SAMPLE

The PCE Price Index

The Fed's Own Inflation Gauge — the Quiet Print That Governs Policy

Aiman Sterling

Free sample: Introduction & Chapter 1

putrabooks.github.io

Buy the full book:

Ko-fi

Payhip

Chapter 1: Introduction

Of all the events in this series, this one has the strangest profile: the price index that formally governs Federal Reserve policy — the gauge against which the 2% target is defined, the number projected in every SEP, the measure the Chair cites when describing progress — is also, on release mornings, among the quietest numbers on the calendar.

The Personal Consumption Expenditures price index moves markets less than the CPI report that arrives weeks earlier. Its release mornings are, most months, non-events: the print lands within the pre-announced range, the tape absorbs it in minutes, and the market moves on to other business. A trader who brought the CPI book's adrenaline to a PCE morning would spend twelve quiet mornings a year wondering where the event went.

The quiet is the puzzle, and the puzzle is the book. The PCE index is not quiet because it is unimportant — it is the most policy-important inflation number that exists. It is quiet because it is forecastable: by the time it prints, the CPI and PPI reports that feed its construction have already published most of its content, and the market's desks have estimated the number within a tenth or two before the Bureau of Economic Analysis ever releases it. The release confirms; it rarely reveals.

This creates two trading realities that this book takes in full:

First, the release itself is a different kind of event — a confirmation event rather than an information event. Its surprises are smaller, its reactions conditional, its playbook structured around what a confirmation can and cannot do. Chapter 9 builds that playbook: when a PCE surprise is tradeable (the genuine deviation from the CPI/PPI-implied estimate — the number surprising relative to what was knowable, not relative to zero), and when the morning is noise dressed in the costume of an event.

Second — and this is the larger half of the book — the gauge's value lives between releases. The Fed's target is on PCE inflation. The SEP's projections are core PCE. The policy path — the dots of Book 2, the reaction function every event in this series trades against — is anchored to this index's trend. A trader who reads the PCE's construction, its wedge against CPI, and its monthly path holds a regime anchor: the standing answer to "where is policy, and where is it going" that every other event in the series prices. Chapter 10 builds the between-events trade: the regime position held across months, marked against the gauge's prints, rather than the release-morning binary.

The series' framework transfers completely and nothing is re-taught: the expectation gap (measured here against the derived estimate, the honest consensus), the regime maps (a confirmation event's reaction structure is regime-dependent like any other), the composition method (PCE has components too, with different weights than CPI's — the healthcare-heavy coverage difference of Chapter 3), and the ledger (twelve quiet

mornings that compound into the era's file).

What this book adds to the series is the target-gauge discipline — the method of reading policy's actual yardstick rather than the market's louder proxy. Traders quote CPI; the Fed targets PCE; the wedge between them (Chapter 4) is a standing source of misreads, and the trader who tracks the gauge the policy is actually anchored to holds an information edge that costs fifteen minutes a month to maintain.

The book's order is the series' learning hierarchy: what the index is (Chapter 2) before why the Fed prefers it (Chapter 3) before the wedge against CPI (Chapter 4) before the target arithmetic (Chapter 5) before the release's forecastability mechanics (Chapter 6) before the policy linkage (Chapter 7) before the era's record (Chapter 8) before the two trades — the confirmation playbook (Chapter 9) and the regime anchor (Chapter 10) — before the sheet (Chapter 11) and the plan (Chapter 12).

The quietest event in the series is, in the long game, the most governing one. Let's take it properly.

What this book assumes, and what it adds

Two frameworks carried from the series, assumed without re-teaching: the expectation-gap (reality priced against what was actually knowable, the discipline Book 1 established) and the two-wave event structure (machines, then humans, from Book 2 onward). The composition method (Book 5's component discipline) is extended here — the PCE's categories differ from the CPI's in weight and scope, and the extension is a lesson in the method's portability rather than a re-teaching.

What this book adds is the series' final structural layer: the denominator discipline. Every rate-path question — every dot, every cut's timing, every "sustained convergence" claim in the Chair's sentences — is arithmetic on this gauge. The trader who holds the denominator (the gauge's construction, its wedge against the banner, its trend's pace) reads every other event in the calendar with a correction the loud numbers do not carry: the market quotes the CPI; the policy moves on the PCE; and the gap between the two closes on the policy's schedule — a schedule this book teaches you to draw.

