Turning a Telescope on Quant: Part A - Seeing the Invisible

Return Stacking, Distant Galaxies, and a Little Bit of Magic

Take a magical trip into the Universe, both financial and cosmological, as we explore how stacking can reveal real but faint signals that no single observation can see. Along the way, we'll pass distant galaxies, bridge astronomical magnitudes to Sharpe ratios, and find out why the real magic in trend following is not prediction, but the disciplined accumulation of weak evidence.

Stacking: Seeing the Invisible

'Return stacking' has become a genuine discussion topic, with Corey Hoffstein of Newfound Research and Rodrigo Gordillo of Resolve Asset Management making persuasive cases for capital-efficient portfolio construction: using futures and overlays to put the same dollar of capital to work in more than one return stream.

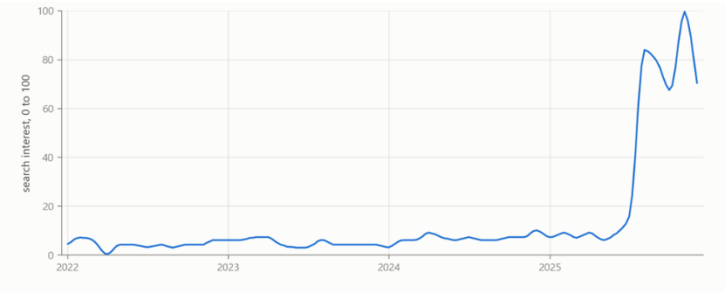

Figure 1.Google Trends for 'Returns Stacking'. Source: Google July 2026.

That is the newer, allocator-level version of the story. Inside a CTA, though, stacking is older and more fundamental. It is the mechanism by which faint market-level signals become tradable: repeated through time, diversified across markets, and combined across genuinely distinct return streams.

In other words, the quiet magic is inside the CTA itself: combining weak, noisy observations without fooling ourselves, then finding out how much of the edge survives the journey from model to market. Stacking creates *potential* edge. Whether it survives depends on the friction it must traverse: turnover, liquidity, implementation drag, and capacity.

Astronomy offers both lessons: true structure can be made visible by stacking weak observations, but low signal-to-noise also gives imagination room to mislead.

In Space, No One Can Hear You ~~Scream~~ Stack

Astronomy has long embraced stacking: combining noisy datasets so uncorrelated *noise* cancels and weak *signal* is boosted. Astronomers stack N multiple images of the same patch of sky to produce a sharper, more sensitive image: random atmospheric distortions (for land-based telescopes) cancel, and CCD noise[1] falls as $1/\sqrt{N}$.

The famous Hubble Deep Field image, containing thousands of distant galaxies in a patch of sky no larger than your thumb held at arm's length, was the result of hundreds of hours-long images taken over the course of ten days. To the naked eye that patch of sky looked empty. Through stacking, it became a window to the early Universe.

Figure 2. The Hubble Deep Field (1995). http://hubblesite.org/image/388/gallery.

When I was still an astrophysicist in the early 2000s, there came an even more magical-seeming development. Some of the most luminous objects in the Universe are the radio-loud, supermassive black hole regions at the centers of distant galaxies. The radio emission

[1] Charge-coupled devices suffer random noise through thermal electrons generated inside their silicon or readout noise when the signal is converted.

is easy to detect, but the host galaxies themselves are too faint to see. Their few arriving photons are lost in CCD noise. Whispering beside a waterfall.

By taking optical images centered on hundreds or thousands of these radio point sources and stacking them together, their common property is boosted relative to the noise. You collect sufficient photons from the host galaxies to measure the average properties of those hidden galaxies, for example whether they are dominated by old red stellar populations or young blue star formation, even though you cannot see them individually. This always seemed like a quiet bit of magic.

When the Hubble Telescope first launched in April 1990, it was discovered it had a severe spherical aberration: its primary mirror was ground 2.2 microns too flat at the edges, causing blurred images. When noise itself is correlated, stacking images does not reduce the noise, which was why in 1993 we sent some brave astronauts to fix it!

