SAMPLE

The Nonfarm Payrolls Report

Two Surveys, One Headline, the Market's Most Explosive Minute — Trading the Jobs Rep

Aiman Sterling

Free sample: Introduction & Chapter 1

putrabooks.github.io

Buy the full book:

Ko-fi

Payhip

Chapter 1: Introduction

Twelve times a year, on a scheduled Friday morning, the market's most explosive minute arrives. At 8:30 a.m. Eastern, the U.S. Bureau of Labor Statistics releases the Employment Situation report — the monthly accounting of American jobs — and within seconds, interest rate futures, the dollar, equity index futures, and gold have all repriced whatever the numbers mean for the Federal Reserve's next decision.

No other recurring event combines this frequency, this scope, and this violence of reaction. Jackson Hole happens once a year. FOMC decisions, eight times. The payrolls report lands monthly, and because the Fed's dual mandate names employment explicitly, every print is read as a direct input to the next policy move. When the report surprises, the market does not get an afternoon to think about it — it gets about five seconds.

This is the fourth book in the Impact News line, and it marks the series' transition from communication events to data events. Books 1 through 3 analyzed words: a speech, a statement, a press conference. This book analyzes numbers — and the machinery of reading them is different in kind. A speech can be diffed, argued with, reinterpreted in real time. A payroll print is a fact that arrives complete, prices into the market at machine speed, and leaves humans to trade the interpretation of a number they cannot change.

Why the jobs report, of all data?

The Federal Reserve's mandate, by statute, covers employment and prices. Of the two, employment is the faster-moving series — inflation changes slowly, but a labor market can turn in a quarter — and the payrolls report is employment's principal measurement. This gives the report a structural role: it is the data event with the most direct, legible line to the Fed's reaction function. A hot inflation print raises questions about which channel matters most; a collapsing labor market report removes the questions.

The market's behavior around the release reflects that directness. The days before a first Friday are noticeably position-protective; options on rate futures and equity indexes price an event premium into the release morning (Book 1's fifth transmission channel, at monthly tempo); and the minutes after 8:30 are, by common desk description, among the most liquid and most violent of the month.

The report in one honest sentence

Before the detail chapters, the newcomer needs one framing sentence, and it is this: the payrolls headline is the number everyone quotes, and the number the market actually trades is a weighted judgment across payrolls, the unemployment rate, wages, hours, and the revisions to all of them. Two of the report's components can

contradict each other — and do, regularly — because they come from two different surveys of two different populations. The trader who knows only the headline is trading one layer of a multi-layer event, which is precisely the layer the machines parse first and therefore the layer with the least human edge.

What you will learn

The chapters ahead cover, in order: how the report is constructed — the two-survey design that explains most of its apparent contradictions (Chapter 2); the payrolls headline and its consensus mechanics (Chapter 3); the unemployment rate and the household-survey quirks that make it fall when jobs disappoint (Chapter 4); the wage and hours data the Fed reads as inflation signal (Chapter 5); the revision process and how to weight it (Chapter 6); the expectation and regime structure — when good news is bad news and why that flips across cycles (Chapter 7); the documented case studies (Chapter 8); the preparation artifact, the jobs-week sheet (Chapter 9); the execution mechanics of the 8:30 release (Chapter 10); the data-event risk architecture (Chapter 11); and the action plan plus the adjacent labor data that surrounds the report — claims, ADP, and the rest of the labor calendar (Chapter 12).

A note on method before the numbers begin: this book describes the report's construction as the Bureau of Labor Statistics publicly documents it, and its historical episodes as they were widely reported at the time. Where a figure or mechanism is described with less than full certainty — sampling error, seasonal adjustment, the estimation methods behind the revisions — the language says so. The employment report is a statistical instrument, not a counting machine, and the traders who respect that fact have a durable advantage over the traders who think the number is a fact rather than an estimate.

Start with the instrument itself: two surveys, both imperfect, welded into one report.