A note for readers entering the series here: this book stands alone at the cost of one prerequisite — comfort with the idea that releases are graded against derived expectations rather than zero, established above and applied throughout. With that single framework in hand, the chapters build in order, and by Chapter 10 the quiet gauge is the center of a complete trading system: the release's playbook, the regime's anchor, and the position the whole calendar grades.

The chapters ahead, in order: the report and its machinery; the preference's reasons; the wedge; the target's arithmetic; the forecastable release and its build; the SEP linkage; the case studies; the release's trading playbook; the regime anchor and the between-events trade; the risks and the sheet; and the action plan. A morning-of quick-start sits in the back matter, as in every book of the series. The whole of it, one

gauge — held to the standard of the loudest numbers in the market, because that is what the policy holds it to.

Chapter 2: The Report Behind the Index

The PCE price index does not get its own release. It is published inside the Bureau of Economic Analysis's Personal Income and Outlays report — the monthly account of what American households earned and spent — and this embedding is the first fact of the event's character: the price index arrives as one line of a broader report, in a release whose headline is often the savings rate or the income month, not the inflation print.

The report, briefly

The Personal Income and Outlays release covers, for the reference month: personal income (wages, transfers, property income), disposable personal income, personal consumption expenditures (the spending itself, in current dollars), the savings rate, and — the trader's business — the PCE price index, headline and core, MoM and YoY. The release lands around the end of the following month, at 8:30 a.m. Eastern, roughly four weeks after the month closes and roughly two weeks after the CPI that partially feeds it.

The report is the national accounts' monthly pulse — the same accounting framework that produces the quarterly GDP, which is why the PCE price index (not CPI) is the inflation line inside the GDP report's own price measures. The index is, structurally, the accounts' gauge: it deflates the spending the BEA measures, and its construction follows the accounts' logic.

The index's construction

The PCE price index measures the prices of the consumption basket that the national accounts observe — and the differences from the CPI's construction begin at the definition of the basket:

Coverage. The CPI covers out-of-pocket spending by urban consumers. The PCE covers all consumption spending on behalf of that same population — including categories households never see on a receipt: employer-paid health insurance premiums, government healthcare programs' spending (Medicare and Medicaid's reimbursement rates flow into the index directly), nonprofit institutions' consumption on behalf of households. The healthcare weight in PCE is roughly double CPI's — the single largest structural difference between the gauges.

Weights. The CPI's weights come from household expenditure surveys; the PCE's come from the business and government side of the national accounts — reported sales and revenues, not consumer recall. The measurement source is, statistically, better grounded; the household survey has well-documented frictions (underreporting at the tails, the response-bias problem), and the accounts' source data is the economy's own

ledger.

The formula. The CPI is (in its headline form) a fixed-basket Laspeyres-type index; the PCE uses the Fisher ideal formula — chaining the index with both base-period and current-period weights, month by month. The Fisher formula captures substitution: as relative prices shift, spending shifts toward what got cheaper, and the aggregate price rise is measured accordingly. The fixed-basket CPI, by construction, ignores substitution — a standing upward bias in CPI relative to PCE.

The net result. A broader, better-sourced, substitution-tolerant index — which is why the Fed, asked in the modern era to name its gauge, names this one. Chapter 3 takes the preference in full; this chapter's work is the anatomy.

The categories, for the composition read

The index's component structure follows the accounts' categories: goods (durables and nondurables) versus services; the healthcare complex at its double-CPI weight; housing (with the same OER-style treatment of owned housing as CPI, but at a lower weight — shelter's share in PCE is roughly half its CPI share, the second-largest structural difference); food and energy at their smaller PCE weights; financial services (a category the CPI barely touches — portfolio management, banking services imputations — which the PCE measures and which adds its own volatility).

For the composition discipline of Book 5: the PCE component read is the same method with different weights — healthcare's doubled weight, shelter's halved weight, financial services' presence — and the desk's PCE estimates are built component-by-component from the CPI and PPI detail, the forecastability mechanism of Chapter 6.

