

JULY 2026

Report Methodology

Primary Commodity Prices

220+

COMMODITIES

7

INDUSTRIES

4

REGIONS

CHAPTER 1

Methodology Overview

SUMMARY

- Coverage spans seven fixed industries and 220+ commodities worldwide, framed as industries → commodities → locations rather than around a fixed set of countries.
- A subscription is purchased by industry: subscribing to one grants every commodity assessment within it, across all the countries and world regions covered for those commodities.
- Three design principles — objective measurement, standardized assessment, and actionable synthesis — govern every price value the program publishes.
- A shared foundation of independent sourcing, common conventions, and two-layer validation underlies every assessment; the full methodology is public and every change is documented in monthly Release Notes.

Primary Commodity Prices is an annual subscription program that delivers monthly, independent price assessments for commodities worldwide. A new set of assessments is released each month, following the published release schedule, so the reported values reflect the latest available data. The program serves traders, procurement and cost teams, analysts, and others who need reliable, independent price benchmarks. The purpose of this document is to make every step from raw inputs to a finished assessment transparent: where the data comes from, how it is validated, how each price is modeled, and how its cadence and revisions are managed.

The full methodology is public and auditable. It requires no login or subscription, and any change to a method or a figure is documented in monthly Release Notes that are themselves publicly available, so the evolution of every method can be traced over time.

What the Program Covers

Assessments rest on a discipline that collects official trade data, transforms it to common bases, and applies dedicated modeling on a monthly cycle. Three principles govern the design of the program:

- * **Objective measurement.** Assessments rest on verifiable, quantitative data — primarily international trade deals reported by official statistics — rather than subjective buyer and seller quotes that reflect expectations instead of actual transactions. This keeps the data transparent, auditable, and unbiased.
- * **Standardized assessment.** The same methodology is applied across every commodity and every month, so month-to-month comparisons remain valid.
- * **Actionable synthesis.** Monthly price assessments and short-term forecasts are organized into a consistent structure that turns raw data into decision-ready insight.

Coverage is framed from the industry down — **industries** → **commodities** → **locations**. Seven fixed industries organize the full scope, and 220+ commodities are assessed across them.

- * **Fertilizers & Gases**
- * **Olefins & Derivatives**
- * **Aromatics & Derivatives**
- * **Alcohols & Organic Acids**
- * **Polymers**
- * **Inorganic Chemicals**
- * **Metals & Mining**

For each commodity, the assessments recognized as most relevant for international trade are presented, spanning multiple countries and world regions. The exhaustive, per-assessment coverage — every commodity and every location assessed within each industry — is maintained in the assessment catalog, which is the authoritative reference for what is covered and is kept current there rather than reproduced in this document.

How Coverage Is Organized

Industry, not country, is the unit of access. A subscription is purchased by industry: subscribing to an industry grants access to every commodity assessment within it, across all the countries and world regions covered for those commodities. For each commodity, a new price value is released monthly for each available series.

NOTE

A subscription to Polymers, for instance, grants access to the assessments for every polymer commodity in that industry — HDPE, LDPE, LLDPE, Polystyrene, and the rest — across all the countries and world regions where each is assessed.

Geographically, coverage is organized around the four Intratec world regions — Americas, Europe, Asia, and Africa & Middle East. Coverage is not limited to a defined universe of countries: it extends to all countries covered for the commodities within each subscribed industry, so the geographic reach follows the trade in each commodity rather than a fixed country list.

A Shared Foundation

The program belongs to a family of Intratec solutions that rest on one common foundation and differ only in what each one measures. That shared base is what makes the price assessments independent and reproducible, regardless of the specific price type behind any single figure.

Table 1.1 The common foundation behind every assessment.

LAYER	WHAT IT PROVIDES
Core principles	Independent sourcing, no commercial ties to market participants, and impartial, auditable methods
Shared conventions	The US dollar anchors all monetary figures with a single set of reference exchange rates; reported quantities follow the International System of Units; one world-regions hierarchy of region, sub-region, and country organizes every analysis geographically
Quality philosophy	A two-layer validation runs before publication — automated processing and cross-referencing of sources, then expert human review as the final quality gate — and the methodology is reviewed continuously

Each solution nonetheless runs its own methodology and its own pipeline; the family is best understood as one set of principles and conventions with several separate methodologies built on top of them, rather than a single model applied across every market. What the common foundation guarantees is consistency of discipline — the same sourcing standards, the same reference conventions, and the same validation gate before release — so a figure from this program can be read on the same terms as any other assessment in the family, whatever the underlying market it describes.

Core Principles

Every assessment is built under the same eight commitments. They govern sourcing, independence, and disclosure alike, and they are the reason a published figure can be traced back to its inputs and the rules that produced it.

Table 1.2 The eight commitments behind every assessment.

COMMITMENT		WHAT IT MEANS
Sourcing		The program is not a primary source for pricing data; the vast majority of data comes from public official sources — statistics bureaus, trade agencies, multilateral organizations, and market exchanges. Private and producer data are used only with the rights holder's permission, and surveys occur only in very exceptional circumstances
Independence		No commercial relationships are held with commodity market participants, removing any single party's incentive to influence results
Impartial, methods	auditable	All data is processed through methods that let each figure be traced to its sources, with no commercial interest entering the calculation
Two-layer validation		Before publication, data passes automated cross-referencing of multiple sources followed by expert human review; both layers run every month
Continuous improvement		Models and methods are reviewed continuously rather than on a fixed cycle, with subscriber feedback feeding directly into corrections
Transparency of changes		Every methodology and data change is documented each month in the publicly available Release Notes
Automated collection at scale		Collection runs on automated, large-scale, resilient systems, gathering a large volume and variety of data consistently every cycle
Continuous strengthening of the source base		The set of sources grows over time, official and recognized institutions are prioritized, and discontinued sources are quickly replaced

Quality is sustained by two mechanisms working together. The monthly validation gate flags inconsistencies automatically and then submits them to market experts as the final review before release. Separately, the source base itself is continuously expanded, prioritized toward official institutions, and repaired whenever a source is discontinued, so coverage and continuity are preserved. Subscribers can also question any published figure through a formal assessment-complaint process, and resolved complaints feed the same continuous cycle of reviews.

What the Program Produces

The methodology turns raw trade data into published price assessments. The modeling path is not uniform: it depends on the price type, which in turn depends on the commodity, the available data, and the point in the supply chain being priced. Five price types cover the program, complemented by the treatment of data shortages and short-term forecasts.

