SAMPLE

The Consumer Confidence & Sentiment Reports

Mood Is Not Money — the Sentiment Surveys, Their Inflation Anchor, and the Gap That I

Aiman Sterling

Free sample: Introduction & Chapter 1

putrabooks.github.io

Buy the full book:

Ko-fi

Payhip

Chapter 1: Introduction

This volume takes the consumer-mind pillar the series has not yet touched: the Consumer Confidence & Sentiment Reports — the Conference Board's confidence survey and the university sentiment survey, the monthly mood prints the spending watch and the policy reaction both run on.

The mood surveys earn their own volume because they measure what the spending books cannot: the consumer's state of mind — the expectations, the worries, the inflation anticipations — the soft data that leads, lags, and occasionally contradicts the hard spending prints the Retail Sales volume covered. That two-survey structure changes the trading physics: the two-print comparison (the confidence board's read and the sentiment survey's read, the monthly cross-check), the inflation-expectations watch (the components the policy reaction cites by name), the sentiment-versus-spending divergence (the mood's honest lead-lag the era's record documents), and the survey noise (the sampling error the monthly print carries). The trader who holds the mood discipline reads the confidence prints as expectations thermometers; the trader who reads them as spending forecasts holds the soft-data tuition.

This book is the twentieth panel of the Impact News series. The differentiation discipline, stated once: the wing's framework is applied here at the mood cadence, not re-taught — the reader new to the series should hold the data wing's earlier volumes (the retail sales and the inflation books especially) first. This volume's own additions: the two-survey vocabulary and release machinery (Chapter 2), the monthly release structure (Chapter 3), the confidence channel map (Chapter 4), the expectation and live-read forms (Chapters 5-6), the survey-noise discipline (Chapter 7), the sentiment-to-spending machinery (Chapter 8), the era's cases — hedged, from public coverage (Chapter 9), the framework doing its work (Chapter 10), the playbook (Chapter 11), the risk sheet and its plan (Chapter 12), and the closing plan (Chapter 13).

What this book is not

The series' standing rule, held at the mood volume: no fabricated episodes, no invented levels, no specific month claimed as a lesson. The era's mood history is built on widely reported public coverage, hedged throughout — the structures the lessons, never the specific prints. This volume takes no position on any current mood reading and predicts no sentiment path; it teaches the survey structure the era's record documents, in the series' hedged form.

The reading order

For the reader short on time: Chapters 3, 4, and 8 — the release structure, the

channel map, and the sentiment-to-spending machinery — are the working core; Chapter 11's playbook works them into routine; the whole book is the durable reference. Both beat the standard read (the mood print traded as the spending forecast, the single survey traded as the consumer's whole verdict), which is the read this book retires — and the read that pays the soft-data tuition.

The wing's softness, stated first

The introduction's first honesty is the wing's defining softness: the confidence and sentiment reports are the calendar's only major releases that measure psychology rather than behavior, and everything the book teaches is the discipline of trading a feeling honestly. The soft series' paradoxes are the wing's own puzzles: the mood at its worst in the survey's history (2022) during one of its strongest spending eras; the mood's political plumbing (the partisan gap) repricing the same economy by tens of points at an election; the mood's strongest single driver (the gasoline pump) being a small share of the household's budget with an outsized share of its attention. The soft series cannot be traded the way the hard series is traded — it must be decomposed, diverged, and confirmed — and the trader who finishes this book will hold the binder's last and subtlest discipline: the measurement of the mood as a component-level instrument whose aggregate number is the least of its contents.

The puzzles, stated

The puzzles the chapters answer, stated plainly: why the mood can collapse while the spending holds (the sentiment-spending divergence); why the two surveys disagree some months (the methodology gap); why the inflation-expectations component moves policy more than the headline (the reaction function's watch); why the sentiment print follows the gas pump more than the payroll (the pocketbook anchor); and why the mood's crash is sometimes the bullish trade (the expectations-reversal structure). Five puzzles, five chapters' answers, and the mood channel's whole honest structure assembled from them.