Can we bridge the interstellar gap from faint galaxies to the weakness of quant-finance signals? The analogy becomes useful only when we ask how faint those financial signals really are.

Eye Spy with My Little Eye

The signal-to-noise comparison is where the analogy becomes useful, but also where it needs care. In astronomy, faintness is measured in *magnitudes*, an inverse logarithmic scale, and the human eye can see objects as faint as magnitude +6 or 7. The most recent Hubble Ultra Deep Field goes almost a billion times fainter, to magnitude +29 or so.

If I've got my strangely underused astronomy-to-finance conversions correct[2], that is equivalent to stacking individual Sharpe ratios of less than 0.01. Coincidentally, CTA-land is in fact not that far off the mark here: A market-level *daily* Sharpe ratio is at a similar level.

Through stacking 252 or so daily observations across a year, that becomes an annual market Sharpe ratio of perhaps 0.1 to 0.2. Then, by stacking 100-plus markets within a portfolio, you may end up with a portfolio Sharpe ratio around 1. And just like stars, some markets are "brighter" than others.

The important point is that the market count is not the magic number. The relevant number is the effective number of independent observations: distinct physical risk factors, different supply-demand constraints, and lower shared beta to mainstream trend. Just like the broken

[2] Five magnitudes correspond to a factor of 100 in flux; equivalently, one magnitude is about a factor of 2.512.

Hubble telescope, a hundred versions of the same common factor do not reveal a hidden galaxy; they just give the noise better lighting.

Stack 'Em High and Watch 'Em Fly

Inside a trend CTA, then, the first stack is internal: not bonds plus commodities, or one fund overlaid on another, but the accumulation of weak directional evidence across time, markets and genuinely different risk factors.

Allocator-level return stacking is a natural extension of the same idea. My colleagues at GQ recently highlighted this in their 'Eras Tour' example: rather than investing $100 in long-only bonds, you could invest $50 in a bond-futures strategy, lever it to $100 notional, and use the other $50 to potentially generate another $100 of commodity-trend exposure.

From $100 of capital, you've created $200 of complementary return streams. That is strategy-level stacking rather than CTA-internal signal stacking, but the family resemblance is clear: more independent bets, less shared noise and disciplined scaling of the diversified return stream.

This is not sorcery, but the outcome can feel magical: a more all-weather defensive portfolio with a potentially better chance of working across inflation regimes. Bonds drag in rising-inflation eras, precisely when commodity trend can excel. Stacked together, they form a defensive engine that runs across all four inflation regimes.

Leverage gives us the tool to arrive at a clearer image, a higher Sharpe portfolio, while also delivering capital efficiency to the investor. But beware! leverage is only benign if the stacked noise is truly uncorrelated. If the "apparent" diversification is actually a common error wearing different market labels, leverage does not create capital efficiency, it turns one mistake into a portfolio.

The Instrument Matters

The implementation analogy is also more literal than it first appears. In astronomy, photons must pass through atmosphere, optics, detector noise, calibration error and sky background before they become evidence. The raison d'être of the Hubble telescope was to cut out much of that noise. But while astronomy stacks faint signals from a physical universe whose laws do not adapt to being observed, finance stacks faint signals from an adaptive social system: agents learn, crowding changes payoffs, transaction costs shift, and capacity constraints bite. For us, noise is built-in, which makes implementation, robustness and restraint part of the signal problem rather than afterthoughts.

In finance, the signal has to pass through turnover, liquidity, bid/ask spread, market impact, execution quality and capacity. Stacking may reveal the object, but the instrument determines how much of it survives measurement.

This is why implementation is not a footnote to the stacking story. It is part of the signal-to-noise calculation. A signal that looks attractive before costs may be untradeable after turnover. A diversified portfolio may be constrained by liquidity. A theoretically elegant stack may become fragile if capacity pressure forces execution into the least forgiving parts of the market.