- * Trade-based prices
- * Formula-based prices
- * Freight-based prices
- * Manufacturing cost-based prices
- * Compiled prices

Each of these is defined and worked through in detail in Chapter 6. The mechanics are introduced here only to orient the reader to the kinds of references the program publishes.

How This Document Is Organized

The chapters follow the journey of a number, from raw inputs to a finished, revised assessment. Cross-references throughout point to the chapter where a topic is treated in full, so any single chapter can be read on its own.

- * Review, validation, and the release cadence that gate every publication — Chapter 2.
- * Collection and integration — where the data comes from and how raw inputs are prepared and made comparable — Chapter 3.
- * The modeling approaches that turn prepared data into assessments — Chapter 4.
- * Publication of the finished assessments, with data status labels and timing — Chapter 5.
- * The detailed price-type mechanics behind each modeled value — Chapter 6.
- * The principles that keep assessments independent, and the known limits of the methods — Chapter 7.
- * The published reference literature behind the methodology — Chapter 8.




Figure 1.1 The end-to-end methodology pipeline: source data is collected, modeled into assessments, published, and presented to readers.

CHAPTER 2

Review, Validation, and Release Cadence

SUMMARY

- Three nested cadences — monthly production, quarterly maintenance, and an annual methodology review — keep every assessment current, and no change is published without advance notice and a Release Notes entry.
- A two-layer validation, automated cross-referencing across official sources followed by risk-based expert review, clears every figure before release.
- Because the assessments are price-only, a single clock governs them: preliminary (P) estimates cover the one-to-three-month source lag and are revised to Final once official data lands, and selected commodities add a roughly six-month forecast.
- Each figure carries one of three status labels — Final, Preliminary, or Forecast — marking how settled its value is.


Every figure published in the Primary Commodity Prices report passes through a discipline built to keep it both current and accountable to real outcomes. Publication is monthly, but the methodology runs on cadences longer than a month as well, each with its own tasks and its own object of review — the individual figures, the data and the models behind them, and the methodology itself. They share one commitment: no change any of them produces is published silently, and every change is recorded in the monthly, publicly available Release Notes.

Three Review Cycles at a Glance

Publishing the assessments is monthly, but the methodology advances on three nested cadences. The shorter cadence keeps each monthly edition current; the longer ones keep the methodology sound as markets evolve.

Table 2.1 The three nested cadences and the tasks each carries.

CADENCE	FOCUS	TASKS
Monthly	Production	Data collection, transformation, modeling, and publication of each edition
Quarterly	Maintenance	Searching for new sources, revising models, and reviewing assessment descriptions
Annual	Methodology	Market research, an assessments-relevance check, and a full methodology review

**Figure 2.1** The methodology's three task cadences: a monthly production cycle, quarterly maintenance reviews, and an annual methodology review.

The Annual Methodology Review

Methodologies are reviewed continuously, with a formal review of all methodologies and methodology documents at least once a year. The review is what keeps every assessment relevant over time rather than merely current. Three kinds of findings can justify a change: alterations 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, prompted by subscriber feedback.

Subscriber feedback is a standing criterion in these reviews, alongside raw-data changes and modeling improvements, so the assessments evolve with what readers report needing clarified or improved. Feedback reaches the methodology through two paths. Individually, a formal assessment-complaint process is available to any subscriber who wishes to question a published figure; a resolved complaint can prompt an adjustment to the affected value. Collectively, the accumulated feedback becomes one of the inputs weighed at review time.

Changes that result from a review are announced to subscribers in advance — before they take effect — together with the duration of the transition period. Before a new or revised method takes effect, it is run through a period of shadow testing, operated alongside the current method so its behavior can be confirmed ahead of official release. The advance notice and defined transition window are what distinguish a methodology change from a routine monthly update: subscribers learn what is changing, and when, before any figure shifts.

Continuous Data Validation and Model Maintenance

Between methodology reviews, a continuous cycle keeps the data sources and the models sound. It rests on a two-layer validation applied to all data before publication: an automated processing stage followed by expert human review. Both layers run every monthly cycle, before any figure is released, and the same process applies uniformly across every Intratec product line rather than varying dataset by dataset.

The automated stage prepares and tests the data. It cross-references multiple official sources — comparing the same data point across independent official records to confirm it agrees — and applies mathematical models to identify inconsistencies between them. Running the same consistent, validated rules each cycle reduces errors introduced through manual handling and removes bias from market participants.

After the automated stage, market experts — analysts with sector-specific knowledge — review the model outputs and serve as the final quality gate, the last check a figure passes before release. The review is risk-based rather than random: any value the automated stage flags as anomalous or inconsistent is always routed to an expert, concentrating attention where the probability of an issue is highest, while a sample of the remaining unflagged values is also reviewed each cycle as an independent check that the automated rules are not silently missing problems.

Model maintenance is the other half of the cycle. Accuracy is evaluated as part of modeling itself, with a built-in back-test supplied by the data lifecycle: every forecast and every preliminary figure is eventually confronted with the final official value once it arrives. Comparing what the models predicted with what the sources later reported shows directly how each model performs. Deviations are flagged for investigation — analysts trace the cause, from erroneous data entry to market changes the model has not yet captured — and the model is recalibrated, improved, or replaced when it no longer reflects conditions.

The insights feed back into the models, particularly during the scheduled quarterly model reviews, so reliability improves over time. This evaluation is internal — it drives recalibration rather than a published accuracy metric.

NOTE

The internal model coefficients and weightings behind an assessment are not disclosed, and neither is a numeric accuracy score. The reliability of preliminary and forecast values is instead maintained through the Preliminary-to-Final back-test — each estimate is measured against the official value that later replaces it, and the model is recalibrated when it drifts — with every resulting correction recorded in the monthly Release Notes.

When a source fails for a prolonged period — its data insufficient, inadequate, or unavailable — one of four approaches keeps the series representative rather than published at a lower standard: estimate the assessment with a model built on good prior or related data; replace the failing source with another that meets the requirements; alter the assessment's definition so it can be produced reliably from the data available; or, if the shortage persists, discontinue the series. The governing commitment is quality over continuity: a series may be initiated, changed, or retired at any time, and one that can no longer be produced reliably is changed or retired rather than published unreliably. Every such addition, change, and retirement is recorded in the monthly Release Notes.

The Monthly Release Workflow

Assessments are published on a monthly cadence, with every assessment refreshed once per calendar month. Each monthly edition supersedes the prior one, building a continuous, regularly spaced record over time.

