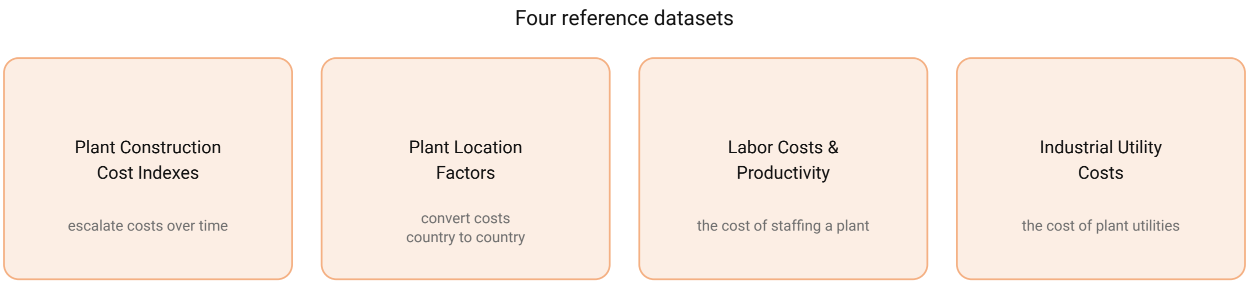

Figure 7.1 Four reference datasets support cost work — construction cost indexes to escalate over time, location factors to convert across countries, labor costs and productivity, and industrial utility costs.

Plant Construction Cost Indexes

The bottom line: the Plant Construction Cost Indexes (the IC Indexes) are monthly composite indicators that reflect the relative cost of building a process plant in a given country at a point in time, and their main use is escalating a capital cost estimate from one period to another.

A **composite indicator** is a single index that combines several distinct cost components into one figure. By condensing labor, materials, logistics, and broader economic conditions into one value, the indexes make construction-cost levels comparable across countries and over time.

Each index is normalized so that its value for **January 2000 is set to 100**, with every other monthly value expressed relative to that reference point. The fixed base year provides a stable reference for measuring change, and each country's index is expressed in that country's local currency.

The index escalates a capital cost estimate over time, adjusting a figure developed in one period onto a current basis: the ratio between the index value at the target date and the index value at the estimate's original date gives the escalation factor applied to the cost. The same series also supports comparison of capital-cost trends across periods and, because of its country coverage, across locations. Each index additionally extends about six months ahead with a forecast, labeled Forecast (F), enabling forward-looking budget adjustments before final investment decisions. This construction-cost machinery is the same one that grounds the Capital & Construction Costs Base Pillar in Chapter 4.

For the United States, the index additionally resolves the materials component into seven equipment sub-categories — heat exchangers and tanks, process machinery, pipes, valves and fittings, process instruments, pumps and compressors, electrical equipment, and structural materials — providing a more detailed decomposition of construction-cost movements. Country editions outside the United States publish the headline index alone.