The mood pulse's one-line identity

The identity the chapters build, stated in one line: the headline discounted by the sampling, graded by the two-thermometer trend, repriced by the components, and sized by the noise — the monthly mood event the market samples, and the consumer pillar's soft-data structure. The mood teacher earns the whole series' complement: the soft thermometers the hard data's own companion, the expectations the policy's named input, and the divergence the desk's honest flag.

Chapter 2: What the Confidence & Sentiment Reports Are

The Consumer Confidence & Sentiment Reports, as this book uses the term, is the

monthly market-moving class of consumer-mood prints: the Conference Board's confidence survey, the university sentiment survey, and the derived indicators — the events that measure the consumer's expectations, and the mood pulse the spending watch runs on.

The two-survey forms

The reports' honest taxonomy: the confidence survey (the board's monthly read — the present situation the expectations index the headline carries), the sentiment survey (the university's twice-monthly print — the preliminary the final, the current-conditions and expectations components), and the derived indicators (the expectation components the inflation anticipations the sub-indices the surveys' own math produces — the policy-watched layers the headline never shows). The forms' honest differences: the board's survey the labor-market anchor (the jobs-plentiful the jobs-hard-to-get components the employment watch cites), the university's survey the inflation anchor (the year-ahead price expectations the policy reaction watches by name), and the derived layers the analysis runs — two surveys, two anchors, and the two-print comparison the monthly read.

The mood vocabulary

The surveys' machinery, the lexicon's foundation: the present-versus-expectations structure (the current conditions the future hopes — the two horizons every mood survey separates, the lead-lag the components carry), the pocketbook structures (the gas prices the grocery bills the personal-finance reads — the pocketbook anchors the mood follows, the macro-mood gap the vocabulary's honest grade), and the methodology structures (the survey sizes the sampling the panel differences — the two surveys' honest differences the monthly comparison discounts). The vocabulary's honest use: the mood surveys are expectations measures (the consumer's anticipations the surveys sample), and the methodology gap the two-print read's own standing caveat — the soft data the whole read's foundation, in the mood edition's standing form.

The event, then

A monthly class of consumer-mood prints — confidence, sentiment, the inflation-expectation components — carrying the consumer economy's soft-data pulse: the mood thermometers the spending watch runs on, the policy reaction's named inputs, and the expectations the hard data only confirms. The series' framework runs it whole; the mood cadence's own specifications this book adds.

The mood's own production

The monthly event's production structure, the machinery's engineering: the survey process (the monthly questionnaires the household panels the responses the surveys sample — the mood's own supply chain), the index process (the responses'

aggregation the components' construction the indexes' own math — the derived layers the methodology builds), and the publication engineering (the monthly clocks the preliminary-final cadence the seasonal notes — the surveys' own honest machinery). The honest read: the mood index is a sampled measure (the household panels the confidence intervals the survey noise the prints carry), and the sampling machinery the noise's own anatomy — the production chain the caveats' own source, in the mood edition's standing form.

The two surveys' architectures, compared honestly

The wing's two main instruments — the Conference Board's Consumer Confidence and the University of Michigan's consumer sentiment — are built differently enough that their disagreements are structural rather than statistical, and the honest comparison is the reader's first discipline. The Conference Board's survey asks about present conditions and expectations in roughly equal weight; the Michigan survey leans on the expectations side and asks its famous price questions (the year-ahead and five-to-ten-year inflation expectations the Federal Reserve cites by name). The reference periods differ: the Conference Board's monthly fieldwork and the Michigan's rotating panel run on different calendars, so the two surveys can print genuinely different months — the 2022 files' record divergence was partly calendar, partly construction, and the desk that knows which is which reads the disagreement as information rather than error. The construction details the file rewards: the Michigan survey's panel structure repeats respondents (the mood's month-to-month change measured partly on the same households), while the Conference Board's monthly sample is fresh — the trend-versus-level question the noise chapter grades.