The cycle follows a fixed sequence. Official government data — figures issued by national statistical and regulatory authorities — is collected during the last week of each month. At the start of the following month, that data is processed, validated through the two-layer approach, and assessed to produce the update. Exact release dates for each series are listed in the published release schedule. The collection and integration of these inputs is covered in Chapter 3.

Every assessment in an edition follows the same monthly collection-and-processing cutoff. Because all inputs are gathered and assessed at the same point in the cycle, the figures within a single edition reflect a common reference moment, which keeps them comparable — no series is more or less current than another within the same edition. Consistency holds even when a particular series' official statistics lag: a model-generated Preliminary estimate fills the gap for that series while the rest of the edition proceeds unaffected, and it is replaced with the Final value once the official figures are released and verified, without shifting the shared reference moment for any other series.




Figure 2.2 The monthly release workflow: Online Chart and API database update, test presentation display, release notes, and communication.

How Current Each Figure Is

Every price series runs through the most recent reported month. Because the assessments are price-only, a single clock governs them all — there is no separate, slower track behind any figure in an edition.

Official trade and market statistics reach national agencies one to three months after the period they describe. Rather than wait for that record to close, preliminary model estimates marked **(P)** stand in for the interval, and each is revised to a final value once the official data lands. An estimate is not a placeholder to be discarded: its accuracy is monitored and refreshed every month until the confirming official data arrives.


Selected commodities also carry a short-term forecast extending about six months beyond the latest reported period — the most probable scenario as of the date it is produced, regenerated every month as new data arrives. How these preliminary estimates and forecasts are built is detailed in Chapter 6. Because a forecast is an outlook rather than a measured value, it is rounded to two significant figures, where historical and preliminary figures keep three; that difference in precision is itself a signal of how firm a given number is.

Data Status Labels

Because official statistics arrive with a lag and some periods have not yet been reported by any source, not every figure in an edition is equally settled. Each figure therefore carries one of three status labels, placed directly alongside the value throughout the data tables and charts so its maturity and basis are clear at the point of reading, without consulting separate notes.

Table 2.2 The three data status labels.

LABEL	CODE	MEANING
Final	Fi	Verified official statistics; the settled value, not expected to change
Preliminary	P	A model-generated estimate used while official data is unavailable or lagging — an interim value, with a revision expected
Forecast	F	A forward-looking model prediction for periods not yet reported by any source

**Figure 2.3** 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.

A figure carries the Preliminary label as soon as the relevant official source data has not yet arrived — typically a source lag of one to three months — and a model estimate stands in its place; it signals that the value is an interim estimate rather than a confirmed statistic, and that a revision is expected. It moves from Preliminary to Final once the official statistics for that period are released and verified: the model estimate is replaced by the confirmed official value, and the label updates from P to Fi. That transition is also what supplies the back-test described above, since each estimate is eventually measured against the official value that replaces it.

Advance Notice and Release Notes

Two communication mechanisms close every cycle. Advance notice applies to methodology changes: before a reviewed change takes effect, subscribers are told what is changing and the length of the transition period, so no change to how a figure is produced arrives unannounced. The monthly Release Notes are the standing public record — every addition, change, and retirement of a series, every correction, and every methodology change is documented there, keeping the path from review or feedback to revision traceable. The treatment of corrections and the disclaimers that accompany them is covered in Chapter 7.

CHAPTER 3

Collection and Integration

SUMMARY

- The vast majority of inputs come from public, official sources; private sources are used only with rights-holder permission and supplement rather than override official records.
- A fixed pipeline of extraction, transformation, and loading — cleaning, standardizing, and filtering every raw figure — puts data on common ground before modeling begins.
- Trade data is filtered and clustered by volume, specification, location basis, and delivery terms, so a single unrepresentative trade cannot skew an assessment.
- Modeled estimates keep coverage continuous through the one-to-three-month lag before official data arrives, and every source change is recorded in the monthly Release Notes.

Before any price is assessed, raw information has to reach the system in a form the analysis can use. Primary Commodity Prices organizes this as a pipeline with three stages — extraction of data from external sources, transformation into a consistent internal form, and loading into the system that feeds the assessments. Around that pipeline sits a fixed collection schedule and a set of publication lags that govern how fresh each figure can be. This chapter describes the path a figure travels from its public source to the point where it is ready for modeling; the modeling stage itself is covered in Chapter 4, and the validation and review cadence in Chapter 2.

Extraction: the source base

The methodology is built on public, official information. The vast majority of inputs come from public official sources, and the program is deliberately not a primary source for pricing data — an approach that minimizes subjective influence and self-reporting bias in the assessments. Collection runs on automated, large-scale, resilient systems that gather a wide volume and variety of data every cycle, so the breadth of the source base does not rest on manual effort and stays consistent from one month to the next. The collection portfolio spans six categories of source.

Table 3.1 The six categories of source feeding the analysis.

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 — from public and private trading platforms
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 of the analysis. Because they are published by national authorities and statistical bureaus, the figures are public, verifiable, and free of any single party's commercial interest, and each result can be traced back to its official source. Data from international and multilateral organizations supplements those trade statistics where broader context or technical reference points are needed. 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.

When more than one source covers the same data point, official and recognized institutions take priority for final figures. Accuracy is favored over speed: a slower official source outranks a faster private one. The source base is not judged once and left alone — reliability is assessed continuously, the number of sources keeps growing past 350+, validation procedures are progressively strengthened, and any discontinued source is replaced quickly so coverage stays uninterrupted.

Extraction: collection schedule

Collection follows a fixed calendar. Automated systems gather data from official government sources during the last week of each monthly cycle. Processing and assessment then take place at the start of the following month, before the edition is published. This repeatable cutoff keeps each edition aligned to a consistent reference point.

Transformation: cleaning and standardizing

Every raw figure passes through a preparation stage before it enters any assessment. In this stage the value is cleaned, its format standardized, 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.

Source data arrives in many file and number formats, each provider following its own conventions. Three things happen to bring them onto common ground:

- * **Format conversion** — all incoming files and number formats are converted into a single internal convention.
- * **Unit standardization** — values are converted to the International System of Units using fixed conversion factors, so a figure means the same thing regardless of where it came from.
- * **Range validation** — each value is checked against expected ranges before any anomaly screening begins.

These consistency checks ensure that differences observed later between figures reflect the market, not formatting artifacts.

Trade data needs a further step. Official trade statistics aggregate transactions that can differ widely in volume, specification, and delivery conditions, and mixing them indiscriminately would distort the assessed price level. Trades are therefore filtered and clustered before use, applying criteria in four categories.

