SEPTEMBER 2026

Agincourt Capital Management, LLC

WHEN THE DATA ITSELF BECOMES THE RISK

Why data quality matters

The reliability of government economic data was once an unquestioned bedrock for investment decisions. Increasingly, however, the reliability of that data has become a front-line risk for bond investors. At Agincourt, our assessment of the economic environment provides important context for our fundamental, value-based investment process. Relative value and security selection remain the primary drivers of our risk-taking, but understanding the backdrop for growth, inflation, and monetary policy helps us evaluate those opportunities and the risks surrounding them. Government statistics are an important part of that assessment: the monthly jobs report, weekly jobless claims, CPI and PCE inflation, GDP, and a long tail of ancillary releases. We recognize that “how good is the government’s data” ranks somewhere below watching paint dry on the list of scintillating *Investment Update* topics, but paint doesn’t move bond yields, so we’ll ask for your patience this month. Even when the data is timely and reliable, interpreting the economic backdrop alongside valuations and market fundamentals is difficult enough. When the data turns noisy, delayed, or subject to outsized revisions, it gets harder still. This isn’t a theoretical worry. Over the past several years, the official data that the Federal Reserve, and every bond investor, relies on has become less reliable. It’s worth walking through what changed, and what the Fed and the Bureau of Labor Statistics (BLS), now under new leadership, are doing about it.

A statistical apparatus under strain

The BLS’s monthly employment report has long been the single most market-moving release on the economic calendar, for good reason: it’s the timeliest broad read available on the labor market. But timeliness has always come at a cost in precision, and that cost has been rising. A decade ago, roughly 60% of the 120,000 sampled businesses in the establishment survey reported in time for the initial release. That response rate is now well under half, which widens the eventual gap between the first published number and the fuller figures that follow months later.

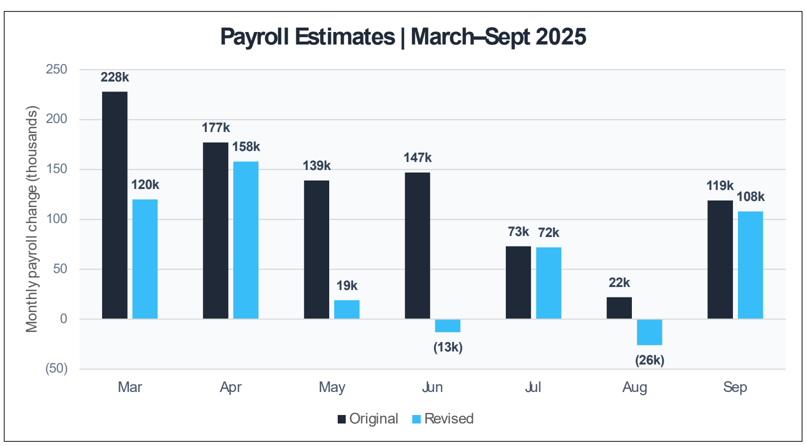

That gap became impossible to ignore in the summer of 2025. The July report carried unusually large downward revisions to the prior two months: May’s job gain was cut by roughly 125,000 and June’s by a similar amount - revisions large enough to flip the narrative on the labor market from resilient to weakening.
That episode turned political. The administration characterized the revisions as evidence of manipulation, and the BLS commissioner was removed from her post shortly after. A subsequent preliminary benchmark revision, tying the survey-based payroll count to slower but more comprehensive state unemployment tax records, suggested job growth in the twelve months through March 2025 had been overstated by roughly 900,000 positions. This was one of the larger benchmark corrections in decades. Independent analyses from the Peterson Institute, the San Francisco Fed, and others have generally concluded that the revisions reflect the ordinary mechanics of a survey-based system under stress from declining participation, not deliberate manipulation. That reassurance doesn’t answer the question that actually matters to us, though: can we trust the initial read on any given month, or do we now need to treat it as a rough first cut?

Two federal government shutdowns within the past year compounded the problem, disrupting the data calendar in ways rarely seen before. The shutdown running from October into November 2025 lasted long enough that no standalone October employment report was ever published. The October unemployment rate will simply never be known, since the household survey couldn’t be collected during its reference week and the BLS says the gap can’t be reconstructed after the fact. A second, shorter shutdown spanning late January into February 2026 delayed several more releases. According to the Government Accountability Office, a report published this June concluded the agency faces genuine and growing risk to data quality, though it also found that data users still generally regard the jobs report as useful.

Staffing, budget, and leadership turnover

The past year’s disruptions didn’t come out of nowhere. They landed on an agency that had already been thinning for over a decade. By outside estimates, the BLS’s budget and headcount have fallen by roughly a fifth since 2009, even as the economy it measures has grown substantially. That erosion left little cushion when a government-wide hiring freeze took hold in 2025. The BLS suspended or scaled back CPI collection, stopping it entirely in Lincoln, Nebraska; Provo, Utah; and Buffalo, New York, and cutting the sample by roughly 15% across dozens of other metro areas, simply because it lacked the staff to visit the usual number of stores and rental units. Producer Price Index detail shrank in parallel, with several hundred wholesale price categories dropped from publication. The staffing shortfall has pushed the BLS to fill a growing share of CPI with estimated prices rather than ones it actually observed that month. When a price collector can't get a real reading — a store didn't respond, or there's no staff to check — the BLS fills the gap by borrowing the price change from a similar item elsewhere. By August

2025, this kind of substitute pricing accounted for 36% of the index, up from just 9% six months earlier.

