SAMPLE

The CPI Report

The Number That Defined an Era — Trading Inflation Data When Every Print Moves the

Aiman Sterling

Free sample: Introduction & Chapter 1

putrabooks.github.io

Buy the full book:

Ko-fi

Payhip

Chapter 1: Introduction

For roughly two years — from late 2021 through late 2023 — the monthly Consumer Price Index release was the single most consequential event in global finance. Not the Fed's meetings, not the jobs report: the CPI print. On release days, equities swung by full percentage points within minutes; on the era's most famous print (Chapter 8), the S&P 500 reportedly staged one of its largest single-day rallies of the decade — on an inflation number that was merely less bad than expected. Chairs of central banks abroad scheduled commentary around it. It became, in the market's own vocabulary, the event.

This book is about that number: what it is, how it is built, why the market prices its details rather than its headline, and how a prepared trader survives and occasionally profits from the monthly eight minutes when it arrives.

It is the fifth book in the Impact News line, and the second of the series' data-event books. Book 4 took the framework to the employment half of the Fed's dual mandate. This book takes it to the price half — and the inflation data differs from the labor data in ways that change the trading entirely: it moves more slowly (inflation has momentum; labor turns faster), it is watched more politically, and its measurement is more contested (the index is a survey-based estimate of a price level, with methodological debates the public follows closely). The market's treatment of CPI reflects all three: the prints are regime-defining, the reactions are regime-dependent, and the components are fought over.

Why CPI, and what this book covers

The chapters ahead, in order: what the index actually measures and how the Bureau of Labor Statistics builds it (Chapter 2); the headline-core distinction and the month-over-month versus year-over-year arithmetic — including base effects, the most underappreciated force in inflation data (Chapter 3); the component structure and the desks' decomposition of the print (Chapter 4); the shelter lag, why the biggest CPI component systematically misleads, and the market signals that lead it (Chapter 5); the expectation machinery — consensus, breakevens, and nowcasts (Chapter 6); the reaction structure and its regime tables (Chapter 7); the case studies of the era (Chapter 8); the preparation artifact — the CPI sheet (Chapter 9); the execution mechanics of the 8:30 release (Chapter 10); the risk architecture, which inherits Book 4's data-event rules and adds the inflation-specific conflicts (Chapter 11); the inflation data map — PPI, PCE, and the alternatives, with the bridge to Book 6 (Chapter 12); and the action plan (Chapter 13).

The three ideas to carry before starting

1. Inflation has momentum. Employment can turn in a quarter; prices are levels, and

their rate of change changes slowly. This is why CPI prints correlate month to month more than payrolls prints do — and why surprises have regime weight: an inflation surprise is evidence about a trajectory, not a month.

2. The market trades the core MoM, not the headline YoY. The bolded number in the headlines is the year-over-year total; the number the desks grade first is the month-over-month core — the fresher, less-volatile read on the underlying trend. A print can produce a screaming headline and a calm market, or the reverse; the gap between the two reads is Chapter 3's arithmetic.

3. Inflation data is Fed data with no equal. The Fed's statutory mandate names prices, its target is defined on a different index (Book 6's subject — the PCE), and its reaction function through the modern era has been transparently inflation-contingent. CPI is not the target gauge, but it is the first-read gauge: it arrives weeks before PCE, and the market prices the Fed's response through it. Understanding that two-gauge structure — CPI as the signal, PCE as the scoreboard — is the inflation era's master key, and this book and Book 6 together build it.

One note on method, as always: the index's construction is described as the Bureau publicly documents it; the historical episodes are described as widely reported, with hedged language wherever a figure or detail is not fully certain. The CPI is among the most-audited statistics in existence — published in full, methodologized in public, revised rarely and transparently — and the trader's edge is never the number itself, only the preparation to read it.

Begin where the number begins: a survey of prices, weighted into an index.

How this book fits the series

The Impact News series has, by this volume, two well-established layers: the event framework (Book 1's expectation gap; the regime maps and reaction tables of Books 2 and 4) and the event-specific mechanics (each book's deep dive into its own event). This book inherits the framework without re-teaching it and spends its pages on the CPI's own mechanics: the index's construction (Chapter 2), the arithmetic that turns the construction into preparation (Chapter 3), the components that turn the average into a verdict (Chapters 4 and 5), the expectations structure that turns preparation into positioning reads (Chapter 6), and the era's record as the framework's demonstration case (Chapter 8).