The release's furniture

The morning's stack, in the trader's order: core PCE MoM (the tradeable number), core PCE YoY (the target arithmetic), headline MoM/YoY (the banner), and the income/spending lines (the report's own headline, occasionally the day's actual story — a spending collapse or a savings-rate print moving the tape while the inflation lines are in line). The trader's standing note: this release carries two stacks — the inflation stack and the activity stack — and the quietest PCE mornings are the ones where both print in line and the market's attention was never there to begin with.

The quarterly meeting

Four times a year, the same BEA machinery prints the quarter's accounts in the GDP report, and the PCE price index's quarterly readings are the series' anchor points: the quarterly core PCE — annualized — is the number the Fed's projections table references, the number the Chair cites in testimony, the measure of the target's progress at the policy cadence. The monthly prints are the quarterly's inputs, and the trader's target-gauge file (Chapter 10) is built at the monthly tempo and read at the

quarterly one.

The report's market mechanics

Three practical notes on the release's machinery, worth the trader's five minutes:

The revisions regime. Unlike the CPI (essentially never revised — Book 5, Chapter 2), the BEA's accounts revise: the monthly income and outlays figures are revised in subsequent releases, and the annual updates re-run whole years. The price index itself is more stable than the spending lines, but the trend read of past months can shift at the revision cycle — a quiet standing source of "the history changed" surprises the sheet's monthly mark should note. The trader's rule: the PCE's past is a draft; the CPI's past is final; grade the present against each accordingly.

The release calendar's logic. The roughly four-week lag exists because the accounts must wait for their source data — the monthly business surveys, the spending panels — and the CPI's two-week head start is the nowcast machine's origin: the same month measured by the consumer-facing index first, the accounts' deflator later. The calendar's consequence for the family map: the PCE print is the inflation month's final reading, the month's account closed — and the next CPI cycle begins before the previous PCE lands, the family's pipeline overlapping itself continuously.

The data access. The full report — the tables, the categories, the revisions — publishes freely at 8:30 on release day, and the desks' algos parse it in seconds. The composition read's practical difference from CPI: the PCE detail is accounts detail (the categories' construction partially from the PPI-side and program data), so the component pull is a heavier lift — most traders rely on the desks' summary tables (the services/goods split, the PCE's "supercore" equivalent) rather than the raw accounts. The practical note: know which of the two component frames you hold — the CPI's categories (which you can build by hand) versus the PCE's (which you read from the summaries) — and grade each accordingly.

The revision calendar, on one page

The gauge's history is not fixed, and its revision schedule is part of the instrument. Monthly PCE arrives with the personal income report at month-plus-four-weeks, but the numbers sit inside the national accounts' three-layer vintage system: the monthly print is preliminary, quarterly updates revise it within the GDP vintages, and each fall the annual update reworks several years of history as source data — tax records, surveys, healthcare claims — replaces the earlier estimates. The September annual updates carry the heaviest weights: benchmark-level revisions to PCE inflation levels have arrived in them repeatedly, moving the core history by tenths of a percentage point that the SEP had already metabolized the old numbers.

The trader's calendar note is short. Monthly morning: treat the print as an estimate with a revision tail the sheet should expect, not an exception. Quarterly vintage weeks: expect the PCE table to inherit GDP's revisions — a change in consumer spending history arrives with the quarterly GDP print, and the PCE revision is in the

same release. September: mark it as the annual update week; before it, the sheet's baseline momentum numbers carry revision risk, and after it, they are cleaner than at any other point in the year. The gauge's reliability is seasonal, and the season is on the calendar.

The accounts behind the index

The PCE price index's deepest structural difference from the CPI is its source: it is not a survey of prices but a by-product of the national accounts — the BEA's estimate of what households actually spent, measured from the demand side (business receipts, program outlays, the accounts' source data) rather than from the CPI's retail-side collection. The consequences run through every chapter of this book: the weights update monthly as the accounts re-measure the consumption mix (the CPI's weights update on a slower schedule); the coverage includes categories the household never sees a bill for (employer-paid premiums, program-paid care, the imputed pieces of financial services); and the revision culture differs — the accounts' source data arrives in vintages, and the monthly prints are, in principle, subject to the accounts' revision cycle, a fact the desks' grading culture treats as noise-level but the purist's sheet notes honestly. The index, in one sentence: the economy's consumption, priced from the accounts' own arithmetic — a fundamentally different instrument from the retail-basket survey, and the reason the two gauges can never quite agree.