The event's own hierarchy of learning

A final orientation, for readers entering at this book. The material ahead sorts into three layers, and the book teaches them in order because each layer protects the next:

The instrument layer (Chapters 2-6): what the report actually measures, how its estimates are built, and why its components contradict each other by design. Without this layer, every later judgment is built on a number the trader does not understand — a draft mistaken for a fact, a ratio mistaken for a level.

The interpretation layer (Chapters 7-8): how the market reads the report — the consensus machinery, the regime decoder, the reaction map, and the record that produced them. This is where the gap framework of the series meets its most quantified expression: the payrolls expectation is a published number, and the gap is arithmetic.

The execution layer (Chapters 9-13): the sheet, the release-minute protocol, the risk

architecture, the labor-market map, and the plan. This is where interpretation becomes position — and where the series' standing truth is enforced at its strictest: the event that arrives twelve times a year, instantly, with a fat tail, is not the place for improvisation.

Read in order, the layers build; skipped, they fail in the order above — first the misunderstandings, then the misreads, then the losses. The instrument is where we start, as every trade should: with what the number actually is.

Chapter 2: Two Surveys, One Report

The Employment Situation report's apparent contradictions — payrolls up while the unemployment rate rises, jobs booming while the labor force shrinks — are not errors. They are the arithmetic consequence of the report's design: two separate surveys, of two different populations, using two different methods, published together as if they were one measurement. Understanding the design is the difference between reading the report and being confused by it.

The establishment survey

The first survey — formally the Current Employment Statistics program — measures jobs. It surveys businesses and government establishments: a large sample of employers is asked how many people were on the payroll during the pay period that includes the twelfth of the month. From those responses, the Bureau estimates the total change in nonfarm payrolls: the headline number of jobs gained or lost.

The establishment survey is the source of: the payrolls change (the headline), average hourly earnings (wages), average weekly hours, hours worked by industry, and the industry-level breakdowns. Its strengths: a large sample, a stable frame, and direct measurement of the thing businesses know best — who they employed. Its limitations, publicly documented by the Bureau itself:

It counts jobs, not people. Someone working two jobs counts twice; a self-employed person counts not at all. The payrolls number is a measure of positions, not of employment in the human sense.

It is an estimate from a sample. The headline carries sampling error — the Bureau publishes its standard error, and the range around any given print is wider than most market commentary acknowledges. A print of +150,000 is, statistically, not cleanly distinguishable from +100,000; the market trades the number as if it were a fact, but the number is a sample estimate with a confidence interval.

It must be seasonally adjusted. Employment has enormous seasonal patterns — retail in December, construction in summer, education in September — which the Bureau removes with models that are themselves periodically revised. Seasonal adjustment is why January reports behave oddly (a new year of adjustment factors arrives) and why some prints are later re-adjusted beyond recognition.

New firms need modeling. Businesses that opened after the survey frame was drawn cannot be surveyed yet, so the Bureau's estimation includes a procedure for net business births and deaths — widely known as the birth-death model. Its mechanics are publicly documented, its inputs estimated; it is the report's most-cited and least-understood component, and it exists because the alternative (ignoring new firms) would bias the count. In strong expansions it can add jobs the survey missed; in downturns its extrapolations can, notoriously, lag the turning points — an episode examined in Chapter 8's 2020 case.

The household survey

The second survey — the Current Population Survey — measures people. It surveys households: a sample of individuals is asked about their own work status in the week containing the twelfth of the month. From it, the Bureau computes the unemployment rate, the labor force participation rate, and the employment-to-population ratio.

The household survey's key concepts: to be counted as unemployed, a person must have actively looked for work in the reference period; those who stop looking are counted as out of the labor force, not unemployed — which is why the unemployment rate can fall in a weak economy when discouraged workers exit the count. The labor force participation rate — the share of the population working or seeking work — is the honest measure of labor-market attachment that the unemployment rate alone cannot give.