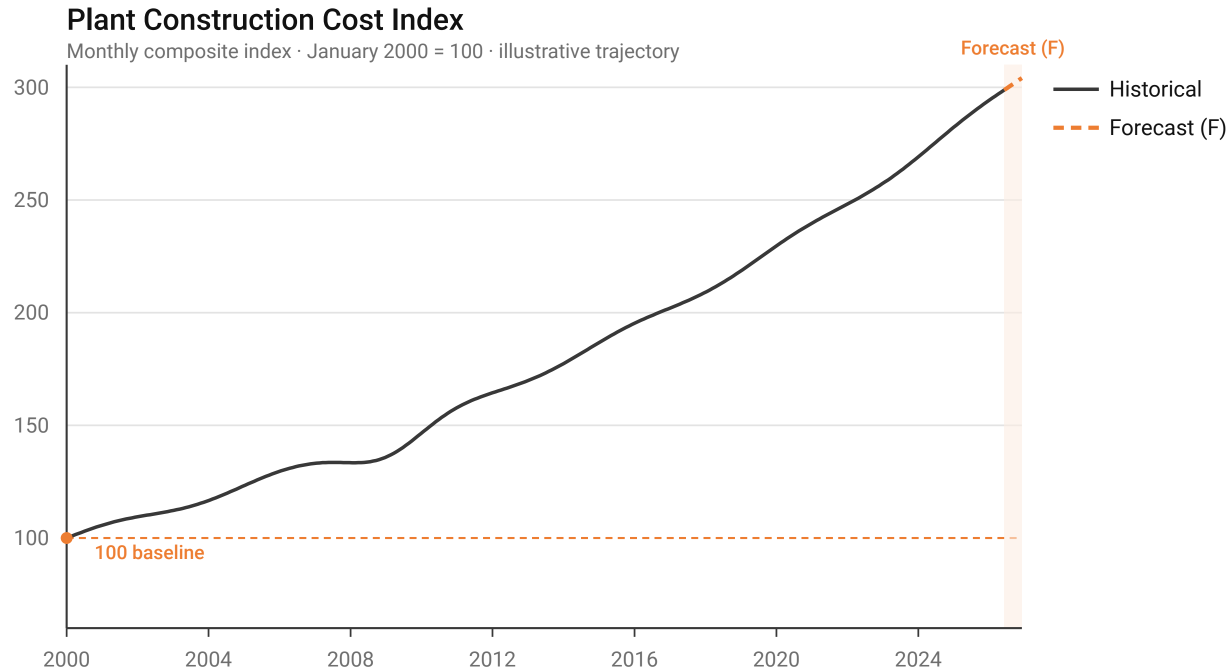

Figure 7.2 An illustrative IC Index trajectory: a monthly composite climbing from the January 2000 = 100 baseline, with the final six-month segment drawn as a dashed Forecast (F) tail, distinct from the historical series.

Plant Location Factors

The bottom line: a Plant Location Factor is a dimensionless multiplier that converts the capital cost of a reference plant from one country to another, answering "what would this plant actually cost to build there?" — entirely outside the scoring framework, with no normalization and no ranking.

Each factor expresses the relative cost of building the same plant in different countries, capturing construction-cost differences across labor, materials, logistics, and the local business environment. The report country is the reference, set to **1.00**, the baseline against which all other locations are measured. Anchoring at 1.00 keeps every other factor directly interpretable as a percentage difference: a factor above 1.00 means the report country is more expensive to build in than that country, and a factor below 1.00 means it is cheaper. Factors are published for all 33 countries.

Each factor is built from four parameter groups, weighted and combined to capture how construction cost differs across countries; the weighting is proprietary, the groups disclosed:

Table 7.1 The four parameter groups feeding each Plant Location Factor.

PARAMETER GROUP	WHAT IT CAPTURES
Labor	Wage rates, directly paid benefits, and other employment expenditures, adjusted for local labor productivity.
Materials	The cost of construction inputs, primarily steel and related materials.
Logistics	Infrastructure-related construction execution costs.
Business environment	Inflation, exchange rates, and producer price indexes.

Applying a factor is a single multiplication: a cost known for one of the other countries, multiplied by that country's location factor, gives the equivalent cost on the report country's basis. A plant costing 100 million USD in a country whose factor is 1.50 therefore implies roughly 150 million USD for the same plant in the report country — a direct cost conversion, with no scoring involved.

Labor Costs and Productivity

The bottom line: this dataset publishes the fully-loaded Total Employer Cost for four labor roles across the 33 countries, alongside two productivity factors that capture how much output a worker delivers.

The four roles covered are **Chemical Plant Operator**, **Chemical Plant Supervisor**, **Manufacturing Labor**, and **Construction Labor**. Each is published as a single **Total Employer Cost** — the complete cost of employing that worker for an hour of labor, combining directly paid wages, benefits, and the mandatory contributions and overhead an employer carries beyond take-home pay. Reporting one fully-loaded figure per role keeps countries comparable: a low headline wage paired with heavy mandatory on-costs does not read as cheaper than it is.