One report, then, carrying the economy's consumption arithmetic — and every property this book has examined flows from that single fact. Hold it in mind through the chapters ahead.

Chapter 3: Why the Fed Prefers PCE

The Federal Reserve's price-stability mandate is, in the modern era, defined on this index: the 2% target is PCE inflation, and the preference is official, documented, and repeatedly stated. The why is the trader's business — each reason for the preference is a fact about how policy reads the economy, and the reaction functions of every event in this series are downstream of that reading.

The stated reasons

The Fed's own explanation of the preference (published in its FAQ and cited in Chair after Chair's communication) centers on three arguments, each mapping to a construction fact from Chapter 2:

Coverage. The PCE measures all consumption on behalf of households — the employer-paid premiums, the government programs, the full healthcare complex. The economy's consumption price level is what the mandate plausibly means, and the PCE measures that level; the CPI measures the out-of-pocket slice. A policymaker defining "price stability" for the economy as a whole rationally prefers the accounts' net.

Weights and source data. The PCE's weights come from business-reported revenues —

the economy's ledger — rather than household surveys with their known frictions. The measurement is better grounded, and the Fed's confidence in the index's path is correspondingly higher.

The formula. The Fisher chain captures substitution; the fixed-basket CPI does not. A target defined on a substitution-biased index would be penalized by the bias itself — the gauge rising mechanically relative to any true cost-of-living concept. The Fisher formula removes that mechanical drift.

The consistency argument. The PCE index is the accounts' deflator — consistent with the GDP framework in which the Fed's other mandates (output, employment) are measured. The three measures — activity, labor, prices — live in one accounting house, and the comparisons across mandates are internally consistent.

The history, briefly

The preference predates the target. The CPI served as the practical reference for decades (it is the index of indexation and public life — COLAs, tax brackets, contracts), but the Fed's research and FOMC deliberations moved steadily toward the accounts' gauge through the 1990s, and the formal 2% target — announced January 2012 — was defined from the start on the PCE. The 2020 framework review reaffirmed it, with the average-inflation overlay defined on the same index. The trader's compressed version: the CPI is the country's gauge; the PCE is the Fed's; the target is on the Fed's.

What the preference means for trading

Three consequences carry trading weight:

The reaction function reads PCE. When the Fed says "inflation is moving sustainably toward 2%," the number in that sentence is core PCE — and the market's CPI-driven narratives regularly run ahead of or behind the gauge the policy is actually anchored to. The gap is the standing misread: CPI months that look like breakthroughs while the PCE trend (healthcare-heavy, shelter-light) has not moved; CPI months that look like stalls while the PCE confirms. Chapter 4's wedge is the standing quantification of that gap, and the trader who tracks it holds the correction to the loudest narratives.

The projections are core PCE. The SEP's inflation column is core PCE — the dots of Book 2 sit on top of this index's projected path. The quarterly meeting of the SEP and the PCE's quarterly anchor (Chapter 2) is the policy calendar's center of gravity: the desks projecting the SEP from the CPI/PPI pipeline (Book 5's Chapter 12) are, in substance, running the PCE nowcast — which is why the PCE release rarely surprises and why the SEP quarters are the PCE's loudest season (Chapter 7).

The Chair's progress claims have a denominator. When the press conference (Book 3) says "progress," the word is measured on core PCE — and the trader grading the claim checks the same trend, not the CPI banner. The target-gauge discipline, in one line: quote the Fed's number to the Fed; quote the market's number to the market; never

confuse which room you are in.

The preference, taken as read, is the book's foundation. Next: the wedge — the standing difference between the two gauges, and the misreads it produces.

The preference in practice: the mandate's other gauges

The preference has a practical penumbra worth filing: the Fed watches a family of gauges, and the trader's target-gauge discipline should know which numbers hold which roles.

The headline PCE — the mandate's literal text, the political number (Chapter 5).

The core PCE — the policy's working reference, the SEP's variable, the gauge this book is about.