What this book adds to the series is the composition discipline — the method of reading a release through its parts rather than its average. Payrolls (Book 4) teaches the two-survey read; the CPI deepens it into the full component method: weights, pipelines, base effects, and the narrative pre-assignment of every category before the release. A trader who has absorbed this book reads any average-number release — CPI, PPI, any index print — as a stack of stories, and the average merely as the stack's headline.

And one era-specific orientation: the inflation years made this report binary and

violent, and that intensity is recent memory rather than permanent condition. The book teaches both regimes — the fight era's outsized reactions and the calm regime's quiet prints — because the trader preparing today trades the regime they are in while building the file for the regime that comes next. The 2021 turn arrived in the tape's reaction changes several prints before it arrived in the narratives; the discipline that catches the next one is this book's.

The event's own hierarchy of learning

The material ahead builds in a fixed order, and the order is the reading protocol: construction (what the number is) before arithmetic (what the number means) before components (what the number is made of) before expectations (what the number is priced for) before reactions (what the number does) before history (what the number did) before preparation and execution (what you do about it). Each layer assumes the ones before it; the trader new to the series can read this book alone with the two frameworks in hand — the expectation gap and the two-wave event structure — and be fully equipped by Chapter 13.

Chapter 2: What CPI Measures

The Consumer Price Index answers a deceptively simple question: how much more (or less) does it cost to buy the same basket of consumer goods and services this month than in a base period? The machinery behind that question — the basket, the weights, the price collection, the index arithmetic — explains both what CPI tells you and what it cannot, and every trading decision about the print rests on that machinery.

The basket and the weights

The Bureau of Labor Statistics constructs CPI from a basket: thousands of goods and services, organized into categories, each with a weight representing its share of consumer spending — the weights derived from household expenditure surveys, updated on a rolling basis in the modern methodology. The major categories and their rough modern magnitudes (figures approximate, and they shift with each weight update): shelter around a third of the index; food and energy together around a fifth; and the remainder spread across vehicles, medical care, apparel, recreation, education, and the rest of the consumption universe.

The weights matter because they determine what the index is: CPI is not a price, it is a spending-weighted average of prices. A 20% surge in a category weighted at 1% moves the index a fifth as much as a 5% surge in a category weighted at 20%. When the market decomposes a print — Chapter 4's work — the weights are the conversion table between category news and index arithmetic.

The price collection

Prices are collected monthly for most categories: the Bureau surveys tens of

thousands of retail establishments and service providers across urban areas, recording the transaction prices of the specific items in the basket. Certain categories are collected less frequently or modeled: the sampling is a genuine survey with genuine error bands; some notoriously difficult categories (medical services, for one, with its complex insurance and billing structure) use extensive modeling whose assumptions the Bureau documents and analysts debate.

Seasonal adjustment then removes the predictable seasonal patterns — apparel's clearing cycles, travel's peaks, the regular annual resets in some services — because the market wants the underlying trend, and seasonal factors are periodically re-estimated, which is why January and other reset-month prints behave with extra noise.

The index arithmetic

From the collected prices, the Bureau computes price relatives for each item, aggregates them within categories using the weights, and produces the index level. The market's working numbers derive from that level:

- The month-over-month change (MoM) — the percentage change from last month's level: the freshest read - The year-over-year change (YoY) — the percentage change from twelve months ago: the politically quoted number - Both, computed twice: for all items (headline) and for all items less food and energy (core)

A standing technical note, documented by the Bureau: the index level itself is occasionally revised (the seasonal adjustment factors), but the published rates — the MoM and YoY numbers the market trades — are effectively final at release. This makes CPI's revision structure the opposite of payrolls (Book 4's Chapter 6): the labor print is a draft revised for months; the CPI print is, for trading purposes, permanent at 8:30. What changes is not the print but the interpretation — and the follow-on gauges that confirm or contradict it.