The household survey's limitations mirror its design: a smaller sample than the establishment survey (with correspondingly larger sampling error), its estimates at the mercy of classification at the household level, and the conceptual quirks — unemployment falling for bad reasons (exits) or rising for good ones (participation recovering, more people seeking work).

Why the two surveys disagree — the design of the divergence

With two independent estimates of the labor market, published in the same minute, disagreement is not a malfunction; it is guaranteed at the margin. The classic divergences:

- Payrolls up, unemployment up. The household survey's population — the self-employed, agricultural workers, unpaid family workers, and those with multiple jobs counted once — differs from the establishment frame. A strong payrolls print alongside a rising unemployment rate usually means the labor force grew faster than household-measured employment. - Payrolls down, unemployment flat or down. Household employment held better than the establishment frame measured, or workers exited the labor force. - Wage strength with weak payrolls. Hours and earnings come from the establishment survey's surviving firms; a slowdown that removes low-wage or part-time positions first can raise average wages mechanically.

The market's desks handle these divergences with a settled hierarchy — the subject of the next four chapters, which take the report's components one at a time. But the design lesson comes first, and it never changes: the report is two estimates, not one truth. Every payrolls morning is two sample-based, model-adjusted, revised-later measurements of two different things, published together and priced as one. The trader who internalizes the design reads the divergences as information; the trader who does not reads them as chaos.

Before leaving the machinery, one practical note: the Bureau publishes the full report — both surveys, all detail, the methodology itself — at the scheduled release minute, on its website, free. The fastest, most complete record of the event costs nothing; the paid advantage in this event is never the data, only the preparation to interpret it.

Chapter 3: The Payrolls Headline

The headline number — the change in nonfarm payrolls for the month — is the report's public face: the number in every push notification, the "jobs number" of market commentary. This chapter covers what it measures, how the market's consensus forms around it, and what the headline does and does not tell the prepared trader at 8:30.

What the number is

The headline is the change in the estimated level of nonfarm payroll employment: this month's estimated total jobs minus last month's, after seasonal adjustment, across the private and government sectors. It is quoted as "nonfarm payrolls rose by X thousand" — a flow number, not a level — which is why a "weak" print of +150,000 in a growing economy still means employment rose.

Its scale context, worth memorizing: the U.S. economy needs roughly 75,000 to 150,000 new jobs a month to absorb population growth, depending on the era and labor-force trend — the market's estimate of that break-even has shifted over the decades with demographics and participation. The market does not judge a print against zero; it judges it against the growth-required pace, against the consensus, and against the Fed's reaction function. A +100,000 print is strong in a recession scare and weak in an expansion — the number is never read raw, always in regime context (Chapter 7).

The consensus and its construction

By release morning, the market's expectation is a published consensus: the median of professional forecasts, gathered by the major data aggregators during the survey week. The consensus forms from: the prior month's trend, the high-frequency labor indicators available before the release (claims, scheduled-hiring indicators, sentiment surveys — the adjacent data of Chapter 12), and each forecaster's model. It is a genuinely honest consensus — published, dated, specific — which makes the payrolls

gap the cleanest gap in the series' whole framework: the market's expected number exists as a single figure, and the actual number arrives as a single figure, and the subtraction is arithmetic.

Around the consensus, two refinements matter. First, the whisper range: the dispersion of forecasts around the median. A consensus of +180,000 built from forecasts clustered tightly between +170,000 and +190,000 is a different event than the same median built from forecasts scattered between 0 and +350,000 — tight dispersion means small surprises have outsized reactions, because fewer participants were prepared for anything else. Second, the post-release revision expectation: sophisticated desks forecast the revision pattern alongside the headline (Chapter 6), because a strong headline with heavy downward revisions to prior months is a weaker print than a modest headline with strong upward revisions.

What the headline does tell you

At its own layer, honestly: the direction and rough magnitude of the month's job creation, at sample precision. In normal months, in normal regimes, the headline is the event — the market's first reaction is the arithmetic gap between print and consensus, transmitted through the rate-path channel (a hot print raising rate expectations: dollar up, yields up, equities down in a "good news is bad news" regime — the full regime table is Chapter 7's subject).