The four roles span the workforce a process plant depends on. The Chemical Plant Operator and Chemical Plant Supervisor figures capture the cost of running a plant; Manufacturing Labor represents the broader production workforce; and Construction Labor covers the trades that build the facility. Manufacturing Labor and Construction Labor are the two roles that feed the labor pillar scores in Chapter 4, while the operator and supervisor figures support operating-cost work directly.

Figures are expressed as a monetary cost per unit of labor time — USD per hour or per year — stated in USD on a common basis to keep the 33 countries comparable. Reported as one fully-loaded figure per role, a country's full labor picture can be read role by role.

Alongside the cost figures, the dataset publishes two productivity factors — one for manufacturing labor and one for construction labor — that scale each country's labor cost by output per worker, with the most productive position set at 1.0×. The factors are what let the scoring compare countries on cost per unit of output rather than raw wages, so a country with higher wages but greater productivity is not unfairly penalized relative to a low-wage, low-output peer. These figures are the foundation for the

Manufacturing Labor Costs and Construction Labor Costs pillar scores described in Chapter 4; Chapter 6 covers how the adjusted figures translate into pillar rankings.

Industrial Utility Costs

The bottom line: this dataset provides estimated costs for ten common industrial utilities across the 33 countries, derived from consistent cost models rather than direct surveys so the figures stay comparable from one country to the next.

The ten utilities priced are those frequently consumed in chemical and petrochemical manufacturing:

Table 7.2 The ten industrial utilities priced in the dataset.

NO.	UTILITY	NO.	UTILITY
1	Compressed air	6	Steam
2	Process water	7	Oxygen
3	Demineralized water	8	Nitrogen
4	Cooling water	9	Hydrogen
5	Chilled water	10	Carbon monoxide

Figures are derived from cost models rather than direct surveys: each model combines national labor cost data, construction-cost inflation, and energy cost inputs to estimate what a utility would cost in a given country. Because every country is run through the same model structure, the resulting figures are comparable across locations, and a standard assumed capacity — a fixed reference production rate per utility — is applied so the figures reflect a comparable scale of supply in every country. The dataset also extends about six months ahead with a price forecast, labeled Forecast (F), supporting forward-looking operating-cost projections.

Those country-level inputs adapt each model to local conditions, but what a model meters differs by how the utility is produced:

Table 7.3 The main cost components each utility model meters, beyond the shared country-level inputs.

UTILITY	MAIN MODELED COST COMPONENTS
Compressed air	Electricity to drive the air compressors, air-filter replacement, and maintenance.
Process and demineralized water	Raw water, treatment chemicals, resin replacement, and electricity.
Cooling water	Clarified water make-up, chemicals, and electricity for the cooling tower and pump motors.
Chilled water	Refrigerant make-up and electricity for the chillers and pump motors.
Steam	Fuel, boiler feed-water make-up, chemicals, and boiler maintenance.
Oxygen and nitrogen	Maintenance, labor, and electricity.
Hydrogen and carbon monoxide	Natural gas or coal feedstock, maintenance, labor, and electricity.

Because the unit cost of a utility falls as the producing system grows, each figure is tied to a fixed assumed capacity, defined on two bases: an on-site system sized for a single industrial consumer, and a larger off-site system sized to supply several nearby plants. The assumed capacities and key specifications are fixed design points of the dataset:

Table 7.4 Assumed unit capacities for utility generation.