The sampling truth, honestly stated

The CPI's measured level is a statistical estimate of an unobservable quantity: "the" cost of living. The sampling is excellent by survey standards and still an estimate — the Bureau publishes variance estimates, and the monthly MoM core print's sampling error is not zero, a fact the market's commentary rarely acknowledges when it celebrates or panics over a 0.1% difference from consensus. The honest practice: treat differences of one tenth within the sampling noise as ties, and let the pattern across months carry the conviction. The desks' habit of rounding grades — "core 0.2 versus 0.2 expected: in line" — is not laziness; it is statistical respect.

What CPI is not

Three standing misreads to retire before the detail chapters:

Not the cost of living. CPI measures a fixed-concept basket's price change, not any

individual's experience — your personal inflation depends on your basket, and the index's weights are the average household's. The public debate about CPI's fidelity to lived experience is a real and interesting debate; it is not a trading input. The market prices the index, the index is defined, and the definition is the trade.

Not the Fed's target. The Fed's 2% target is defined on the PCE price index — Book 6's subject — with CPI running persistently hotter than PCE by roughly a few tenths of a percentage point in the modern era, for structural reasons Chapter 12 and Book 6 cover. CPI is the early read on what the PCE will show, and the market prices the translation.

Not a rent measure. Shelter, the largest component, measures housing cost as an owner or renter experiences it — dominated by owners' equivalent rent, a survey-based construct explained in Chapter 5 — and not market rents, which are a different, faster series. The single most consequential fact about modern CPI trading is this gap, and the next chapter's arithmetic sets it up.

How the index is actually built

The Consumer Price Index measures the average change over time in the prices paid by urban consumers for a market basket of goods and services — a sentence worth unpacking, because each clause carries trading meaning.

"Urban consumers." The index covers the urban population — the large majority of households — via two overlapping populations (all urban consumers; wage earners and clerical workers historically), and the coverage choice is why the index is a consumer gauge: it prices what households buy, not what businesses pay (PPI's world) or what the economy produces (PCE's broader net).

"A market basket." The basket is constructed from consumer expenditure surveys, updated on a rotating chain basis in the modern methodology — the weights reflecting several years of spending patterns, with a well-known lag: the basket you trade against is a picture of spending as surveyed in prior years. The trading consequence is structural: CPI weights change slowly while economies change fast, and the pandemic era's violent category swings (goods spikes, services collapses and rebounds) were partly a weights story — the index priced the old basket through the new world until the rotations caught up.

"Prices collected." The Bureau of Labor Statistics collects prices monthly across categories and geographies — a mix of the traditional pricing samples and (increasingly, and with well-publicized friction in the shelter categories) administrative and transaction data. The collection machinery's calendar is the release calendar's ancestor: the release prints on the month's data, which is why the mid-month timing, and why the release lag is roughly a half-month — the market trades December's inflation in mid-January.

"Indexes." The output is an index level (1982-84 = 100 in the headline series), not a price. The trading numbers are the changes: MoM, YoY, seasonally adjusted for the

short horizon. The index structure is what makes the base-effect arithmetic of Chapter 3 possible — the YoY is literally this month's index over last year's same-month index, and last year's number is published, known, and sitting in your calendar.

The two production lines, briefly

The CPI program publishes two families: the national series (the headline event: the all-items and core indexes, the detailed expenditure categories, the regional breaks) and the special series (owner's equivalent rent, the chained CPI, the trim/median variants of Chapter 12's map). The trader's daily driver is the national series' first page: all-items MoM and YoY, core MoM and YoY, and the category detail — the four numbers and the stack of components that Chapters 3 and 4 build the read upon.

The production lines matter for one risk-specific reason: the series' revisions policy. The CPI is essentially not revised. Unlike payrolls (Book 4's two-month revision cycle), the printed index stands, with only rare, announced methodological corrections. The trading consequence is clean: what prints at 8:30 is the number — the gap you graded is the gap that stands — and the "wait for the revision" discipline of the jobs report does not transfer. The release is terminal; grade it once.

Chapter 3: Headline, Core, and the Arithmetic

Two distinctions organize every CPI conversation: headline versus core, and month-over-month versus year-over-year. Both pairs sound like technical trivia and are in fact the entire trading structure — they determine which number the market moved on, and why the same print can generate opposite narratives in different months.