Table 3.2 The four filtering criteria applied to trade data.

CRITERION	PURPOSE
Minimum-volume thresholds	Exclude trades too small to represent the market level
Specification and grade	Keep together only trades matching the assessed specification
Location basis	Group trades by the location basis they reflect
Trade size and delivery terms	Prevent distortions from mixing trades of different sizes or delivery conditions

The filter criteria categories are disclosed; the specific numeric thresholds and the statistical techniques applied within them are proprietary. How the filtered and clustered trades are then turned into an assessed price — through the program's price types — is detailed in Chapter 6.

> **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.

Preparation is the first half of the transformation stage. The second half — normalization to a common comparison basis — leads into the modeling work described in Chapter 4.

Loading: into the assessment system

Figures that clear preparation are loaded into the internal system as a consistent, vetted foundation for analysis. From this point the data is on common ground — single file convention, standardized units, validated ranges, trades clustered by specification and basis — and ready for the modeling stage to act on it.

Source timing and freshness

Freshness is bounded by how quickly official data becomes available. Official trade statistics are generally released one to three months after the period they describe. This delay is a normal feature of official data: governments require time to compile, validate, and publish national figures, and the analysis is structured around that reporting rhythm rather than treating it as a shortcoming.

To keep coverage continuous through that lag, internally developed mathematical models generate estimates for periods still awaiting official release. Internally developed models are not a data source — they fill the gap until the official figure arrives. Every modeled value carries a data-status label, either Preliminary or Forecast, so an estimate is always distinguishable from a settled official value. Once the official data is released and verified, the estimate is replaced and the figure becomes Final.

The same mechanism absorbs disruptions to the source base. A source may be temporarily unavailable, still lagging, or permanently discontinued; in every case the published series does not break. Model estimates keep coverage continuous while a substitute source is incorporated quickly under the same official-first priority. Every such change is documented in the monthly, publicly available Release Notes, so the lineage of an affected series remains traceable — the record shows how and why the sourcing behind a figure changed over time.

The Collection-and-Integration Path

The figure below summarizes the path, from source extraction through transformation to loading into the system that feeds modeling:




Figure 3.1 Data collection and integration: raw source data is extracted, transformed, and loaded into the data warehouse.

CHAPTER 4

Modeling

SUMMARY

- The model applied to a given commodity and location is chosen case by case from five approaches, weighing data availability, timeliness, and fit to the targeted specification; the trade-based approach is the default wherever suitable data exists.
- When market data is scarce or absent, qualified analysts estimate values from related assessments, producer references, completed deals, and market fundamentals, keeping estimation a disciplined step rather than free-form judgment.
- Forecasts extend the most recent path about six months ahead, refreshed monthly as point values without confidence intervals.
- Normalization corrects for differences in delivery location, data quality, quantity, specification, and timing, so figures from diverse markets sit on a common, comparable basis.

Once raw information has been cleaned, standardized, and loaded, the modeling stage turns it into published price assessments. Primary Commodity Prices does not apply a single modeling recipe: the appropriate model is chosen per commodity and location, the stage absorbs cases where suitable raw data is scarce or absent, it extends recent trends a short way into the future, and it expresses every figure on a common, comparable basis. The five specific price models — and the mechanics that distinguish them — are detailed in Chapter 6; what follows is the ground common to all of them.

Regardless of price type, the shared modeling layer follows one flow — warehouse data is modeled into assessments, normalized to a common basis, and stored in the published database.



Figure 4.1 Modeling flow: warehouse data is modeled into assessments, normalized to a common basis, and stored in the published database.

Selecting the Model

Selection follows from what the market actually offers for a given commodity at a given location, weighing three things together: how available the data is, how timely it is, and how well it fits the targeted specification — the exact grade and form being priced. The choice is made case by case rather than scored: the trade-offs are conceptual, and there is no published ranking or error rate that elects one model over another.

Each model carries its own data requirement, and the model is chosen to match the conditions the market presents.

Table 4.1 What each modeling approach requires, and the conditions that favor it.

APPROACH	REQUIRES, AND BEST WHEN
Trade-based	Official trade statistics that are timely and suitable for the targeted specification — the preferred approach wherever such data exists
Formula-based	A stable relationship with related commodity prices and economic variables — used where the price moves predictably with measurable drivers
Freight-based	An existing assessment at another location, plus freight and insurance rates — used to carry a value to a location another assessment already anchors
Manufacturing-cost-based	Known production economics: input costs, utilities, labor, overhead — used where market prices are scarce but the cost of producing the commodity is known
Compiled	Several public-source price series that can be validated — used where a validated monthly average represents the market

The trade-based approach is the default wherever suitable trade data exists; the others come into play as the data picture thins. The individual mechanics of each are covered in Chapter 6.

When Raw Data Is Insufficient

When the available inputs are insufficient, inadequate, or unavailable — or when a purely data-driven approach would not yield a representative assessment — qualified analysts estimate values from a wide range of factual market information: related assessments, producer references, completed deal reports, and supply-and-demand fundamentals. Analysts verify all data used and may propose adjustments to the underlying models, so estimation stays a disciplined modeling step rather than free-form judgment.

The most common gap is timing. Some official statistics are released one to three months after the period they describe, so model-based estimates stand in for the period up to the current month, derived from related assessments that carry no such lag. These interim values are marked Preliminary and are expected to be revised once the official figure arrives. Every estimate is distinguishable from a settled official value by its data-status label, described later in this chapter.

New series and assessments, and changes that come out of a methodology review, are not switched on without a trial. New series and assessments undergo shadow testing for a period before official release, and methodology changes are announced to subscribers in advance — together with the duration of the transition period — before they take effect. This keeps a new or revised model under observation against live conditions before it carries published figures.

Keeping Models Accurate

Model evaluation runs continuously alongside the modeling itself, with a formal review of all methodologies at least once a year. The accuracy of the models — forecasts included — is evaluated during modeling, and computational algorithms and regression models are adjusted when necessary. When a series shows atypical behavior or an inconsistency surfaces, the data behind it is flagged for verification; this can trigger a recalibration of the model, or its replacement when it no longer reflects the conditions of the market it covers. Analysts also investigate the cause of each inconsistency — from an erroneous data entry to a market change the model has not yet accounted for — and feed those findings back into the models, so accuracy and reliability improve over time.