The comparison's discipline for the sheet: the two surveys are one instrument's two strings (the trend chapter's cross-check), and the divergence file is graded by its cause — the calendar (the surveys measuring different weeks of a moving month), the construction (the expectations-heavy versus the present-conditions-heavy), or the economy itself (the genuine mood signal the convergence read waits for). The 2022 record case ran a genuine economic divergence (the price questions collapsing while the labor questions held), and the desk that decomposed the two surveys' components held the read the aggregate disagreement obscured. The honest comparison, closed: never average the two surveys into one "consumer mood" — the average of two different instruments is a number that describes neither, and the divergence itself is the information.

The weekly anatomy

The mood lexicon's honest depth: the headline layer (the composite indexes the reflex trades — the mood's first sample), the component layer (the present the expectations the inflation anticipations — the honest read the headline buries), and the comparison layer (the two surveys' prints the methodology holds — the two-thermometer read the desk runs). The honest anatomy: the three layers the three clocks (the reflex the seconds, the components the repricing, the convergence the signal), and the layer's

identification the read's own horizon — the mood anatomy, in the series' standing file.

The class's last line: the confidence the labor anchor, the sentiment the inflation anchor, the components the second life — and the two-thermometer comparison the mood read's own whole first structure.

The indexes' own math

The derived machinery's honest depth: the composite construction (the components' aggregation the headline indexes the methodology builds — the present the expectations the weighted composite), the diffusion structure (the responses' balance the improving-minus-worsening the survey math runs — the direction the diffusion reads), and the component decomposition (the sub-indices the layers the release's own tables carry — the inflation anticipations the jobs reads the big-ticket intentions). The honest math: the headline is the built layer (the composite the components' own weighted echo), and the components the honest depth the headline aggregates — the index math, at the mood edition's standing form.

Chapter 3: The Release Structure

The mood reports' structure is a two-clock calendar: the board's monthly print and the university's twice-monthly cadence — the survey clocks the desk schedules around. This chapter builds the structure.

The two clocks

The releases' standing anatomy: the board's clock (the monthly confidence print the month-end window carries — the last-Tuesday convention the coverage grades), the university's clock (the twice-monthly sentiment cadence — the preliminary mid-month the final end-month, the two-print rhythm the month's mood runs on), and the calendar's interplay (the two surveys' dates the month's own sequencing — the preliminary's lead the confidence's confirmation the final's settlement). The honest anatomy: the mood calendar runs two clocks (the survey dates the month's own mood tempo), and the desk's monthly file the two-print sequencing the whole schedule builds around.

The print's contents

The releases' honest contents: the headline prints (the confidence index the sentiment index — the first prints the consensus prices), the component prints (the present-situation the expectations the current-conditions the personal-finance — the layers the honest analysis runs), and the policy-watched prints (the inflation-expectation components the year-ahead anticipations — the sub-indices the reaction function cites). The honest contents: the headline the reflex trades, the components the analysis runs, and the inflation expectations the policy watch holds — the three layers the single print's whole anatomy.

The survey's honest machinery

The mood prints' honest caveats: the sampling noise (the household panels the sample sizes the confidence intervals — the survey error the monthly print carries), the methodology differences (the board's questions the university's questions the panel designs — the two surveys' honest divergence the monthly comparison discounts), and the revision structures (the preliminary's update the final's correction — the two-print cadence the first print's own discount). The honest machinery: the mood print is a survey (the sampled noise the monthly print carries), and the caveats' discipline the whole mood read's own foundation — the noise's honest anatomy, in the series' standing form.

The preliminary-final mechanics, Michigan edition

The Michigan survey's two-print structure is the wing's only preliminary-final release, and the mechanics deserve their own section because the two prints are genuinely different instruments. The preliminary print arrives mid-month on a partial sample — the month's first half's fieldwork, the mood's fastest public read — and the final print completes the month's panel with a revision that has historically averaged a modest fraction of the preliminary's move. The structure's trading content: the preliminary print is the surprise-carrier (the market trades the month's first read at reflex speed, the revision's prior priced at the historical band), and the final print is the confirmation-carrier (the trend discipline's data point, the divergence read's settling print — graded at the noise chapter's tempo rather than the reflex's).