Headline versus core

Headline CPI ("all items") is the full basket: food, energy, shelter, everything, weighted by expenditure share. It is the honest cost of the basket — the number a household experiences — and the politically quoted one.

Core CPI ("all items less food and energy") strips the two most volatile categories. The rationale, standard in central banking since the inflation-targeting literature: food and energy prices swing on supply shocks (weather, oil politics) that say little about the underlying inflation trend — the demand-driven, wage-price, monetary-relevant part. Core, smoothed of the noise, is the better estimator of where inflation is heading; headline, unsmoothed, is where inflation is.

The market's hierarchy at 8:30, in the modern era: core MoM graded first (the trend read), headline MoM second (the reality check), YoY numbers last (the political read, the slowest-moving). A print with a hot headline (an energy spike) and an in-line core is a calm print for the market — the desks dismiss the energy noise and price the trend. A print with a soft headline (an energy collapse) and a hot core is a hawkish print — the underlying trend is running hot regardless of the headline's relief. The

headline-core conflict is the release's most reliable whipsaw generator: the machines price the headline in the first second (it is the bolded number), and the desks reprice on the core over the first minutes.

Month-over-month versus year-over-year

The YoY number — "inflation is running at X percent" — is the average of the last twelve MoM prints, plus the twelve before that as its base. This produces the single most underappreciated mechanism in inflation trading: base effects.

The arithmetic, worked once, is permanent knowledge. YoY = (this month's index level ÷ last year's same-month level) − 1. The YoY change between this month and next month therefore depends on two things: next month's new MoM print, and which old MoM print is falling out of the calculation — the one from twelve months ago that exits the window. When last year's month was abnormally high, its exit mechanically lowers next month's YoY even if the new MoM is ordinary. When last year's month was abnormally low, its exit mechanically raises the YoY.

The trading consequences:

YoY turns are predictable months ahead. The base-effect calendar — which old prints enter and exit the YoY in coming months — is computable from published data, and analysts publish it. The disinflation of mid-2023, for instance, owed a large mechanical share to the exit of 2022's giant MoM prints — a fact the era's coverage debated in real time ("how much of this disinflation is base effects?"). A prepared trader knows the calendar: this month's YoY will drop mechanically by roughly X tenths regardless of the new print. The market knows it too — the base effects are in the consensus — but the surprise arithmetic still runs through them, and the YoY-versus-consensus comparison is only interpretable with the calendar in view.

MoM surprises are the real news. Because base effects are public arithmetic, the genuine information in every release is the new MoM prints — headline and core — versus their consensuses. The MoM is the monthly increment of the trajectory; the YoY is the smoothed average of history. When commentary says "inflation surprised," the tradeable version is "core MoM surprised" — everything else is arithmetic or politics.

The YoY's momentum. Because the YoY averages twelve months, it moves slowly and resists reversal: a single soft MoM print barely dents a hot YoY, which is why the market — needing to forecast the Fed's response over quarters — extrapolates the MoM trend rather than the YoY level. The desk shorthand "the trend is the target" means: the market prices where the MoM sequence implies the YoY will be in six months, not where the YoY is today.

The four numbers, graded in order

The release-morning grade, on the CPI sheet in exactly this sequence:

1. Core MoM vs. consensus — the trend read; the market's first-grade number 2.

Headline MoM vs. consensus — the reality check; conflicts with (1) produce the whipsaw 3. Core YoY and headline YoY vs. consensus — the political read, checked against the base-effect calendar 4. The components beneath the core number — Chapter 4's work: what made core hot or soft

A trader who internalizes the order cannot be ambushed by a screaming headline over an in-line core, or soothed by a soft headline over a hot core — the two classic retail mistakes of the release, monthly, for years.

Next: the components themselves — the decomposition the desks run in the first minutes, and the categories that carry the print's meaning.