Short-Term Forecasts

For periods not yet reported by any official source, the modeling stage extends the most recent path forward. A forecast is a forward-looking estimate, generated by statistical and machine-learning models trained on historical series together with related economic variables — commodity prices, economic indices, and industry indicators. The models learn patterns from past behavior and from variables that move together with the series being projected.

A forecast reaches about six months ahead of the latest reported period. This horizon balances analytical usefulness against the rising uncertainty of longer projections, keeping forward-looking values close enough to observed conditions to stay informative. Because the underlying market parameters can shift quickly, forecasts are refreshed monthly, each one reflecting the most probable scenario as of its update date. Forecasts are published as point values; they do not carry confidence intervals.

The forecasting approach is stated openly — the model families, the variables they draw on, the horizon, and the monthly refresh. What stays proprietary are the parameters that produce individual values: the exact model architecture, the specific input features and their relative weights, and the underlying coefficients.

Normalization and Comparability

Normalization is the adjustment of gathered market information so it matches the base standard reflected in the published assessments. Markets are diverse — trades differ in delivery location, data quality, quantity, specification, and timing — and merely averaging them would not represent the assess-

ment value accurately. Normalization happens concomitantly and iteratively with the modeling work, not as a strictly sequential step before it. Each kind of difference has its own correction.

Table 4.2 How differences across sources are corrected during normalization.

DIFFERENCE ACROSS SOURCES	HOW IT IS CORRECTED
Delivery location	Freight rates
Data quality	Factors, indexes, and cross-reference with other sources
Quantity and specification	Averages and outlier detection
Timing between assessment curves	Mathematical formulas

Location is the correction most visible in the published series. Trades reflecting different delivery points are adjusted using freight rates, so every figure is expressed at the published location basis — typically the location where trading of that commodity is most common. The same underlying market information is carried to each published location through these freight adjustments, so a commodity priced at more than one point in the trade rests on a single comparable basis rather than a patchwork of unadjusted quotes.

Two further conventions keep the dataset internally consistent. Series published less often than the monthly cadence — quarterly or annually — are aligned to the monthly grid through interpolation, so every month can be compared across the full dataset. And at the end of the transformation stage, figures are rounded by maturity: historical and preliminary data to three significant figures, forecast data to two. The coarser rounding on forecasts is deliberate honesty about precision — forward-looking estimates carry greater uncertainty than observed values, and the published figure reflects that.

> **Example:** a price observed at a loading terminal and a price at a destination terminal describe the same commodity market at different points in the supply chain; freight adjustment brings both onto the commodity's published location basis so the two can sit side by side in the same series.

Data-Status Labels

The modeling stage produces the values in a series that are not yet verified — preliminary estimates standing in for lagging official data, and forward-looking forecasts. Each such figure carries a status label (Final, Preliminary, or Forecast) signaling how settled it is; the labels and the Preliminary-to-Final transition are defined in Chapter 2.

The reliability of preliminary and forecast values is tracked internally by back-testing each estimate against the final official value that later replaces it; that accuracy tracking is covered in Chapter 6.

CHAPTER 5

Information Publication

SUMMARY

- Every assessment clears a two-layer validation — automated cross-referencing followed by expert human review — before it is released.
- Expert review is risk-based: a flagged value is always checked, while a representative sample of unflagged values is reviewed each cycle as an independent check on the automated rules.
- Finished assessments are grouped into monthly editions, delivered in several formats, each edition superseding the prior one on a published release schedule.
- The methodology behind every figure stays publicly accessible with no login or subscription, and each change is recorded in the monthly Intratec Release Notes.

Before any price assessment produced by Primary Commodity Prices reaches subscribers, it passes a final review and is then made available in a defined form. Two steps stand between a finished assessment and the subscriber: a last cross-check on the value, and delivery in a defined format — while the methodology behind every figure remains openly accessible.

The Final Review

Every assessment clears a two-layer validation before release: an automated processing stage followed by expert human review. The automated stage prepares and tests the data, and market experts — analysts with sector-specific knowledge — act as the final quality gate, the last check a figure passes before publication.

The automated stage cross-references multiple official sources and applies mathematical models to identify inconsistencies between them. Cross-referencing compares the same data point across independent official records to confirm it agrees. Running the same consistent, validated rules every cycle reduces errors introduced through manual handling and removes bias from market participants.

Expert review then examines the model outputs, adding contextual judgment that automated rules alone cannot provide. It keeps processing consistent while confirming that outputs remain sensible in real market conditions. This review is risk-based rather than exhaustive: it covers a representative sample of each cycle's output.

Table 5.1 How the final review allocates expert attention

WHAT THE VALUE IS	HOW IT IS REVIEWED
A value flagged as anomalous or inconsistent by the automated stage	Always routed to expert review, concentrating attention where the probability of an issue is highest
A value the automated stage did not flag	A sample is reviewed each cycle, an independent check that the automated rules are not silently missing problems

The combination directs expert judgment where it matters most while keeping the entire pipeline under human oversight. Only after a finished assessment clears both layers does it move to publication. The cadence on which review and release recur, and the status labels a released figure can carry, are described in Chapter 2; the modeling that produces the underlying prices is described in Chapter 6.

Clearing that gate is the first of two points at which a figure is held to account, not the last. A preliminary or forecast value is later measured against the official statistic that eventually replaces it, so the same figure is tested twice — before release, against independent sources, and after it, against the outcome it anticipated. Publication does not close the book on a figure; it releases it into a record that keeps tracking how well the estimate held up. Because that comparison runs for every preliminary and forecast value, the reliability of the modeling is re-tested against real outcomes on every cycle rather than asserted once and assumed to hold.

How Assessments Are Published

Published assessments are organized into monthly editions. Each edition presents the assessments together at a common reference moment, so the figures it contains can be read and compared as a coherent set, and each new edition supersedes the prior one to build a continuous, regularly spaced record over time. Exact release dates for each series are listed in the published release schedule at <https://help.intratec.us/account-and-subscriptions/release-schedule/#release-schedule>. The detailed collection-and-processing timing that defines an edition — and how revisions are handled after release — are covered in Chapters 2 and 7.

Once an edition clears the final review, it is published and made available in several delivery formats — Online Charts, the web API, and spreadsheet and BI add-ins — so the assessments can be read, queried, and worked with in the form each task requires. The access and delivery options tied to each subscription are documented in the Intratec Help Center.




Figure 5.1 The publication flow: reviewed database content is published and delivered through Online Chart, API, and spreadsheet or BI formats.

Open Access to the Methodology

The reasoning behind every published assessment stays publicly available. All methodology documents are accessible with no subscription or login required, so any reader can trace each step from raw inputs to finished assessment, independent of access to the assessments themselves.