The practical question, then, is not merely whether stacking creates edge. It is whether the edge survives measurement, execution and scale. Robust design lives where stacking amplifies signal faster than implementation dissipates it.

Conclusion: Making the Invisible Visible

The magic of stacking is not that it manufactures return from nowhere. It is that, under the right conditions, it allows faint but real structure to become visible. In astronomy, the result is a galaxy where the eye saw darkness. In trend following, it is a portfolio-level edge where individual daily signals look almost indistinguishable from noise.

But the conditions matter. The observations must contain coherent signals. The errors must not all be the same error. The leverage must scale diversification rather than disguise concentration. And the instrument, whether telescope or trading process, must be good enough to preserve what the stack reveals. The magic isn't in seeing things that aren't there. The magic is in making the invisible visible.

In Part B, we will turn to the darker side of the same problem: when the data are faint enough, stacking can reveal hidden galaxies, but imagination can also draw canals on Mars.

Important Disclosures

This material is for general informational and educational purposes only and is not investment, legal, tax, accounting, or other professional advice; an offer or solicitation to buy or sell any security, commodity interest, derivative, advisory service, or private fund interest; or a recommendation regarding any instrument, strategy, or product. It does not consider any reader's objectives, financial circumstances, risk tolerance, restrictions, tax status, or other needs. Investment decisions should be made only after consulting appropriate advisers and reviewing applicable offering, advisory, account, and risk materials.

Trading commodity interests, including futures, options, swaps, and other derivatives, is speculative and involves a high degree of risk. These instruments may be volatile, leveraged, or illiquid and may cause substantial losses, including losses exceeding the amount invested where applicable. Using derivatives to increase notional exposure or combine return streams can magnify losses, liquidity demands, financing costs, and adverse-correlation effects. Past performance is not necessarily indicative of future results.

Sharpe ratio is a risk-adjusted measure calculated by subtracting a risk-free rate from an investment's return and dividing by return volatility over the applicable period. A higher ratio generally indicates more return per unit of measured volatility. Results depend on the period, return frequency, risk-free rate, annualization, fees and expenses, and return distribution; comparisons may not be meaningful unless methodologies are consistent. Sharpe ratio does not capture all risks, including liquidity, leverage, concentration, tail, counterparty, and operational risk.

References to trend following, return stacking, thematic or factor exposures, signals, portfolio construction, Sharpe ratios, and related concepts describe general observations and research considerations only and, unless expressly stated, do not describe the actual or expected performance, risk, holdings, trading, or investment process of any Gresham product, fund, account, or strategy.

Any signals, factors, strategies, backtests, modeled or hypothetical results, academic studies, numerical examples, or other non-actual results are illustrative only and, unless expressly stated, do not represent actual Gresham trading. Such information has inherent limitations, may use hindsight or assumptions that prove inaccurate, and may omit material economic, market, and implementation effects, including fees, expenses, taxes, financing, transaction costs, slippage, liquidity, position limits, changing correlations, market impact, and operational constraints. No investor, account, fund, or strategy is represented as likely to achieve similar results.

Gresham, its affiliates, clients, funds, accounts, or personnel may hold or trade, directly or indirectly and long or short, instruments, sectors, themes, issuers, or strategies discussed herein. Positions and exposures may change without notice, and this material does not describe all such activity.

Information from academic research, market commentary, index providers, data vendors, public websites, and other third parties is believed reliable but has not been independently verified; no representation is made as to its accuracy, completeness, timeliness, or reliability. References and hyperlinks are for convenience and context, do not imply endorsement, adoption, or sponsorship, and Gresham is not responsible for third-party website content, availability, or privacy practices.

Forward-looking statements, identified by terms such as "may," "will," "should," "expect," "anticipate," "project," "estimate," "intend," "continue," or "believe," involve risks and uncertainties; actual events, results, or performance may differ materially. Nothing herein is a guarantee, promise, assurance, or representation as to the future.