The mechanics' discipline for the sheet: the two prints get different trades. The preliminary's surprise is traded at the addressed instrument's reflex (the expectations components' rates trade, the headline's equity fade) with the revision prior attached — the position sized to survive the final's historical band; the final's confirmation is traded at the trend chapter's tempo — the genuine mood signal confirmed only when both prints agree with the two-survey convergence. The most common error in the wing's file is trading the final's revision as news (the revision's size is inside the preliminary's noise band more often than not) and the least-common discipline is holding the preliminary's position through the final with the revision prior's honest size — the two-print structure's whole edge, held by whichever desk respects the difference between a first read and a settled one.

The month's mood sequencing

The two clocks' monthly interplay: the preliminary's lead (the university's mid-month print — the month's first mood sample the desk prices), the confidence's confirmation (the board's month-end print — the second thermometer the month's mood verifies), and the final's settlement (the university's end-month revision — the first print's own correction the cadence delivers). The honest sequencing: the mood month runs sample-verify-settle (the three-print rhythm the desk schedules), and the sequence's position the print's own context — the monthly interplay, in the mood edition's

standing form.

The release's last line: the two clocks the cadence, the components the depth, the sampling the discount — and the month's three-print rhythm the mood read's own honest structure.

The reference-week physics, mood edition

The mood surveys' timing carries a physics the data wings' releases do not: the survey asks about feelings, and the feelings' reference window is the fieldwork's own weeks — which means the print is a transcript of a specific fortnight, not a smoothed month. The Michigan survey's preliminary print (the month's mid-month fieldwork) captures the month's first half's headlines and pump prices; the final print revises toward the full month; and the Conference Board's fieldwork window can bracket entirely different headlines than the Michigan's — the 2022 files printed survey disagreements that were, partly, two different weeks of the same inflation summer. The reference-week discipline: before grading any mood print, check the fieldwork window against the calendar — a survey fielded through an FOMC week, a debt-ceiling deadline, a market crash, or a gasoline spike carries that week's fingerprints, and the print's move is the week's transcript (recoverable at headline speed) rather than the household's structural change (persistent at pocketbook speed).

The trading translation of the physics: the mood print's surprises caused by fieldwork-week news are fade candidates at the divergence chapter's odds (the August 2011 file's V — the headline week's transcript retracing at headline speed), and the moves caused by pocketbook changes are trend candidates at the spending machinery's odds. The desk's pre-release homework is a calendar check, free in minutes: what happened during this print's fieldwork that the household could not have missed? The answer sets the print's prior — the transcript's expected size, the fade's odds — before the release morning grades the surprise. The one-line entry: the mood survey is a fortnight's diary, and the desk that knows which fortnight holds the diary before reading it.

The components' release anatomy

The print's honest layering, the desk's standing map: the present-conditions layer (the current reads the jobs the finances — the now-conditions the release's first table carries), the expectations layer (the forward reads the six-month the year-ahead — the anticipations the release's second table carries), and the buying-plans layer (the intentions the autos the homes the durables — the big-ticket forward reads the release's own tables). The honest anatomy: the mood release is a layered document (the now the next the intentions — the three horizons the honest parse separates), and the layer's identification the read's own horizon — the components' anatomy, in the weekly-vs-monthly edition's standing form.

Chapter 4: The Confidence Channel Map

The mood reports' channel map runs the consumer economy's three routes: the spending channel, the inflation-expectations channel, and the policy channel. This chapter builds the map.

The spending channel

The first channel, the mood's home route: the expectations structure (the mood the spending leads-or-lags — the consumer's state of mind the retail complex prices), the big-ticket structure (the confidence the durable purchases — the mood's honest tightest link the autos the appliances the discretionary complex the era documents), and the cross-check structure (the mood versus the retail sales — the soft lead the hard confirmation the divergence the desk's honest flag). The honest read: the mood is the spending's soft leading edge (the anticipations the purchases follow — at the lag the era's record grades), and the big-ticket link the channel's honest center — the spending channel the map's foundation.