That openness extends to how the record changes over time. Every data revision, methodology change, assessment addition or retirement, and correction is documented each month in the publicly available Intratec Release Notes, so a figure never changes silently and the lineage of every series remains traceable. The aim is to make the path from raw data to published figure transparent: every step is documented and open to inspection.

CHAPTER 6

Prices Modeling

SUMMARY

- Five price types cover the program — trade-based, formula-based, freight-based, manufacturing cost-based, and compiled — each matched to the commodity, the available data, and the point in the supply chain being priced.
- Preliminary estimates and short-term forecasts fill the one-to-three-month gap left by delayed official statistics, keeping every assessment current between consolidated data releases.
- The reliability of preliminary and forecast prices is tracked internally by back-testing each estimate against the final official value, which drives model recalibration rather than a published accuracy figure.

Primary Commodity Prices converts raw market data into published price assessments through a modeling stage that validates price, volume, specification, and location, removes outliers, and applies the model matched to each assessment. The path depends on the price type, which depends on the commodity, the available data, and the point in the supply chain being priced. Five price types cover the program, complemented by the treatment of data shortages and short-term forecasts.

Trade-Based Prices

Trade-based prices come from international trade information reported by countries — what was traded, at what volume, and at what declared value. Candidate records are admitted by a criterion of **data homogeneity**, a check that a set of trade records all refer to the same assessment: the same commodity, with similar specifications and comparable trade conditions. A perfectly homogeneous dataset would show no variation, but real trade data carry some spread even within a single specification, so a tolerable homogeneity level is accepted rather than demanded absolutely. A clustering algorithm then groups similar transactions by their identified statistical modes and data dispersion, discerning underlying market patterns, dampening anomalies, and yielding a single price assessment per cluster — avoiding the “basket” effect that arises when heterogeneous trades are averaged together. A raw dataset in which no homogeneity can be found is treated as unqualified for price calculation and is not used.

From homogeneous data, two trade-based price types follow, differing mainly in the statistical treatment applied:

- * **Unit value** — the total traded value divided by the total quantity for a given country and month, taken directly from customs authority reports with no statistical treatment.
- * **Transaction price** — a filtered version that improves data quality through two sequential filters. First, a minimum-volume filter removes low-volume records: because trade data are rarely reported at the individual-transaction level, the minimum is set as a multiple of a typical trade size — often two or three times the standard cargo size for an internationally traded commodity, or a multiple of the most common transport unit for domestic and small-quantity trades. Second, gross outliers — prices that deviate sharply from the median of all transactions — are removed using the Median Absolute Deviation, a dispersion measure more resilient to outliers than the standard deviation.

A raw unit value is the simplest of these figures; its lack of statistical treatment carries three limitations against the transaction and compiled prices published alongside it:

- * **Aggregation bias.** Customs records group heterogeneous products under broad classifications, so a unit value blends distinct commodities into one average that can drift away from the price of any specific grade actually traded.
- * **No outlier removal.** Because no statistical treatment is applied, a unit value does not exclude implausible or erroneous transactions — the filters that transaction prices apply at both the volume and the price level.
- * **Lack of market specificity.** A unit value averages across market destinations, price types, and commodity specifications at once — segmentation (wholesale versus retail versus industrial; spot versus long-term contract; grade) that compiled prices preserve.

A unit value therefore serves as a transparent, unadjusted reference, while transaction prices and compiled prices give the more robust basis for tracking a specific commodity market. Because official trade statistics publish with a delay, **preliminary prices** fill the interim gap, standing in as model-based best estimates until the consolidated figures arrive.

Formula-Based Prices

Formula-based prices are calculated as a function of the prices of related commodities, economic indicators, currency rates, and similar variables; they are used where direct market data are unavailable or sparse. Market research first identifies which commodities and economic factors most likely correlate with the assessment — typically its raw materials or primary derivatives. A regression model is then fitted, its parameters calibrated so the model reproduces historical assessment data. The fitted model is applied to current market prices and indicator values to generate the assessment, which keeps the currency and quantity unit of the series it belongs to, so it stays directly comparable with the other prices in that series. Assessments produced this way are identified as formula-based.

Freight-Based Prices

Freight-based prices carry a price already established at one location to a different point in the supply chain, by adding or subtracting the cost of moving the commodity between two delivery points — letting a price set at one location stand in for the same commodity at a second point when direct trade data there are thin or absent. Two assessment types follow from the direction of travel:

- * **Netforward** — from a price at the loading terminal (an FOB basis), maritime freight and insurance are added to express the delivered cost at the destination terminal (a CIF basis).
- * **Netback** — from a delivered price at the destination terminal (a CIF basis), insurance and maritime freight are subtracted to recover the value at the loading terminal (an FOB basis).




Figure 6.1 Netforward builds a delivered (CIF) price by adding maritime freight and insurance to an FOB price; netback works in reverse, recovering an FOB price by subtracting insurance and freight from a CIF price.

The freight rate reflects the transport form most commonly used for each assessment, even where other modes are sometimes used. Freight rates and insurance are themselves assessed values entering the calculation alongside the source assessment, and the result is reported in that assessment's currency and quantity unit — so a freight-based price sits on the same basis as the other prices in its series: a function of another assessment, the applicable freight rates, and insurance.

Manufacturing Cost-Based Prices

Manufacturing cost-based prices estimate the operating cost of producing a commodity from its process inputs — the cash outlay required to run a plant. The build-up spans raw materials net of by-product credits, utilities, labor, maintenance, operating charges, plant overhead, and local taxes and insurance:

- * **Raw materials and by-product credits** — the cost of feedstocks and other inputs, net of revenue from saleable by-products or waste streams recovered during the process.
- * **Utilities** — electricity, steam, fuel, water, and other energy sources required for production.
- * **Labor** — direct and indirect operating staff.
- * **Maintenance** — upkeep, repairs, and replacement of process equipment and infrastructure.
- * **Operating charges** — catalysts, consumables, and other process-dependent costs.
- * **Plant overhead** — facilities management, administration, and shared services allocated to production.
- * **Local taxes and insurance** — statutory and operational taxes, plus coverage for equipment and liability.




Figure 6.2 Manufacturing cost build-up: gross raw materials cost less by-product credits and net utilities form the variable cost; operating labor through plant overhead, taxes, and insurance form the fixed cost; together they total the operating cost.