Leadership was equally unsettled. Commissioner Erika McEntarfer, confirmed 86-8 in early 2024, was dismissed in August 2025 within hours of the weak jobs report described above, amid claims of political manipulation. Her first proposed successor, a think-tank economist with no government statistical experience, drew bipartisan criticism over his qualifications and was withdrawn weeks later; the agency then ran in an acting capacity for the better part of a year. A second nominee fared better: career BLS economist Brett Matsumoto made data integrity the center of his June confirmation hearing, and the Senate confirmed him in early August, 51-47, as the BLS's 17th commissioner. That same day, another employment report carried large downward revisions to prior months, a reminder that a credible commissioner doesn't by itself eliminate the underlying data-quality problem, even if it helps restore confidence that the numbers aren't managed for political effect.

The Fed's response: five task forces, one squarely about data quality

New Federal Reserve Chairman Kevin Warsh, who succeeded Jerome Powell in May, has moved quickly on this front. In July he announced five outside task forces reviewing core aspects of monetary policy: communications, balance-sheet policy, inflation frameworks, productivity and jobs (with a notable focus on artificial intelligence), and, most relevant here, the quality of economic data itself.

The data task force is led by Harvard economist Raj Chetty, a leading voice on using large administrative datasets in place of traditional surveys; former Walmart CEO Doug McMillon, who brings a private-sector view of tracking real-time activity at a scale government statistics can't match; and University of Chicago economist Kevin Murphy. Their mandate is to examine whether the Fed's data inputs can be modernized against the kind of survey-response decline that produced last year's benchmark shock. With a permanent, Senate-confirmed commissioner now in place at the BLS, the task force finally has a stable counterpart at the agency to work alongside.

We think this is the right instinct. A Fed that actively rethinks its data infrastructure, rather than treating survey-based releases as unquestioned ground truth, is more likely to weigh incoming data appropriately and communicate its uncertainty honestly. But task forces take time, and an agency still rebuilding a depleted staff isn't going to modernize overnight. In the meantime, the practical burden of noisier official statistics falls on investors today, not on whatever improved framework may exist a year or two from now.

Our response: looking beyond the BLS

This is why we've spent time expanding our use of private and high-frequency data sources that don't depend on a monthly household or establishment survey. Many of these sources are broadly available to investors, and none is a perfect substitute for official statistics. Most capture a narrower slice of the economy, and some carry their own biases. Used alongside official releases, though, they provide valuable additional perspective and can help produce a more complete and timely read on the economic backdrop.

On the labor market, we track private payroll data drawn from actual payroll records for tens of millions of workers, available before the BLS's release. On inflation, we monitor high-frequency, transaction-level price indices built from tens of millions of online and point-of-sale price observations daily, updating continuously rather than monthly. Bank of America credit card spending data gives us a real-time weekly read on consumer demand that predates the retail sales report by weeks. We also track corporate earnings trends across market capitalizations and geographies: aggregate revenue and margin trends across thousands of public companies amount to a continuous, crowd-sourced read on economic activity, one that needs no government survey and is hard to manipulate or delay. Forward-looking management guidance and the direction of analyst earnings revisions add a further layer. Companies raising or cutting outlooks in real time often signal turning points before they show up in official statistics. And in credit markets, we watch corporate cash-flow data, commercial real estate delinquency trends, and asset-backed performance across consumer segments to gauge the health of the broader economy.

This framework was tested just this month. The August employment report, released in early September, showed payrolls up 162,000 — more than triple the consensus forecast, and a sharp reversal from a nearly stagnant early summer. Weekly ADP private payroll and spending data had been pointing to firmer conditions through August, providing useful context for interpreting the stronger official report and suggesting that underlying conditions may have been somewhat firmer than the earlier data implied. The swing from a "cooling, rate-cut-friendly" narrative to one now debating a possible rate hike happened almost entirely on the back of two data points a month apart. That kind of whipsaw is exactly why triangulating across sources, rather than reacting too strongly to any single release, is the more prudent approach.

Still anchored to valuations

None of this changes our fundamental investment philosophy: valuations are the anchor to everything we do. But it does reinforce the importance of considering any single official data release in the context of a broader set of information. In an environment where an initial economic signal can be meaningfully revised only a few months later, patience and corroboration become increasingly important. The Fed's decision to formally study its own data infrastructure, under a chairman willing to bring in outside academic and private-sector expertise, tells us we're not alone in recognizing this problem. The confirmation of a permanent career-BLS commissioner is a welcome step toward stabilizing the agency itself. We expect the task force's recommendations, whenever they arrive, to influence how the Fed reacts to incoming data and the broader market's confidence in the numbers that move it. Until then, our approach remains one of triangulation. Official statistics are still the baseline and the benchmark against which everything else is measured, but they are no longer the only lens through which we view the economy, and we don't think they should be for any other market participant.