The inflation-expectations channel

The second channel, the surveys' policy-critical route: the year-ahead expectations (the consumer's price anticipations — the components the policy reaction cites by name), the long-run expectations (the consumer's structural inflation hopes — the anchored-or-unanchored watch the whole credibility grade runs), and the break-evens' echo (the consumer expectations versus the market's own priced anticipations — the two-thermometer comparison the desk runs). The honest read: the inflation expectations are the surveys' whole second life (the components the headlines bury the policy watch centers), and the unanchoring watch the channel's own tail — the expectations channel the mood's signature route.

The policy channel

The third channel: the reaction-function route (the mood prints the policy communication cites — the confidence the sentiment the components the speeches quote), the labor-watch route (the jobs components the employment read the board's survey carries — the labor-market soft thermometer the confidence print includes), and the narrative route (the mood the consumer story the coverage builds — the confidence print the consumer narrative's own headline). The honest read: the mood surveys are the reaction function's named inputs (the expectations components the chairs' speeches cite), and the policy watch the mood's whole third life — the policy channel, in the mood edition's standing map.

The policy channel, graded from the anchor file

The confidence channel map's policy channel is the wing's bridge to the rates complex, and its grade comes from the anchor file: the mood surveys' policy relevance is not the mood at all — it is the expectations components, the household's inflation forecast that the anchoring thesis and the modern Fed rhetoric cite by name.

The channel's transmission, from the 2021-2023 file: the year-ahead expectation's prints ran through the tightening cycle's rhetoric (the transitory debate's household evidence), the five-to-ten-year expectation's stability was the anchor's own defense (the de-anchoring risk's standing reassurance), and the long-run print's band discipline became the policy conversation's reference — the Fed's leadership quoting the household's long-run expectation as the currency's credibility gauge. The channel's grade: the expectations components move the front end at hard-data speed, the mood's other components move nothing in the policy complex, and the channel's whole weight rides on the price questions the survey's architects wrote decades before the policy relevance arrived.

The channel's discipline: the policy trade keys to the expectations print's address (the front end, the anchor watch) and decomposes the driver before the trade (the pump-price transcript check — the year-ahead expectation's gasoline arithmetic — before pricing the de-anchoring narrative at policy weight). The 2022 file's demonstration: the year-ahead expectation's spike was the pump's transcript at its most visible, and the desk that faded the transitory in the expectation's spike (the pump's reversal prior) held the anchor watch honestly while the tape priced the household's pump memory as the currency's structural repricing. The one-line entry: the policy channel is the survey's price questions wearing the mood's clothing — trade the questions, decompose the answers, and hold the channel's weight at the anchor's true gauge.

The channel file's honest construction

The map's standing construction, the desk's own build: the trend file (the two-survey levels the phases the multi-month direction — the always-current cycle the desk pre-loads), the pocketbook file (the gas the wealth the jobs — the drivers the mood follows the desk pre-maps), and the component file (the inflation expectations the anchoring the break-evens' echo — the repricing layers the desk watches). The construction's honest use: the confidence block is the soft-data sheet the binder carries (the monthly prints the files' compounding), and the prepared desk's edge the accumulated trend file the reflexes' own gap — the monthly build the mood trade runs.

The channels' last line: the spending the soft lead, the expectations the policy repricing, the reaction the named cite — and the components the whole map's honest second layer.

The gasoline price, as the mood's strongest single driver

The channel map's most documented single driver deserves its own entry: the pump price. The file's finding, stable across decades and both surveys: the gasoline price moves the consumer mood with an outsized weight relative to its share of the household budget — a visible, frequent, painful price that the household transacts weekly and remembers vividly, and the mood surveys' print its movement with a correlation the rest of the pocketbook cannot match. The 2022 record-collapse file

was partly a gasoline file (the pump prices of the summer's peak coinciding with the mood's all-time lows, and the mood's partial recovery with the pump's decline); the historical files run the pattern through every energy cycle (the mood's dips at the pump spikes of the 2000s and 2011, the mood's boosts at the oil collapses' cheap-gas quarters), and the desk that reads the mood without the pump's chart is reading half the instrument.

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.