UTILITY	ON-SITE	OFF-SITE	SPECIFICATION
Compressed air	360,000 Nm^3/h	—	8 bara
Process water	36 m^3/h	36,000 m^3/h	—
Demineralized water	36 m^3/h	3,600 m^3/h	—
Cooling water	1,000 m^3/h	10,000 m^3/h	—
Chilled water	800 kW	—	5 °C
Steam (HP)	36 mt/h	360 mt/h	45 barg
Steam (MP)	36 mt/h	360 mt/h	8 barg
Steam (LP)	36 mt/h	360 mt/h	3 barg
Oxygen	10,000 Nm^3/h	20,000 Nm^3/h	99.5 vol%
Nitrogen	10,000 Nm^3/h	20,000 Nm^3/h	99.7 vol%
Hydrogen	30,000 Nm^3/h	60,000 Nm^3/h	99.9 wt%
Carbon monoxide	30,000 Nm^3/h	60,000 Nm^3/h	99 wt%

Steam is priced at three pressure grades — high, medium, and low — that together count as one of the ten utilities.

> **NOTE**
>
> The capacities and specifications above describe the standardized modeling basis shared by every country — a generic overview, not a site-specific estimate. The detailed assumptions behind each individual country assessment — production route, source selection, and country-specific inputs — are documented separately, in a Specific Assessment Guide published for each covered country.

Each utility carries up to three figures, each representing a different cost basis. A **cash cost** is the recurring out-of-pocket operating cost of producing the utility, excluding capital recovery.

Table 7.5 The three cost bases on which each utility figure can be reported.

COST BASIS	MEANING
On-site cash cost	Cost to produce the utility on the plant site.
Off-site cash cost	Cost to produce the utility at a separate, off-site facility.
Contract price	Price to purchase the utility under contract from a third party.

Cash-cost bases exclude capital-recovery items — depreciation, corporate overhead, and return on capital; the contract price includes them. The distinction is useful when comparing self-supply against purchasing: cash costs benchmark the operating-cost floor, while the contract price benchmarks what the market charges. Not every utility is reported on all three bases — the applicable basis reflects how that utility is typically supplied. Gases (hydrogen, oxygen, nitrogen, carbon monoxide, and compressed air) and steam are reported on the on-site cash-cost basis; process water is reported as a contract price.

The energy cost inputs to these models flow from two Intratec data sources — Primary Commodity Prices and Energy Price References — a deliberate cross-product link that keeps utility estimates consistent with broader energy benchmarks. The resulting utility costs then feed the Energy & Utilities Costs Base Pillar (Chapter 4), where they help quantify how competitively a country can supply the energy and utility services that industrial production depends on.

NOTE

The reference datasets carry absolute figures — real costs and index levels — outside the 0–100 scoring. The fixed design points stated here (January 2000 = 100, the reference country at 1.00, the ten utilities, the four parameter groups) are part of the published structure; the underlying cost and index values, and the weighting that combines the parameter groups, are the proprietary data the report supplies.

CHAPTER 8

Timing, Cadence, and Revisions

SUMMARY

- The report publishes monthly, each edition measuring one country against the full set of 33.
- Scored pillars carry a four-month analytical lag, and a status label marks every value's maturity.
- Values settle from estimate toward verified statistic, with a six-month forward horizon.

A competitiveness score is a snapshot of a moving target, so when a figure was measured matters as much as what it says. This chapter follows the life of a single monthly edition: the rhythm on which it publishes, the four-month analytical lag the scored pillars carry, the status label every value wears so its maturity is plain at a glance, the way a value settles from estimate toward verified statistic, and the six-month horizon over which the report looks ahead.

Publication cadence

The bottom line: the Industry Economics & Competitiveness report publishes monthly, and each edition focuses on one country measured against the full set of 33, regenerated from the latest available data so the analysis tracks shifting conditions rather than resting on an annual snapshot.

Every monthly cycle runs the same fixed sequence. Official government data — figures issued by national statistical and regulatory authorities — is collected in the last week of the month. At the start of the following month, that data is processed and assessed to produce the update, and each new edition supersedes the prior one, building a continuous, regularly spaced record over time. Because every country runs through the same monthly collection-and-processing cutoff, the figures within a single edition share a common reference moment: no country's data is more or less current than another's inside the same edition, which is what keeps cross-country comparison valid.