The base effect, worked

A worked example of Chapter 3's arithmetic, because the ledger and the sheet both use it:

- Suppose last year's same-month MoM was +0.4% — a hot base month - Then this month's print of +0.2% MoM lowers the YoY rate by roughly a fifth of a percentage point — a clean YoY fall on a merely ordinary MoM - Suppose instead the base was 0.0%: the same +0.2% MoM raises YoY by two tenths — an "acceleration" headline on the same ordinary print

Same fresh inflation, opposite YoY headlines, both fully predictable weeks in advance. In the 2023 disinflation this arithmetic produced the era's recurring theater: YoY headlines falling "faster than expected" on scheduled base months, and desks with the calendar shrugging while commentators narrated. The sheet's base-effect table (Chapter 9) exists so you are always on the right side of that theater.

The decimals, taken seriously

Two rounding facts with outsized trading weight:

MoM is quoted to one decimal, and the trading range of ordinary months is narrow — roughly 0.1 to 0.4 in most regimes. A single tenth of a MoM is therefore a large relative move: 0.2 versus 0.3 consensus is a 50% miss on the marginal number, and the machines price it accordingly. The inflation release's surprises look small in absolute terms (a tenth here, a tenth there) and are enormous in the terms that matter — the surprise is calibrated against the range, not the level, and the range is hundredths-wide in normal regimes.

YoY thresholds are narrative cliffs. Crossing round numbers — the era's famous milestones, in each direction — moves coverage, sentiment, and positioning independently of the underlying arithmetic: 8% YoY, 9% YoY on the way up; the fall back through 6, 5, 4, 3 on the way down — each crossing a headline event with follow-through beyond the print's economics. The trader's use is positioning's, not fundamental's: threshold months carry crowded option positioning and outsized first-minute volatility — the calendar knows the crossings in advance, and so does the

street.

Chapter 4: The Components

The core MoM number is an average, and averages conceal. The desks' first minutes of a CPI release are spent not on the number but inside it — decomposing the core print into the categories that produced it, because the market prices categories, not averages. This chapter is the decomposition guide: the categories, their weights and behaviors, and the desk taxonomy that turns a print into a narrative.

The desk taxonomy: four buckets

Modern inflation trading groups the basket into four functional buckets, and the print's meaning is read through them:

1. Food and energy (headline-only). Excluded from core by construction; volatile on supply; politically loud but analytically discounted. The exceptions matter historically — the energy shocks of the era (the 2022 oil complex's surge) fed into other prices over time — but on release morning, these categories are the noise layer, and a headline print driven by them is dismissed by the desks as such.

2. Core goods (goods less food and energy): apparel, used and new vehicles, furniture, appliances, recreation goods. The goods categories are globally traded and supply-chain sensitive — the 2021-2022 era's goods inflation was, in the common account, substantially a supply-chain phenomenon (shipping costs, semiconductor shortages, demand rotation into goods), and its 2023 disinflation the unwinding of the same. Goods prices can deflation-spike and disinflation-spike; the desks read them as the flexible part of core, and their prints carry information about supply normalization more than demand strength.

3. Shelter. The largest component — roughly a third of the total index and the largest share of core — and the lagging one. Its measurement (Chapter 5) and its lag structure deserve their own chapter; in the taxonomy, shelter is the slow component: mechanically steady in the short run, trend-defining over quarters, and the source of the market's most persistent CPI forecasting problem.

4. Core services less shelter — the famous "supercore." Services excluding shelter and energy: medical services, transportation services, recreation services, and the rest. This bucket is the Fed's modern obsession: services prices are labor-intensive (wages are the dominant cost), so sustained supercore inflation is the wage-price channel made visible — the Fed's public commentary through 2022-2024 repeatedly framed the inflation fight's "last mile" as a services-ex-shelter problem. A hot core print driven by supercore is the most hawkish composition a release can offer; a hot core print driven by a one-off goods category is the least.

The decomposition protocol

The release-day exercise, ten minutes with the BLS tables (or the fast summaries of them):

ABOUT THE AUTHOR

Aiman Sterling writes the Impact News series -- a trader's field guides to the market-moving events on the calendar, from the Jackson Hole symposium to the monthly data releases. The series' method is preparation over prediction: written expectations, regime maps, and honest ledgers, applied event by event.

The Impact News books follow a shared framework and stand alone in depth: each event gets its own mechanics, its own history, and its own action plan, with no repeated filler between volumes.

Aiman writes from the practitioner's side of the terminal: systems that survive the first minute, risk rules that survive the first loss, and the discipline of grading every event against the expectations written before it.