The trimmed and median variants — the Cleveland Fed's trim/median constructions applied to the PCE family exist alongside the CPI versions, and the Fed's research commentary leans on them when the composition is noisy — the same satellites of Book 5's map, in the accounts' house.

The market-based and non-market splits — the BEA's distinction between market transactions and imputed/program categories: a research-level cut, occasionally quoted when the healthcare-program dynamics move the headline in ways the market categories did not — the scope effect of Chapter 4, visible in one table.

The long-run expectation measures — the surveys and market measures of expected inflation (the expectations family the Fed monitors alongside the realized gauges): not price indexes at all, but part of the mandate's dashboard, and the trader's regime file carries them as the target's forward column alongside the gauge's realized one.

The discipline's payoff, stated once more: the Fed's reaction function reads this family with a known hierarchy — core PCE as the working number, the satellites as context, the expectations as the forward check — and the trader who holds the same hierarchy grades every inflation print, of every gauge, in the same order the policy does.

A caution on the preference

The preference is about measurement, not about a commitment to mechanical policy on the gauge's prints. The Fed does not target the MoM; it does not react to tenths mechanically; the framework (Chapter 8) interprets the gauge, and the interpretation is the policy. The trader's version of the caution: the target-gauge discipline is about holding the correct denominator for the policy's arithmetic — not about pretending the arithmetic runs itself. The denominator corrects the commentary's errors; the reaction function still needs the event books (the whole rest of the series) to read.

The preference, tested

The preference is not abstract — it is visible in the record's sharpest moments, and

testing it against the banner is a standing exercise for the anchor page. In the 2021-22 surge, the commentary era ran largely on the CPI's banner readings (the 9% headline prints were the public's inflation), while the policy's own assessments (the SEPs' projections, the framework's average-inflation arithmetic, the eventual tightening's calibration) ran on the PCE frame — core PCE peaking in the mid-6s YoY against core CPI in the 6s-to-7s, a full point of banner-versus-gauge gap at the cycle's worst. The desks that graded the surge on the gauge's arithmetic were, by the wedge's structure, always reading "cooler" than the headlines — and were, by the policy's own schedule, reading the number that mattered. The preference's test repeats every cycle: the banner leads the narrative, the gauge leads the policy, and the wedge's sign (which gauge is hotter, and why) is a standing line on the anchor page — the market's most quotable number and its most consequential one, separated by a measurable, trackable, tradeable gap.

Chapter 4: The CPI–PCE Wedge

The two consumer inflation gauges measure, roughly, the same economy — and print, persistently, different numbers. Over the modern record, the CPI has typically run a few tenths of a percentage point above the PCE at the annual horizon, and the difference — the wedge — is one of the most stable facts in macro measurement. The trader who knows the wedge's anatomy holds a standing correction to half the inflation commentary published.

The wedge's sources, ranked

Chapter 2 and 3's construction differences, quantified in the BEA's own published decompositions of the gauges:

The weight effect (the largest). The categories where the gauges' weights differ most — shelter (roughly double in CPI what it is in PCE) and healthcare (roughly double in PCE what it is in CPI) — move the wedge whenever their inflation rates diverge from the basket's average. In the era of shelter's pipeline lag (Book 5, Chapter 5), the wedge widened on schedule: CPI's shelter weight held core CPI up while PCE's lighter shelter line let core PCE fall first. The disinflation's "PCE arrived at target before CPI" arithmetic was substantially this weight effect, visible in advance to anyone holding the decomposition.

The formula effect (the persistent). The Fisher chain's substitution tolerance removes the CPI's mechanical upward drift — a standing few tenths at the annual horizon, the wedge's floor. The trader's compressed version: even with identical category inflation, the CPI prints hotter, because the fixed basket ages while the Fisher basket adapts.

The scope effect (the structural). The categories PCE covers and CPI does not — the employer-paid premiums, the program reimbursements, the financial services imputations — carry their own inflation rates, largely invisible in consumer-facing commentary. The healthcare complex's reimbursement dynamics (the government-set rates) make the PCE's healthcare line a policy-administered number

in significant part — smoother than market prices, and a stabilizer of the gauge's trend.

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.