The four-month analytical lag

The bottom line: the analytical pillars carry a data lag of roughly **four months**, so each score reflects conditions about four months prior — the time it takes for official statistics to arrive and be validated before they enter a score.

The lag is a direct consequence of an official-data-first approach. The pillars draw on official trade and economic statistics, which national sources themselves release one to three months after the period they describe. Intratec then applies its own processing and validation before scoring. Together, the source lag and that preparation step put each published score roughly four months behind the period it describes, and the payoff is that every competitiveness score rests on verified official data rather than provisional estimates.

The reference datasets — the absolute-value figures covered in Chapter 7 — are not subject to this analytical lag. Because they are built largely from modeled cost inputs rather than from lagged trade statistics, they refresh on a more current basis and carry six-month forecasts, staying ahead of the scored pillars rather than trailing the reporting period. The result is a deliberate division of labor across one edition: the scores sit on a settled, four-month-old official footing, while the reference datasets supply current and forward-looking magnitudes for cost work.

Status labels: Final, Preliminary, and Forecast

The bottom line: every figure carries one of three status labels that signal how settled it is, so an interim estimate is always distinguishable from a confirmed official value at the point of reading.

The label is shown directly alongside each value throughout the report's tables and charts, so the maturity and basis of any number can be read without consulting separate notes. The three labels mark a clear progression in certainty:

Table 8.1 The three data status labels and what each signals.

LABEL	CODE	WHAT IT SIGNALS
Final	Fi	Verified official statistics — the settled value, not expected to change.
Preliminary	P	A model-generated estimate standing in while official data is unavailable or lagging; a revision is expected.
Forecast	F	A forward-looking model prediction for a period not yet reported by any source.

A figure carries the Preliminary label when the relevant official source data has not yet arrived — typically a source lag of one to three months — and a model estimate is used in its place. Marking it Preliminary makes clear the value is an interim estimate awaiting confirmation, rather than a settled statistic.

Revisions and status transitions

The bottom line: each figure moves through a three-stage lifecycle — **Forecast → Preliminary → Final** — and each transition replaces the prior value as better data becomes available, with every change recorded publicly so a number never shifts silently.

The progression follows the calendar. A Forecast (F) is a model prediction for a period not yet reported. Once that period becomes current but official data has not yet arrived, the figure becomes a Preliminary

(P) estimate from the same model. When verified official statistics are released, the figure becomes Final (Fi) — the model estimate replaced by the confirmed official value. That settling from estimate toward verified statistic is the normal path of almost every value in the report.

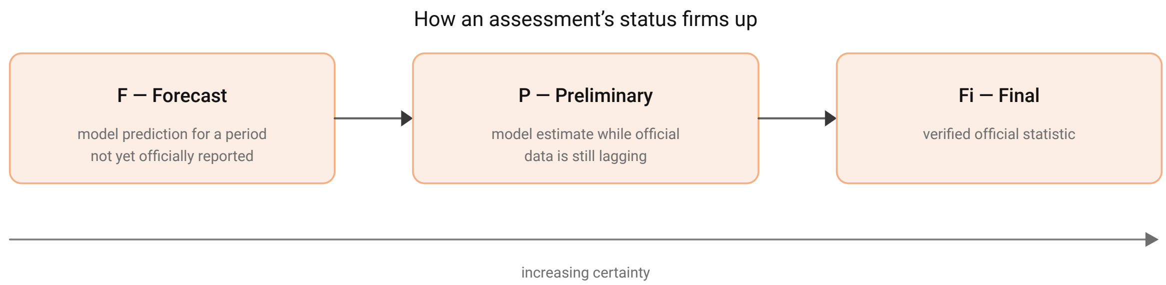

Figure 8.1 An assessment firms up through three statuses — Forecast (F), Preliminary (P), and Final (Fi) — as certainty rises from a model prediction to a verified official statistic.