What the headline does not tell you

The preparation work is mostly learning what the headline is not:

Not the Fed's number. The Fed reads the full report — and, more importantly, a moving average of reports, plus the claims series, the JOLTS data, and all the survey evidence. No single print moves a committee that meets on six-week cycles; the market knows this, which is why the pattern of prints across months matters more than any single one, and why the report's market impact is regime-dependent in the first place.

Not the quality of the jobs. The headline counts positions; it says nothing about full-time versus part-time mix (a household-survey question), industry composition, or the temp-help and hours data that desks read as leading indicators of labor demand.

Not the inflation signal. The wage-and-hours layer (Chapter 5) is where the print's inflation information lives; a huge payrolls print with flat hours and soft wages is a different Fed input than the same print with overtime surging.

Not final. Every print will be revised twice more, then benchmarked annually (Chapter 6). At the extreme, history has seen initial prints in the hundreds of thousands revised away, and turning points where the initial print pointed one direction and the revised record another. The market prices the print as if final because it must; the trader who knows it is provisional reads follow-through differently.

The headline in the reaction hierarchy

At 8:30, the market's machines price the headline first — within a second — and the professional desks grade it, in the first minutes, against a settled hierarchy that the following chapters build:

1. The headline versus consensus (the gap that moves the market) 2. The revisions to prior months (the trend adjustment — a strong print with downward revisions is discounted) 3. The unemployment rate and participation (the household layer — can confirm or contradict) 4. Wages and hours (the inflation layer — can invert the equity reaction) 5. The composition detail (industry, temp help, diffusion — the qualitative read)

A prepared trader holds this hierarchy written on the jobs-week sheet (Chapter 9) before 8:29; an unprepared one receives it, in fragments and out of order, from headlines for the next hour. The difference between the two traders at 8:31 is not information — it is hierarchy.

The next chapter takes the hierarchy's third item, the number that moves markets second only to the headline and contradicts it most often: the unemployment rate.

The birth-death model, honestly stated

Part of the headline is not sampled at all. The birth-death model is the Bureau's statistical adjustment for employer firms too new to be in the survey's frame and too dead to have been counted out of it — net business formation, estimated from history and applied to each month's raw count. In ordinary times it contributes a modest, stable drift, and the market ignores it correctly. In turning points it is the release's most controversial ingredient: the model extrapolates formation from the recent past, so it runs hot at the top of a cycle, contributing phantom jobs just as real formation stalls, and ran famously cold during the pandemic's collapse before the benchmark caught the rebound the model missed. Its errors are always in the direction of the previous regime — which is precisely why the benchmark revisions of Chapter 6 and the tax-record truth serum exist.

The headline trader does not need to model birth-death; the professional stance is to know its sign and its season. The Bureau updates the adjustment factors annually, and the factor tables are published with the release, so the sheet can carry the year's monthly pattern in advance — some months' headlines carry structurally larger adjustments than others, and a beat or miss delivered inside a large birth-death month deserves a credibility discount that the market rarely applies in the first minutes. Treat the model as the headline's least-known contributor: unmodellable in real time, wrong at every turning point, and corrected on a schedule. The traders who remember that in the turning-point months — and forget it in all the others — are pricing the headline one layer deeper than the room.

Reading the detail tables: the ten-minute upgrade

The headline and its consensus are in every summary; the report's detail tables are free and read by a small minority. For the prepared trader, three tables repay ten minutes monthly:

The industry breakdown. Where the jobs actually came from — goods versus services, and the specific categories. Its uses: composition-checking the wage read (Chapter 5), grading breadth (the diffusion logic in miniature), and tracking the cycle's sector rotation (government, healthcare, leisure, professional services each carry different cyclical information). A print concentrated in one or two categories is fragile; a print spread across the table is robust — and the same +150,000 is a different print in each case.