The raw-material and utility inputs feeding the cost model are themselves Intratec-assessed figures, drawn from three sources:

- * **Raw materials and by-product credits** — Intratec Primary Commodity Prices, the commodity price assessments described throughout this document.
- * **Energy references** — Intratec Energy Prices & Markets, which assesses feedstock and fuel prices such as naphtha, natural gas, and power.
- * **Industrial utility costs** — Intratec Industry Economics & Competitiveness, which publishes industrial utility cost estimates.

Because every raw material, by-product, and utility cost feeding the model traces back to its own published assessment, a manufacturing cost-based price stays auditable end to end — and because those input series also carry forward-looking estimates, the price can be projected using them, as a function of process indicators, raw-material and utility prices, labor, and an industrial-cost index.

Compiled Prices

Compiled prices are built from price data published by public sources — price series stated as a value per quantity over a month. The raw data first pass a formatting step adapting them to the internal data structure and units, then statistical validation detects and removes anomalies and outliers. These assessments are typically computed as a monthly average of closing prices for daily or weekly published data, then normalized on the same footing as every other assessment.

Preliminary Prices and Forecasts

When official data are insufficient, inadequate, or unavailable — or when a data-driven approach would not yield a representative assessment — analysts estimate values from a wide range of factual market information: other published assessments, producer or export indices, labor costs, completed transactions, spread and exchange trades, and supply-and-demand fundamentals. Each estimate follows a rigorous procedure: the market data used are collected and verified, and a mathematical model grounded in observed market behavior is proposed.

Preliminary data are the most common case. They are produced by proprietary models that address the one-to-three-month lag with which official statistics are typically released, providing the best available estimate for the period up to the current month rather than waiting for the consolidated release. Preliminary assessments are marked with “(P)” in publications.

Short-term price forecasts are provided for selected commodities, derived from models that combine the prices of related commodities, economic indices, and industry indicators. A forecast reflects the most probable scenario for the upcoming six months as of the date it is generated. Because the underlying market parameters can shift significantly within short periods, forecasts are refreshed monthly. The reliability of both preliminary and forecast prices is tracked internally rather than published as a metric: each estimate is later compared against the final official value that replaces it, and that Preliminary-to-Final back-test drives model recalibration whenever an estimate drifts from what the official data ultimately reports.

Normalization



Figure 6.3 Whichever of the five price types produces an assessment, a shared normalization stage restates it on one common, comparable basis — the same currency, unit, and timeline as every other assessment in the program.

As a shared final stage, every assessment — whichever price type produced it — is normalized to a common, comparable basis, so a compiled price arrives in the same form as a trade-based, formula-based, freight-based, or manufacturing cost-based one. The mechanics of that adjustment are covered in full in the comparability article on the methodology site rather than repeated here.

CHAPTER 7

Disclaimers and Limitations

SUMMARY

- Every assessment change, correction, and retirement is documented in the monthly, publicly available Release Notes, keeping the published data auditable rather than silently revised.
- No commercial relationships are maintained with market participants, and the modeling approach is fully public while the calibration parameters stay protected.
- Published figures carry real boundaries — dependence on official source data, modeled estimates where data is missing, and short-term forecasts issued without confidence intervals.
- No single “right” price exists for a commodity, so the assessments are best read as consistent, robust estimates and combined with other indexes when used as a contract reference.

Primary Commodity Prices is built to be auditable end to end: assessments evolve as the underlying data evolves, every change is documented in a public record, and the methods that produce each price are disclosed openly. Series are added, changed, and retired as their data allows; corrections are made and traced; independence and impartiality are preserved throughout; and the uncertainty in each figure is signaled to its reader — all within defined boundaries on how a published price should be read.

Assessment Additions, Changes, and Retirements

A series may be initiated, terminated, or changed at any time, driven by the availability and quality of raw data from the sources in each country. Because the vast majority of data comes from public official sources, the continuity of an assessment ultimately depends on those sources continuing to publish data that meets the assessment’s requirements. The governing commitment is quality over continuity: an assessment that can no longer be produced reliably is changed or retired rather than published at a lower standard.

When data for an assessment is insufficient, inadequate, or unavailable for a prolonged period, one of four approaches is adopted to keep the published series representative.

Table 7.1 The four approaches when an assessment's data falls short over a prolonged period.

APPROACH	WHAT IT DOES
Use a model	Estimate the assessment with a model built on good previous and/or related data.
Change the data source	Replace the failing source with another that meets the requirements.
Alter the assessment	Adjust the assessment's definition so it can be produced reliably from the data available.
Discontinue the assessment	If the lack of suitable data persists, retire the series rather than publish unreliable figures.

Short-term gaps — a source that is merely lagging or temporarily unavailable — are not handled this way. They are addressed earlier in the pipeline through model-generated Preliminary estimates and quick source replacement, as covered in Chapter 6.

Relevance over time is maintained through the methodology review process itself, and subscriber feedback is a standing criterion in those reviews — alongside changes in the raw data and modeling enhancements — so the assessments evolve with what readers report needing clarified or improved. The review cadence and the steps a revised method passes through before it takes effect are detailed in Chapter 2.

Data Corrections and Their Documentation

Each price follows a three-stage lifecycle, and the transitions between stages are themselves a form of revision. A **Forecast (F)** is a model prediction for a period that has not yet been reported. Once that period becomes current but official data has not yet arrived, the figure becomes a **Preliminary (P)** estimate produced by the same model. When verified official statistics are released, the figure becomes **Final (Fi)**. Each transition replaces the prior value as better data becomes available.

Beyond that normal progression, a published figure 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; or a submitted assessment complaint is resolved, prompting an adjustment to the affected figure. The full validation and complaint mechanism is described in Chapter 2.

Every change is documented in the Intratec Release Notes, published each month and publicly available. The Release Notes record data revisions, methodology changes, assessment additions and retirements, and corrections, so any reader — subscriber or not — can trace how and why a figure changed over time. This monthly public record keeps the published data auditable: a number never changes silently.

Each monthly release runs the same workflow before anything reaches subscribers: the Online Chart and API databases are updated, a test-presentation check confirms the revised figures display correctly, the Release Notes entry is prepared, and the changes are communicated. Routing every correction through

this fixed, repeatable path is what keeps a revision from slipping out unrecorded — the same discipline that governs a routine monthly update governs an out-of-cycle correction.




Figure 7.1 The monthly release workflow: Online Chart and API database update, test presentation display, release notes, and communication.