Beyond that normal settling, a published number is revised for three main reasons:

* Official statistics arrive and replace earlier estimates — the standard Preliminary-to-Final transition.
* Continuous model reviews and methodology improvements introduce corrections that refine previously published values.
* A submitted assessment complaint is resolved, prompting an adjustment to the affected figure.

When methodology or data is revised, the historical series may be restated so that trends stay comparable over time. Every change — data revisions, methodology changes, assessment additions and retirements, and corrections — is documented in the monthly, publicly available Intratec Release Notes, which keep the published data auditable for any reader.

The forecast horizon

The bottom line: forecasts extend about **six months** ahead of the latest reported period, each value labeled Forecast (F), so recent trends carry forward past the last officially reported point without straying so far that the projection loses its footing.

A Forecast value is a forward-looking estimate for a period not yet reported by any official source. The six-month horizon is a deliberate balance: it keeps forward-looking values close enough to observed conditions to stay informative while holding back from the rising uncertainty that accompanies longer projections. The forward path is generated with statistical and machine-learning models trained on historical series and related economic variables — commodity prices, economic indices, and industry indicators — which learn patterns from past behavior and from variables that move together with the series being projected. The modeling approach is stated openly; the parameters that produce individual values — the exact model architecture, the specific input features and their weights, and the underlying coefficients — remain proprietary.

To signal the lower precision of estimated periods honestly, forecast figures are rounded to **two significant figures** — coarser than the three significant figures applied to historical values. The same rounding convention applies to Preliminary estimates while the underlying official data is still absent.

The methodology cadence

Publication is monthly, but the methodology itself runs on **three nested cadences**, each carrying its own maintenance tasks:

Table 8.2 The three cadences at which methodology work is carried out.

CADENCE	TASKS
Monthly	Core production: data collection, transformation, modeling, and publication.
Quarterly	Scheduled maintenance: searching for new sources, revising models, reviewing assessment descriptions.
Annual	Broadest pass: market research, an assessments-relevance check, and a full methodology review.

The monthly cadence keeps each edition current. The quarterly and annual cadences keep the methodology itself sound as markets evolve — ensuring that sources remain representative, models reflect current market dynamics, and the assessment architecture stays aligned with what the commodity industries require. When quarterly maintenance identifies a better source or a model improvement, the change is announced in the monthly Release Notes before it takes effect, and the historical series is restated as needed so that trends remain comparable. A permanently discontinued source is replaced as quickly as possible under the same official-first selection priority used for all source management; model estimates under a Preliminary or Forecast label keep the series continuous in the interim.

NOTE

A figure's status label travels with it: Final (Fi) for a verified official statistic, Preliminary (P) for an interim model estimate, Forecast (F) for a value up to six months ahead of the last reported period. The four-month analytical lag applies to the scored pillars; the reference datasets of Chapter 7 are current and carry their own six-month forecasts. Every revision is logged in the monthly public Release Notes.

CHAPTER 9

Independence, Review, and Limitations

SUMMARY

- The analysis rests on public official data and impartial, auditable methods, with no commercial ties to the markets measured.
- A clear line separates what is disclosed from what stays proprietary.
- Known limitations and uncertainty are signaled at the point of reading, and the methodology is open to review and challenge.

A methodology earns trust through two things: independence from the parties it measures, and candor about what its numbers can and cannot claim. This closing chapter states the independence commitments behind the analysis, draws the line between what is disclosed and what stays proprietary, sets out the main limitations of the results, explains how uncertainty is signaled at the point of reading, and describes how the methodology is reviewed and how a published figure can be challenged.

Independence

The bottom line: the analysis rests on public official data processed through impartial, auditable methods, with no commercial ties to the parties whose markets it measures — so no single participant has either the standing or the incentive to influence a result.

The vast majority of the data comes from public official sources — national statistics bureaus, trade and regulatory agencies, multilateral organizations, and market exchanges. Private sources and producers' data are used only with permission from the rights holders, and the analysis is not a primary source for pricing data. There are no commercial relationships with the producers, traders, and other firms active in the markets covered, and surveys are used only in very exceptional circumstances. Removing voluntary submissions from the core of the method removes the channel through which any one party could shape the figures in its favor.