The diffusion index. The share of industries adding jobs, published in the release's summary section — the breadth measure in its official form. Its levels and changes lead the payroll trend at turning points, and one glance grades the headline's quality in a way the headline number itself cannot.

The household-survey detail. The tables behind the unemployment rate: labor force level and participation, the reasons for nonparticipation, duration of unemployment, and the part-time-for-economic-reasons count. This is where the four-case decoder of Chapter 4 gets its evidence, and where U-6's components live.

The ten-minute habit — headline read, then the three tables — is the difference between consuming the report and reading it. Within a year, the habit builds a private dataset: your own sense of which industries are rotating, what breadth looks like in this cycle, and how the household detail has been behaving — context no summary provides and no headline can price.

Chapter 4: The Unemployment Rate

The unemployment rate is the report's second headline: the one number from the household survey that gets bolded next to payrolls in every summary, quoted by politicians, and watched by the Fed. It is also the report's most misread statistic — a number that can fall for terrible reasons, rise for good ones, and contradict the payrolls print with perfect statistical innocence. This chapter is the field guide.

What the rate actually measures

The unemployment rate is unemployed people as a share of the labor force. Both halves are household-survey classifications, made at the person level, in the survey week:

- Employed — worked at all during the reference week (even an hour), or had a job but didn't work (vacation, illness) - Unemployed — did not work, and actively searched for work in the reference period - Not in the labor force — did not work and did not search: retirees, students, caregivers, and the discouraged

The rate's denominator — the labor force — is what gives the number its deceptive

depth. The participation rate (labor force as a share of the civilian adult population) determines how the same number of unemployed people translates into a rate: unemployment of 7 million in a labor force of 160 million is 4.4%; the same 7 million in a 150 million labor force is 4.7%. The rate moves for denominator reasons as often as for numerator reasons, and the denominator moves with demographics, retirement waves, and search behavior — none of which appear in the bolded number.

The rate's four ways of moving

Every possible unemployment-rate move decomposes into one of four cases, and each has a different economic meaning:

1. Falls because unemployment fell (people found jobs). The healthy decline: household employment rose faster than the labor force. Confirming, bullish for the growth read.

2. Falls because the labor force shrank (people left). The unhealthy decline: unemployed people stopped searching and exited the count. The rate improves while the labor market worsens. This is the "discouraged worker" dynamic, and its signature in the detail is a falling participation rate alongside the falling unemployment rate. The market's desks check participation before celebrating any unemployment-rate decline — a rate improvement with participation falling is read as weakness wearing a good headline.

3. Rises because unemployment rose (jobs lost or search failing). The honest deterioration — the case that dominates recession headlines.

4. Rises because the labor force grew (more people looking). The good rise: people re-entered the labor force — perhaps drawn by improving prospects — and are counted as unemployed while they search. Its signature: rising participation alongside rising unemployment, and it can accompany a strong payrolls print. Markets have traded this case as bad news anyway (the bolded number rose), then reversed as the composition was read — a recurring whipsaw pattern for headline-only traders.

The practical rule for the prepared trader: the unemployment rate is a ratio, and the ratio's meaning lives in which side moved. Participation is the tell. Rate down + participation down = weak; rate up + participation up = strong-ish; rate down + participation up = the strongest combination; rate up + participation down = the weakest. Write the four cases on the sheet; grade the print into one of them in the first minute.

U-3 and its alternatives

The official headline rate is technically U-3, one of a series of alternative measures the Bureau publishes. The most-watched alternative, U-6, is the broadest: it adds discouraged and marginally-attached workers (people who want work and have searched recently but not in the reference period) and those working part-time for economic reasons (who want full-time work). U-6 runs meaningfully above U-3 and

captures the slack that the headline misses — it is the number to check when the headline rate looks too clean, and the Fed's own commentary has, in various eras, leaned on underemployment language that tracks U-6's content.

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.