The same record makes the path from complaint to correction fully traceable. A formal assessment-complaint process is available to subscribers who wish to question any published figure. A submitted complaint is investigated, and when the investigation confirms an issue, the affected figure is corrected and the change is documented in the same monthly Release Notes — so the route from a raised question to a published correction can be followed by anyone.

Methodology changes — as distinct from individual figure corrections — are additionally announced to customers in advance, with a defined transition period, before they take effect.

Independence and Impartiality

The analysis rests on impartial, auditable methods applied to data drawn, for the vast majority, from public official sources — statistics bureaus, trade agencies, multilateral organizations, and market exchanges. Private sources and producers' data are used only with permission from the rights holders, and Primary Commodity Prices is not a primary source for pricing data.

No commercial relationships are maintained with the participants of the markets covered — the producers, traders, and other firms active in those markets — and surveys are used only in very exceptional circumstances. This removes the incentive for any single party to influence the results, and keeps the assessments independent of the interests they measure.

Transparency is maintained without exposing internals. The methodology approach is fully public, requiring no login or subscription: the data sources, the five modeling approaches, the assessment structure, the baselines, and the data status labels are all disclosed. What remains proprietary is the precise calibration — the model parameters, regression coefficients, and normalization constants. The approach is disclosed; the parameters are protected.

How Uncertainty Is Communicated

Every figure carries a data status label — Final, Preliminary, or Forecast — signaling how settled the underlying value is, displayed directly alongside the value in data tables and charts so the maturity of any number is visible at the point of reading. Precision itself also communicates uncertainty: historical and Preliminary values are published to three significant figures, while Forecasts carry two, reflecting the greater uncertainty of projecting a period not yet reported. And when official data arrives, a Preliminary figure is replaced by the Final value, so an interim estimate is never silently left standing.

Limitations of the Analysis

Published figures should be read against the boundaries of the data and methods that produce them.

- * **Dependence on official source data.** Figures depend on official source data, which can lag or be revised later. Where official statistics are delayed, Preliminary estimates fill the gap until verified data arrives.
- * **Modeled estimates.** Some values are modeled estimates, used where official data is missing or insufficient. The data status label on each figure signals how settled the underlying value is.
- * **Forecasts without confidence intervals.** Short-term forecasts are estimates of the most probable scenario for the period ahead, published without confidence intervals; because the underlying market parameters can shift quickly, they are refreshed each cycle. How Preliminary and Forecast figures are produced and how their reliability is tracked is described in Chapter 6.
- * **Coverage follows the trade.** Assessments are produced for the countries and world regions where official trade in a commodity supports them, not for a fixed universe of locations. Where trade at a specific location is too thin to assess reliably, its value is modeled from related data or left unpublished for that location rather than reported at a lower standard.

Two further points frame how the prices should be interpreted. No single “right” price exists for a commodity: the price actually paid depends on factors such as order size, trade terms, geographical location, and contract type. Rather than claim one definitive value, every published figure is constructed through a consistent, robust methodology designed to deliver reliable estimates — which is also why accuracy is measured against consolidated historical values produced by that same methodology for the specific assessment, not against any single market price. And when the data serve as an index for structuring contracts between suppliers, processors, and end customers, they are best combined with other price indexes rather than relied on as the sole reference: this reflects no concern about accuracy, but ensures a contract benefits from the broader perspective and robustness that a wider range of market insight brings.

Usage limitations specific to each product — what each dataset should and should not be used for — are documented in the Intratec Help Center.

CHAPTER 8

Methodology References

SUMMARY

- The methodology stands on an established body of published reference literature — statistical manuals, textbooks, and technical papers on commodity markets and robust price estimation.
- Recurring terms are defined in the online Glossary, and the methodology site always carries the most recent edition of the methodology.

The methodology rests on an established body of reference literature — textbooks, statistical manuals, and technical papers on the prices and markets of chemical and petrochemical commodities, and on the robust-statistics methods used to estimate prices from trade data. This bibliographical foundation, drawn on throughout the development of Primary Commodity Prices and the elaboration of its methodology, is listed below, alongside where recurring terms are defined and where the up-to-date methodology can be found.

Bibliography

Table 8.1 Reference literature behind the Primary Commodity Prices methodology.

NO.	REFERENCE
1	Atkinson, A. C. et al. 2017. Robust Bayesian Regression with the Forward Search: Theory and Data Analysis. Test, vol. 26, pp. 869–886.
2	International Monetary Fund. 2009. Balance of Payments and International Investment Position Manual (6th ed.).
3	International Monetary Fund. 2009. Export and Import Price Index Manual. Organisation for Economic Co-operation and Development.
4	Mankiw, G. 2020. Principles of Economics (9th ed.). Cengage.
5	Miao, G.; Wegner, E. 2022. Using Unit Value Indices as Proxies for International Merchandise Trade Prices. SDD Working Paper No. 111. Organisation for Economic Co-operation and Development, Paris.
6	Perrotta, D. et al. 2020. The Robust Estimation of Monthly Prices of Goods Traded by the European Union. EUR 30188 EN. Publications Office of the European Union, Luxembourg. ISBN 978-92-76-18351-8. doi:10.2760/635844.

NO.	REFERENCE
7	Riani, M. et al. 2018. The Use of Prior Information in Very Robust Regression for Fraud Detection. International Statistical Review, vol. 86(2), pp. 205–218.
8	Rais, S. 2008. Outlier Detection for the Consumer Price Index. In Proceedings of the Survey Methods Section: Statistical Society of Canada Annual Meeting, vol. 110.
9	Schofield, N. C. 2021. Commodity Derivatives: Markets and Applications (2nd ed.). John Wiley and Sons.
10	Spyros, A. et al. 2015. The Estimation of Fair Prices of Traded Goods from Outlier-free Trade Data. EUR 27696 EN. doi:10.2788/57125.
11	Temere, Dawit Sisay. 2016. Disaggregated International Trade Prices. Nationaløkonomisk Tidsskrift.
12	United Nations. 1981. Strategies for Price and Quantity Measurement in External Trade: A Technical Report. Series M, No. 69.

Defined Terms

Recurring terms used throughout the methodology — the price types and data status labels — are defined in full in the Glossary on the methodology site: methodology.intratec.us/glossary.

The Up-to-Date Methodology

This document reproduces the methodology as of the edition date printed on its cover. The methodology itself is reviewed continuously — data sources are checked and replaced, models are recalibrated, and clarifications from subscriber feedback are folded in, with every change documented in the monthly Release Notes — so a downloaded or printed copy gradually falls behind. The most recent version is always the one published on the methodology site, publicly available with no subscription or login required: methodology.intratec.us.