Disclosure without exposing internals

The bottom line: the approach is fully public; the calibration is not. Anyone can read how a score is built — what cannot be read are the precise parameters that turn the raw inputs into a number.

The methodology approach is open and requires no login or subscription. The data sources, the 14-pillar structure, the 0–100 scoring framework with its baselines, and the data status labels are all disclosed and explained in the preceding chapters. What remains proprietary is the precise calibration — the pillar weights, the regression coefficients, and the normalization constants. This is a deliberate boundary rather than an omission: the logic of the framework is auditable, while the tuned parameters that would let the index be trivially reverse-engineered are protected.

> **NOTE**
>
> Disclosed: the data sources, the 14 pillars, the 0–100 framework (50 = the global mean), the baselines, and the Final / Preliminary / Forecast status labels. Proprietary: the pillar weights, the regression coefficients, and the normalization constants. The structure is public; the calibration is not.

Limitations

The bottom line: the results rank relative position within a fixed 33-country set — they do not certify that any one country is a viable place to invest, and they inherit the limits of the official data and the modeling assumptions behind them.

Several boundaries should stay in view when reading a score:

* Scores are **relative** to the 33-country set, not absolute investment thresholds. They place a country against its peers; they do not establish that a project there will succeed.
* Figures depend on official source data, which can lag or be revised later, and some values are modeled estimates where official data is missing.
* Reference-dataset forecasts, covering roughly six months ahead, are estimates published without confidence intervals.
* Industry-specific results depend on the dominant production route assumed for a country; a different route would shift the margin and commodity-price pillars accordingly.

Usage limitations specific to each dataset — what a given figure should and should not be used for — are documented in the Intratec Help Center rather than restated here.

Communicating uncertainty

The bottom line: every figure wears its maturity openly, so an interim estimate is never mistaken for a settled official value.

Each figure carries a data status label — Final, Preliminary, or Forecast — signaling how settled the underlying value is; the full lifecycle behind these labels is described in Chapter 8. Forecasts are shown with reduced significant figures — two rather than three — so their lower precision is visible in the number itself. And comparisons are accompanied by benchmark anchors at the 10th and 90th percentiles, which convey the spread of the distribution at the point of reading rather than leaving a single figure to stand without context.

How the methodology is reviewed

The bottom line: the methodology is reviewed continuously, with a formal review of every methodology and methodology document at least once a year, and any change is announced before it takes effect.

Between formal reviews, model reviews and methodology improvements run continuously, anchored by a quarterly model review that feeds accuracy findings back into the models. A formal review can act on three kinds of finding: changes in the raw data used for the calculations, modifications to the calculation methodology such as modeling enhancements, and clarification or simplification of the methodology documents themselves. Changes resulting from a review are announced in advance, together with the length of any transition period, and are documented in the monthly, publicly available Intratec Release Notes — so a reader can always trace what changed and when. Any new or revised information undergoes a period of shadow testing before its official release, run alongside live conditions so its behavior is confirmed first.

Challenging a published figure

The bottom line: a subscriber who doubts a figure has a formal route to question it, and a confirmed issue is corrected on the public record — so the path from challenge to correction stays traceable.

A formal assessment-complaint process is available to subscribers who wish to question any published figure. A submitted complaint is investigated; when the investigation confirms an issue, the affected figure is corrected and the change is documented in the monthly Intratec Release Notes. The correction process covers clerical and calculation errors, technical glitches, and methodology misapplication, as well as retroactive adjustments when new information arrives or a model is enhanced.

Beyond individual corrections, subscriber feedback is one of the standing inputs to methodology reviews — considered alongside changes in raw data and modeling enhancements, including the clarification and simplification of the methodology documents. Every resulting change is logged in the monthly public Release Notes, keeping the published data auditable for any reader.